

Dimensional Inconsistency Measures and Postulates in Spatio-Temporal Databases

John Grant

GRANT@CS.UMD.EDU

*Department of Computer Science and UMIACS
University of Maryland, College Park, USA*

Maria Vanina Martinez

MVMARTINEZ@DC.UBA.AR

*Institute for Research in Computer Science (UBA-CONICET)
Department of Computer Science, FCEyN
University of Buenos Aires, Argentina*

Cristian Molinaro

CMOLINARO@DIMES.UNICAL.IT

Francesco Parisi

FPARISI@DIMES.UNICAL.IT

*Department of Informatics, Modeling,
Electronics and System Engineering
University of Calabria, Italy*

Abstract

The problem of managing spatio-temporal data arises in many applications, such as location-based services, environmental monitoring, geographic information systems, and many others. Often spatio-temporal data arising from such applications turn out to be inconsistent, i.e., representing an impossible situation in the real world. Though several inconsistency measures have been proposed to quantify in a principled way inconsistency in propositional knowledge bases, little effort has been done so far on inconsistency measures tailored for the spatio-temporal setting.

In this paper, we define and investigate new measures that are particularly suitable for dealing with inconsistent spatio-temporal information, because they explicitly take into account the spatial and temporal dimensions, as well as the dimension concerning the identifiers of the monitored objects. Specifically, we first define natural measures that look at individual dimensions (time, space, and objects), and then propose measures based on the notion of a *repair*. We then analyze their behavior w.r.t. common postulates defined for classical propositional knowledge bases, and find that the latter are not suitable for spatio-temporal databases, in that the proposed inconsistency measures do not often satisfy them. In light of this, we argue that also postulates should explicitly take into account the spatial, temporal, and object dimensions and thus define “dimension-aware” counterparts of common postulates, which are indeed often satisfied by the new inconsistency measures. Finally, we study the complexity of the proposed inconsistency measures.

1. Introduction

Managing spatio-temporal data is an important problem in many applications, ranging from location-based services and geographic information systems to air traffic control and environmental monitoring. For this reason, the representation and processing of spatio-temporal data has attracted much attention by the AI community (Cohn & Hazarika, 2001; Gabelaia, Kontchakov, Kurucz, Wolter, & Zakharyashev, 2005; Yaman, Nau, & Subrahmanian, 2004, 2005; Knapp, Merz, Wirsing, & Zappe, 2006; Bona, Grant, Hunter,

& Konieczny, 2019) as well as by the database one (Pelaniš, Saltenis, & Jensen, 2006; Tao, Papadias, & Sun, 2003; Kollios, Gunopulos, & Tsotras, 1999; Agarwal, Arge, & Erickson, 2003; Pfoser, Jensen, & Theodoridis, 2000; Hadjieleftheriou, Kollios, Tsotras, & Gunopulos, 2002).

In this paper, we focus on databases representing atomic statements of the form “object id is/was/will be inside region r at time t ”, denoted as atoms of the form (id, r, t) . This allows the representation of information concerning moving objects in many application domains. For instance, a cell phone provider is interested in knowing which cell phones will be in the range of some towers at a given time (Bayir, Demirbas, & Eagle, 2010). A transportation company is interested in predicting the vehicles that will be on a given road at a given time in order to avoid congestion (Karbassi & Barth, 2003). A retailer is interested in knowing the positions of shoppers moving in a shopping mall in order to offer customized coupons on discounts (Kurkovsky & Harihar, 2006). These are just a few examples of a common issue.

Most previous work on spatio-temporal data assumes that a *consistent* version of the data is available before processing queries or performing updates. However, in real life scenarios, this assumption is rather simplistic in general, but especially when the information changes dynamically and in a distributed manner. Inconsistency in those settings can appear for various reasons, such as sensing errors that can affect the estimation and generation of spatio-temporal information. For instance, in smart video surveillance (Sedky, Moniri, & Chibelushi, 2005), an identification system may fail to assign the same identifier to the same object monitored by cameras having different views (An, Chen, Kafai, Yang, & Bhanu, 2013; Zhang, Kalashnikov, Mehrotra, & Vaisenberg, 2014), thus generating wrong object identifiers that do not correctly model the monitored situation. Likewise, different objects may be assigned the same identifier after exiting and re-entering the monitored space (Bedagkar-Gala & Shah, 2014), yielding an inconsistent representation of the situation being monitored. Other errors can be due to imprecise estimation of regions occupied by objects, and inaccurate time-stamps associated with sensor readings, because of the inherent imprecision of the systems originating spatio-temporal data (Pfoser & Jensen, 1999).

Moreover, as data are often collected from heterogeneous sources, such as on-board GPS devices and roadside sensors (Zhang & Trajcevski, 2014), the same monitored situation (an object being in a given place at a given time) can result in different sets of spatio-temporal data corresponding to the different sources. Thus, data integrated from different sources are likely to generate inconsistent spatio-temporal databases, where an inconsistency is due to the presence of data representing conflicting information on the monitored scenario. For instance, an on-board GPS device could detect the car in which it is installed quite far away from the place where it actually is at the given time (e.g., on a route parallel to that being used), while a roadside sensor could detect (the license plate number of) that car on a different street at the same time. The information provided by these two sources, which are independent of each other, may be inconsistent since they entail the presence of the same car in different places at the same time. For that matter, each of the two sources may provide wrong information: the detection region of the GPS device may be inaccurate and/or the license plate number detected by the roadside sensor may be wrong.

Considering the specific aspects and dimensions that spatio-temporal information contains, as indicated in the examples, we observe that defining, identifying, measuring, and

deciding how to deal with such data appear to be more complex than dealing with theories given in propositional logic, where there has been an extensive body of work on inconsistency measures. In this paper, we study the use of inconsistency measures in spatio-temporal databases. An inconsistency measure is a mathematical function used to determine the quantity of inconsistency in a database. The measure zero implies that no inconsistency is present, while the larger this number, the “dirtier” is the information.

Contributions. This paper substantially extends our preliminary attempt to measure inconsistency in the spatio-temporal context (Grant, Martinez, Molinaro, & Parisi, 2018a). In particular, we make the following contributions.

- We propose new inconsistency measures for spatio-temporal databases that, differently from classical ones, explicitly take into account the spatial, temporal, and object information characterizing such data. Specifically, we first define natural measures that look at individual dimensions (time, space, and objects), and then propose measures based on the notion of a *repair*. The basic idea of the latter is to measure inconsistency as the cost of restoring consistency (in a “minimal” way).
- Rationality postulates are properties defined for propositional inconsistency measures in order to describe general desirable behaviors of the measures, though it is not universally agreed upon which postulates should be satisfied by an inconsistency measure to be considered as a “good” one (Besnard, 2014). In fact, as we discuss later, not all the postulates are jointly satisfiable, as there may exist properties (e.g., MI-Normalization and Attenuation) expressing conflicting desiderata. Nevertheless, analyzing the compliance of inconsistency measures w.r.t. postulates is an important way to evaluate the quality of an inconsistency measure and for this reason it is one of the main problems investigated in the area of inconsistency measurement, even if the set of desirable postulates is not universally accepted (Thimm, 2017, 2018). Given the importance of postulate satisfaction, we analyze which of eight applicable postulates are satisfied by the new inconsistency measures. It turns out that the classical postulates are not satisfied in many cases. The reason is that the latter have been devised for propositional knowledge bases and thus neglect the spatial, temporal, and object information altogether.
- In light of the aforementioned issue, we introduce “dimensional” postulates, that is, postulates inspired by the classical ones but explicitly taking into account the time, space, and object dimensions. The new measures satisfy them in most of the cases.
- Finally, we characterize the complexity of deciding whether the inconsistency value is less than or equal to a given threshold, for each of the inconsistency measures. While this decision problem is generally intractable for the classical measures for propositional knowledge bases (Thimm & Wallner, 2016; Thimm, 2018), we find that it becomes generally feasible, particularly under some reasonable assumptions, for the new measures for spatio-temporal databases.

We point out that this work provides insights into measuring inconsistency also in different (or more general) settings where “dimensions” (e.g., space, time, objects in our case)

characterize the data at hand: our results suggest that both inconsistency measures and postulates should treat such dimensions as first-class citizens in order to properly deal with inconsistent “dimensional” data.

Organization. Section 2 introduces the syntax and semantics of spatio-temporal (ST) databases. Section 3 defines the notion of an inconsistency measure along with classical postulates drawn from the propositional setting. Section 4 defines new inconsistency measures tailored for spatio-temporal data that take into account the three dimensions of objects, space, and time, and analyze their behavior w.r.t. (classical) postulates. Section 5 introduces inconsistency measures based on the notion of repair along the three dimensions. Specifically, four new inconsistency measures are proposed (including a measure based on deletions resulting in maximal consistent subsets) and their satisfaction of the classical rationality postulates given in Section 3 is investigated. It turns out that most of the classical postulates are not satisfied by these measures. So in Section 6 we modify some of the postulates and propose “dimension-aware” postulates specifically devised for spatio-temporal data and analyze in which cases they are satisfied by the new inconsistency measures. Finally, in Section 7 we explore the computational complexity of all the new measures introduced in Sections 4 and 5, and find that in general the new measures have lower complexity than the standard propositional measures. The paper is concluded and future work is outlined in Section 8.

2. Spatio-Temporal Databases

This section introduces the syntax and semantics of Spatio-Temporal (ST) Databases (DBs).

2.1 Syntax

Throughout the paper, we assume the existence of three finite sets: ID is the set of object ids, T is the set of integer time values, and $Space$ is the set of point locations; below we explain in further detail what these sets represent. We assume that an object can be in only one location at a time, but a single location may contain more than one object. This means that objects have no size and the space is discrete. However, the granularity of the space can be arbitrarily chosen to fit the application’s needs. As a matter of fact, real-world geographic information systems (GISs) generally assume that space is discretized into a grid and most geographic data structures (e.g., quadtrees, R-trees, etc.) supported by GISs make the same assumption (Samet, 2006).

A *region* r is a nonempty subset of $Space$, that is, $\emptyset \subset r \subseteq Space$.

Definition 1 (ST atom/database). *An ST atom is a tuple (id, r, t) , where $id \in ID$ is an object id, r is a region, and $t \in T$ is a time value. An ST database is a finite set of ST atoms.*

Intuitively, the ST atom (id, r, t) says that the location of object id belongs to region r at time t . Hence, ST atoms can represent information about the past and the present, such as that generated by techniques for interpreting RFID readings (Fazzinga, Flesca, Furfaro, & Parisi, 2014, 2016), but also information about the future, such as that derived from methods for predicting the destination of moving objects (Mittu & Ross, 2003; Hammel,

Rogers, & Yetso, 2003; Southey, Loh, & Wilkinson, 2007), or from querying predictive databases (Akdere, Cetintemel, Riondato, Upfal, & Zdonik, 2011; Agarwal, Chen, Ji Lin, Shanmugasundaram, & Vee, 2010; Parisi, Sliva, & Subrahmanian, 2013).

We point out that *Space* is defined as an arbitrary set of point locations; it can be as large as a user needs, and one can choose the granularity of space according to her needs. The same comments can be applied to time, that is, to *T*.

In the definition a region may be any non-empty set of points in *Space* (not necessarily connected). But in our examples, for convenience, we use a square grid for *Space*. In the grid, each location can be written as (x, y) , where x and y are integers and $0 \leq x, y \leq N$ for some integer N . Thus, *Space* contains $(N + 1)^2$ points. Also, we will use *isothetic* rectangles, by which we mean rectangles oriented in the standard way, that is, with horizontal and vertical sides. In the sequel, when we will refer to the *restriction of isothetic rectangular regions*, the considered ST databases are implicitly defined on a set *Space* corresponding to a square grid and the regions of their atoms are isothetic rectangles. An isothetic rectangular region can be specified by 2 points, its lower left corner, (x_1, y_1) , and its upper right corner, (x_2, y_2) , where x_1, y_1, x_2 , and y_2 are integers, $0 \leq x_1 \leq x_2 \leq N$, and $0 \leq y_1 \leq y_2 \leq N$. It is worth noting that we also consider degenerate isothetic rectangles consisting of single points, where $x_1 = x_2 = y_1 = y_2$, or sets of (axis-)aligned points, where either $x_1 = x_2$ or $y_1 = y_2$. Sometimes it will be useful to refer to the size of a region r , denoted $|r|$, by which we mean the number of points in r . Other (more complex) representations of spatial regions and time are clearly possible. Depending on the application domain there might be a requirement for more complex representations: this may include the modeling of vectors for spatial-temporal information. These knowledge representation issues are outside the scope of this paper and left to future investigation.

The following example illustrates a scenario where an ST database is used to model spatio-temporal data.

Example 1. *Consider a farm where the tasks of irrigation and fertilization are performed automatically by means of drones. The drones have limited memory and unilateral data transfer capabilities; therefore, they do not report about their location in time. In this example $ID = \{id_1, id_2, id_3\}$, $T = \{1, 2, 3\}$, and $Space = \{(x, y) \mid 0 \leq x \leq 16, 0 \leq y \leq 16\}$, where x and y are integers.*

As shown in Figure 1(b), the field of the farm is divided into 7 regions; each region is defined by the scanning range of a sensor in the terrain. The terrain sensors, besides monitoring soil and crop properties, such as humidity, etc., also capture and report the position of all the drones that fly over them. The traces of drones captured by terrain sensors are reported to a central system that stores the information.

This data is represented as tuples such the ones in Figure 1(a): every tuple consists of the id of a drone, the region of the field where some sensor identified the position of the drone, and the time at which the location of the drone was recorded. For instance, the tuple in the first row of the table in Figure 1(a), representing the ST atom $(id_1, a, 1)$, says that drone with id_1 was in region a at time 1. We use \mathcal{S}_F to denote the ST database consisting of the ST atoms represented by the tuples in Figure 1(a). Drone id_1 was recognized at time 1 in both regions a and b . In some cases, such as for id_1 at time 2, \mathcal{S}_F contains no information.




Figure 1: (a) (Tabular representation of the) ST database \mathcal{S}_F of Example 1; (b) Regions of the farm field for \mathcal{S}_F ; (c) Interpretation I of Example 2.

Given an ST database \mathcal{S} , an object id , and a time t , we use the notation $\mathcal{S}_{id,t}$ to refer to the set $\mathcal{S}_{id,t} = \{(id', r', t') \in \mathcal{S} \mid id' = id \wedge t' = t\}$, that is, the set of ST atoms in \mathcal{S} that refer to the specific object identifier id and time value t .

2.2 Semantics

The meaning of an ST database is given by the interpretations that satisfy it.

Definition 2 (ST interpretation). *An ST interpretation I is a function $I : ID \times T \rightarrow Space$.*

An interpretation specifies a trajectory for each $id \in ID$. That is, for each $id \in ID$, I says where in $Space$ object id was/is/will be at each time $t \in T$. In particular, this means that an object can be in only one location at a time.

Example 2. *An interpretation I for the ST database \mathcal{S}_F of Example 1 is shown in Figure 1(c).*

We now define satisfaction and ST models.

Definition 3 (Satisfaction and ST model). *Let $a = (id, r, t)$ be an ST atom and I an ST interpretation. We say that I satisfies a (denoted $I \models a$) iff $I(id, t) \in r$. I satisfies an ST database \mathcal{S} (denoted $I \models \mathcal{S}$) iff for all $a \in \mathcal{S}$, $I \models a$. We say that I is a model for a (resp., \mathcal{S}) iff I satisfies a (resp., \mathcal{S}).*

We use $\mathbf{M}(\mathcal{S})$ to denote the set of all models for an ST database \mathcal{S} .

Example 3. *In our example, interpretation I is a model for the ST atom $(id_3, c, 1)$, as id_3 at time value 1 is assigned to the spatial point $(9, 10)$, which is a point in region c . Reasoning analogously, it is easy to see that I is a model for all of the atoms in Figure 1(a). Hence, I is a model for \mathcal{S}_F .*

Example 4. *Let I_1 be the interpretation that is equal to I except that $I_1(id_3, 1) = (12, 6)$. It is easy to check that I_1 is not a model for the ST atom $(id_3, c, 1)$, because object id_3 is not assigned to a point in region c at time 1. Since $(id_3, c, 1)$ is in \mathcal{S}_F (see the last row of the table in Figure 1(a)), I_1 is not a model for \mathcal{S}_F .*

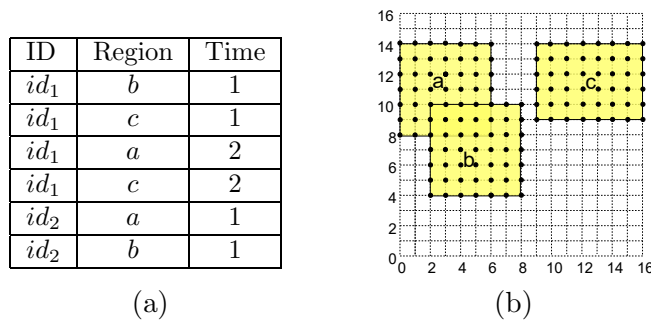


Figure 2: (a) (Tabular representation of the) ST database \mathcal{S}_e of Example 5; (b) Regions of the farm field for \mathcal{S}_e .

In the following, we will use the symbol M to refer to interpretations that are models, that is, elements in $\mathbf{M}(\mathcal{S})$.

Definition 4 (Consistency). *An ST database \mathcal{S} is consistent iff $\mathbf{M}(\mathcal{S}) \neq \emptyset$.*

Example 5. *The model of Example 3 shows that \mathcal{S}_F is consistent. Consider next the ST database \mathcal{S}_e in Figure 2, where a simpler scenario of the farm field is shown. In this case, $ID = \{id_1, id_2\}$ and $T = \{1, 2\}$, while Space is still the set $\{(x, y) \mid 0 \leq x \leq 16, 0 \leq y \leq 16\}$, where x and y are integers. Note that there is no model for \mathcal{S}_e , since it is impossible to define a function that assigns a spatial point to id_1 at time values 1 and 2, and can satisfy at the same time all the ST atoms in which it is involved (there is no overlap between regions b and c , and between a and c , respectively). Thus, \mathcal{S}_e is inconsistent.*

There is an important connection between the consistency of \mathcal{S} and its subsets $\mathcal{S}_{id,t}$, that is, \mathcal{S} is consistent iff for each $id \in ID$ and $t \in T$, $\mathcal{S}_{id,t}$ is consistent. Then, checking whether \mathcal{S} is consistent can be accomplished in polynomial time as follows: for each non-empty subset $\mathcal{S}_{id,t}$, we check whether the intersection of the regions of the ST atoms in $\mathcal{S}_{id,t}$ is non-empty. We also mention that consistency checking remains solvable in polynomial time in the more general setting of *probabilistic spatio-temporal knowledge bases* of Grant, Molinaro, and Parisi (2018b)—we refer the reader to (Grant et al., 2018b) for more details.

The set of all minimal (under set-inclusion) inconsistent subsets of an ST database \mathcal{S} is denoted as $\mathcal{MI}(\mathcal{S})$. For the ST database \mathcal{S}_e of Example 5,

$$\mathcal{MI}(\mathcal{S}_e) = \{\{(id_1, b, 1), (id_1, c, 1)\}, \{(id_1, a, 2), (id_1, c, 2)\}\}.$$

Notice that each set in $\mathcal{MI}(\mathcal{S})$ is included in exactly one $\mathcal{S}_{id,t}$.

As stated in the following proposition, when the regions are isothetic rectangles (with *Space* being a square grid) every minimal inconsistent set has size 2. We will use this result in Section 7 when characterizing the complexity of inconsistency measures.

Proposition 1. *Under the restriction of isothetic rectangular regions, every minimal inconsistent subset of every ST database \mathcal{S} has size 2.*

Proof. We have previously explained that the size of a minimal inconsistent subset of \mathcal{S} must be at least 2. Hence we need to show that it cannot be greater than 2. As mentioned




Figure 3: An ST database with convex regions and a minimal inconsistent subset of cardinality 3.

before, for inconsistency we need to deal only with ST atoms that have the same id, t pair, see (Grant et al., 2018a); hence we assume that such a pair is fixed and we deal only with the regions.

Let $|\mathcal{S}| = n + 1$ where $n \geq 2$. We write $\mathcal{S} = \{R_1, \dots, R_n, R_{n+1}\}$ where for all i , $1 \leq i \leq n + 1$, $R_i = ((x_i^\ell, y_i^\ell), (x_i^u, y_i^u))$, using the standard representation for isothetic rectangles. If \mathcal{S} is minimal inconsistent then $\mathcal{S}' = \{R_1, \dots, R_n\}$ is consistent and hence $\bigcap_{i=1}^n R_i \neq \emptyset$.

Find R_h such that $x_h^\ell \geq x_i^\ell$ for all i , $1 \leq i \leq n$. That is, going from left to right, no rectangle has its leftmost x value after R_h . Also, find R_k such that $x_k^u \leq x_i^u$ for all i , $1 \leq i \leq n$. That is, going from left to right, no rectangle has its rightmost x value before R_k . If $x_k^u < x_h^\ell$ then R_k is to the left of R_h , so $R_h \cap R_k = \emptyset$ which is not allowed by our assumption. Hence $x_h^\ell \leq x_k^u$. In an analogous way, find R_t such that $y_t^\ell \geq y_i^\ell$ for all i , $1 \leq i \leq n$ and R_z such that $y_z^u \leq y_i^u$ for all i , $1 \leq i \leq n$. Then $\bigcap_{i=1}^n R_i = ((x_h, y_t), (x_k, y_z))$. As \mathcal{S} is inconsistent, $R_{n+1} \cap ((x_h, y_t), (x_k, y_z)) = \emptyset$. There are 4 possibilities: R_{n+1} is either to the left, or to the right, or below, or above $\bigcap_{i=1}^n R_i$. Suppose that R_{n+1} is to the left of $\bigcap_{i=1}^n R_i$. Then $x_{n+1}^u < x_h^\ell$ and so R_{n+1} is also to the left of R_h and hence $R_{n+1} \cap R_h = \emptyset$. This violates the assumption that \mathcal{S} is minimal inconsistent. The other 3 cases are handled similarly. This shows that no set of ST atoms of size greater than 2 can be minimal inconsistent. \square

This result does not hold for arbitrary convex regions in the plane, as shown by the (inconsistent) ST database reported in Figure 3. Notice that the point $(2, 2)$ (resp., $(1, 1)$, $(3, 1)$) belongs to both r_1 and r_2 (resp., r_1 and r_3 , r_2 and r_3). Thus, every subset of the ST database having cardinality at most 2 is consistent. However, the three ST atoms together are inconsistent, and thus there is a minimal inconsistent subset of cardinality 3.

One may wonder for which integers it is possible to obtain a minimal inconsistent subset of that size. The next result answers this question.

Proposition 2. *For every integer $n \geq 2$ there is an ST database that has a minimal inconsistent subset of size n .*

Proof. For a given integer $n \geq 2$ we show how to construct n ST atoms that are minimally inconsistent. Let $|Space| \geq n$. Choose a region R such that $|R| = n$. Let $\mathcal{R}^{n-1} = \{r \subset R \text{ such that } |r| = n - 1\}$ and $\mathcal{S} = \{(id, r, t) \mid r \in \mathcal{R}^{n-1}\}$ for a fixed id, t pair. Then $|\mathcal{S}| = |\mathcal{R}^{n-1}| = n$ and $\bigcap_{r \in \mathcal{R}^{n-1}} r = \emptyset$. Also, for every point $p \in R$, p must belong to $n - 1$ (all except one) regions in \mathcal{R}^{n-1} . Hence the intersection of any $n - 1$ regions is not empty. Therefore $\{(id, r, t) \mid r \in \mathcal{R}^{n-1}\}$ is a minimally inconsistent set of size n . \square

ST databases can be viewed as a special case of Spatial PrObabilistic Temporal (SPOT) databases (Parker, Infantes, Grant, & Subrahmanian, 2009; Parisi, Parker, Grant, & Subrahmanian, 2010; Grant, Parisi, Parker, & Subrahmanian, 2010) which allow the representation of atomic statements of the form “object id is/was/will be inside region r at time t with probability in the interval $[\ell, u]$ ”. A SPOT statement is written as $(id, r, t, [\ell, u])$. Both ST databases and SPOT databases deal with uncertain information about the position of objects at a given time, but in different ways, as explained below. Uncertainty in ST databases is modeled by allowing the presence of inconsistent ST statements, such as (id, r, t) and (id, r', t) , with $r \cap r' = \emptyset$, which say that object id was detected in two different disjoint regions at the same time. In contrast, SPOT databases are assumed to be consistent, and uncertainty is modeled by statements having *compatible* probability intervals, which entails that the sum of the lower bound probabilities of being in different positions for the same object at the same time must not exceed 1. For instance, the SPOT statements $(id, r, t)[0.5, 1]$ and $(id, r', t)[0.5, 1]$, with $r \cap r' = \emptyset$, say that object id was detected in two different disjoint regions at time t with probability greater than or equal to 0.5 in both cases. Herein, the probability intervals are compatible, as it is possible to find a probability distribution on the positions of id at time t . This is not the case, for instance, for the SPOT database consisting of the statements $(id, r, t)[1, 1]$ and $(id, r', t)[1, 1]$, with $r \cap r' = \emptyset$, which is inconsistent (id cannot be in two disjoint regions with probability 1 in both cases).

Notably, the SPOT framework has been implemented and tested on real US Navy databases containing ship location data by Parker et al. (2009) and Parisi et al. (2010), where probability intervals were added to the original data in order to model spatio-temporal uncertainty by means of *consistent* SPOT databases, but the original data were in the form of an (inconsistent) ST database. Thus, ST databases can actually be used to store data coming from real-world applications. Examples of such applications include video surveillance (Sedky et al., 2005; Bedagkar-Gala & Shah, 2014) and tracking systems (Zhang & Trajcevski, 2014) discussed in the Introduction.

3. Inconsistency Measures and Rationality Postulates for ST Databases

In this section, we provide a general definition of inconsistency measure along with the standard postulates. (In the following two sections, we will propose concrete inconsistency measures treating the time, space, and object dimensions as first-class citizens).

3.1 The Definition of Inconsistency Measure

An inconsistency measure (IM) is a function that assigns a nonnegative real value or infinity to every database. Inconsistency measures can be classified in various ways and may satisfy certain properties. One distinction is between absolute measures that measure the total amount of inconsistency and relative measures that use a ratio to determine how big a portion of the database is inconsistent. We will not deal in this paper with relative inconsistency measures. For absolute measures the following definition is appropriate, where $\mathcal{R}_{\geq 0}^{\infty}$ is the set of nonnegative real numbers and the infinity symbol.

Definition 5 (Inconsistency measure). *Let \mathcal{L} be the set of all ST databases. An inconsistency measure $\mathcal{I} : \mathcal{L} \rightarrow \mathcal{R}_{\geq 0}^{\infty}$ is a function such that the following two properties hold:*

1. (Consistency) For any $\mathcal{S} \in \mathcal{L}$, $\mathcal{I}(\mathcal{S}) = 0$ iff \mathcal{S} is consistent.
2. (Monotony) For any $\mathcal{S}, \mathcal{S}' \in \mathcal{L}$, if $\mathcal{S} \subseteq \mathcal{S}'$, then $\mathcal{I}(\mathcal{S}) \leq \mathcal{I}(\mathcal{S}')$.

These two properties ensure that all and only consistent ST databases get a measure of 0 and that the measure is monotonic for subsets.

3.2 Some Rationality Postulates

We now report several properties that some inconsistency measures possess. These are just some of the properties that have been discussed as desirable for inconsistency measures. We point out that some of the postulates we omit are not useful for ST databases. For example, several deal with the case of logically equivalent formulas, but in our case all the formulas are ST atoms and no pair of distinct ST atoms are logically equivalent. That is why we deal only with the postulates given below. See (Thimm, 2018) for a thorough study of these properties, usually called rationality postulates. Each property has some intuitive rationale, even though some properties are not compatible with one another.

In Definition 5 we have already given two such properties that we believe all inconsistency measures must have. In Definition 6 below, we present a list of eight additional properties that have counterparts for propositional knowledge bases. As we are considering these properties directly for ST databases, we write the postulates as they apply to ST databases rather than propositional knowledge bases.

Before giving the definitions of the postulates, we introduce the notation and terminology for problematic and free ST atoms. For an ST database \mathcal{S} , the ST atoms that appear in some minimal inconsistent subset are called *problematic* ST atoms. The ST atoms that are not problematic are called *free*. Formally, we write $Problematic(\mathcal{S}) = \bigcup_{MIS \in \mathcal{MI}(\mathcal{S})} MIS$ and $Free(\mathcal{S}) = \mathcal{S} \setminus Problematic(\mathcal{S})$.

Definition 6. Let \mathcal{I} be an inconsistency measure, and $\mathcal{S}, \mathcal{S}'$ be two ST databases. We consider the following postulates:

1. (Free-Formula Independence) If $(id, r, t) \in Free(\mathcal{S})$ then $\mathcal{I}(\mathcal{S}) = \mathcal{I}(\mathcal{S} \setminus \{(id, r, t)\})$.
2. (Penalty) If $(id, r, t) \in Problematic(\mathcal{S})$ then $\mathcal{I}(\mathcal{S}) > \mathcal{I}(\mathcal{S} \setminus \{(id, r, t)\})$.
3. (Dominance) If (id, r, t) and (id, r', t) are ST atoms such that $r \subseteq r'$ then $\mathcal{I}(\mathcal{S} \cup \{(id, r, t)\}) \geq \mathcal{I}(\mathcal{S} \cup \{(id, r', t)\})$.
4. (Super-Additivity) If $\mathcal{S} \cap \mathcal{S}' = \emptyset$ then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$.
5. (Attenuation) If $MIS, MIS' \in \mathcal{MI}(\mathcal{S})$ and $|MIS| < |MIS'|$ then $\mathcal{I}(MIS) > \mathcal{I}(MIS')$.
6. (Equal Conflict) If $MIS, MIS' \in \mathcal{MI}(\mathcal{S})$ and $|MIS| = |MIS'|$ then $\mathcal{I}(MIS) = \mathcal{I}(MIS')$.
7. (MI-Normalization) If $MIS \in \mathcal{MI}(\mathcal{S})$ then $\mathcal{I}(MIS) = 1$.
8. (MI-Separability) If $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$ and $\mathcal{MI}(\mathcal{S}) \cap \mathcal{MI}(\mathcal{S}') = \emptyset$ then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$.

Free-Formula Independence states that free ST atoms do not contribute anything to the inconsistency measure. Penalty states that problematic ST atoms increase the inconsistency.

In the formulation for propositional knowledge bases, Dominance states that for two formulas α and β where $\alpha \models \beta$ and α is not a contradiction, adding α to a database increases its inconsistency at least as much as adding β . For ST databases the only way that an ST atom α can logically imply an ST atom β is if the *id* and *t* values are the same and the region of α is a subset of the region of β . We do not need to check for a contradiction as no ST atom by itself can be contradictory.

Super-Additivity deals with two ST databases that do not intersect, in which case the inconsistency of the union is at least as great as the sum of the inconsistency of the two ST databases. Attenuation and Equal Conflict refer to minimal inconsistent subsets. The former states that inconsistency is inversely related to size; while the latter requires minimal inconsistent subsets of the same size to have the same inconsistency. MI-Normalization states that a minimal inconsistent subset must have measure 1. MI-Separability means that for any two ST databases for which there is no common minimal inconsistent subset and for which the set of minimal inconsistent subsets of the union is obtained by taking the union of the minimal inconsistent sets of both, the inconsistency of the union is the sum of the inconsistencies of the components.

It is worth noting that, from the definitions of the postulates given above, analogously to what holds for their counterparts for propositional knowledge bases (Thimm, 2018), it is the case that:

- (i) MI-Normalization is incompatible with Attenuation (i.e., there is no inconsistency measure that satisfies both these postulates at the same time), and
- (ii) MI-Normalization implies Equal Conflict, i.e., if a measure satisfies MI-Normalization then it satisfies Equal Conflict.

As we investigate these properties for the new inconsistency measures, we will find that in most cases they are not satisfied. We do not consider this to mean that the postulates are not good ones or that our inconsistency measures are not good; it simply means that these specific postulates are not appropriate for the new measures. Intuitively, the reason is that the postulates do not take into account the time, space, and object dimensional information within ST atoms. This turns out to be too coarse-grained an approach for ST databases. To cope with this issue, we will define “dimensional” postulates (in Section 6), which are suitable for ST databases and the new inconsistency measures.

4. Dimensional Inconsistency Measures for ST Databases

Several inconsistency measures have been proposed in the literature for measuring inconsistency in propositional knowledge bases (Hunter & Konieczny, 2008; Grant & Hunter, 2013, 2011; Thimm, 2016; Döder, Raskovic, Markovic, & Ognjanovic, 2010). Therefore, one natural approach to measuring inconsistency in ST databases is to translate them into semantically equivalent classical propositional knowledge bases and then apply known measures. This approach has been considered by Grant et al. (2018a); however, its weakness is that it does not take into account the three dimensions of spatio-temporal data. The reason is that the translation introduces a propositional variable $x_{id,p,t}$ for each $id \in ID$, $p \in Space$, and $t \in T$, so that the three dimensions get condensed into a single propositional variable,

which makes it impossible to measure inconsistency along one (or more) dimensions with known inconsistency measures.

In this section, we propose several inconsistency measures that are not possible to define for propositional knowledge bases but are relevant for ST databases, as they explicitly take into account the dimensions characterizing such data. We use the fact that ST databases can be considered along three dimensions: objects, time, and space. This allows us to measure the inconsistency along one or a combination of dimensions.

Separating the dimensions of ST databases requires looking inside the formulas. Consider what such a step means for propositional knowledge bases. The formulas there contain propositions and logical connectives (as well as parentheses). As absolute inconsistency measures typically use only the problematic formulas, a natural way of measuring inconsistency is to count the number of distinct propositions in the problematic formulas. Let us call a proposition p *problematic* if p appears in a problematic formula. Then we can define $\mathcal{I}_P(K) = |\{p \mid p \text{ is a problematic proposition}\}|$. Actually, we did not find this definition in the literature on inconsistency measures. However, it is the absolute version of a relative inconsistency measure studied by Xiao and Ma (2012). So \mathcal{I}_P is our inspiration for measuring inconsistency along the three dimensions.

4.1 Measuring Inconsistency along the Object and Time Dimensions Individually

In some cases we may just be interested in how many objects or how many time values are involved in an inconsistency. Those are the two cases we deal with below.

4.1.1 MEASURING INCONSISTENCY ALONG THE OBJECT DIMENSION

An inconsistency measure based strictly on objects is:

$$\mathcal{I}_O(\mathcal{S}) = |\{id \in ID \mid (id, r, t) \in MIS \in \mathcal{MI}(\mathcal{S})\}|.$$

Thus, \mathcal{I}_O counts how many objects are contained in some minimal inconsistent subset, that is, the number of objects involved in an inconsistency. For instance, in the US Navy databases containing the ship location data used by Parker et al. (2009) and Parisi et al. (2010), \mathcal{I}_O could be used to get the number of ships that, at the same time, are located in different regions having no common point.

Example 6. For \mathcal{S}_e of our running example, we have $\mathcal{MI}(\mathcal{S}_e) = \{(id_1, b, 1), (id_1, c, 1)\}, \{(id_1, a, 2), (id_1, c, 2)\}$. Therefore, we have $\mathcal{I}_O(\mathcal{S}_e) = 1$, as only id_1 is involved in an inconsistency.

The following theorem shows the postulates that are satisfied by \mathcal{I}_O as well as those that are not satisfied. Free-Formula Independence is satisfied because a free ST atom is not involved in any inconsistency, and thus it does not contribute to the number of objects involved in inconsistency. Dominance is satisfied because, for two ST atoms (id, r, t) and (id, r', t) with $r \subseteq r'$, (id, r, t) is involved in the same inconsistency as (id, r', t) or possibly more, and thus the addition of (id, r, t) to an ST database can only possibly increase the total number of objects involved in inconsistency w.r.t. the addition of (id, r', t) . Equal




Figure 4: ST database \mathcal{S}_2 used in the proof of Theorem 1 for Attenuation.

Conflict and MI-Normalization are satisfied because any minimal inconsistent subset MIS involves exactly one object, and thus $\mathcal{I}_O(MIS) = 1$. The remaining postulates are not satisfied by \mathcal{I}_O , and for each of them we provide a counterexample in the proof.

Theorem 1. *The \mathcal{I}_O inconsistency measure satisfies Free-Formula Independence, Dominance, Equal Conflict, and MI-Normalization, but it does not satisfy Penalty, Super-Additivity, Attenuation, and MI-Separability.*

Proof. We consider the postulates individually.

1. Free-Formula Independence. A free formula does not add an inconsistency.
2. Penalty. Let $\mathcal{S}_1 = \{id\} \times \{r_1, r_2\} \times \{t_1, t_2\}$ where $t_1 \neq t_2$ and $r_1 \cap r_2 = \emptyset$. All the formulas are problematic because of the two inconsistencies but only one object is involved. Hence $\mathcal{I}_O(\mathcal{S}_1) = 1 = \mathcal{I}_O(\mathcal{S}_1 \setminus \{(id, r_2, t_2)\})$.
3. Dominance. Let \mathcal{S} be an ST database. Notice that either $\mathcal{I}_O(\mathcal{S} \cup \{(id, r', t)\}) = \mathcal{I}_O(\mathcal{S})$ or $\mathcal{I}_O(\mathcal{S} \cup \{(id, r', t)\}) = \mathcal{I}_O(\mathcal{S}) + 1$. In the first case, the result is immediate. In the second case, there must be a set of ST atoms $(id, r_1, t), \dots, (id, r_j, t)$ in \mathcal{S} such that $(\cap_{i=1}^j r_i) \cap r' = \emptyset$. But then $r \subseteq r'$ implies that $(\cap_{i=1}^j r_i) \cap r = \emptyset$, and therefore $\mathcal{I}_O(\mathcal{S} \cup \{(id, r, t)\}) = \mathcal{I}_O(\mathcal{S}) + 1$.
4. Super-Additivity. Suppose that two ST databases \mathcal{S} and \mathcal{S}' involve inconsistencies for a single common object. Then, $1 = \mathcal{I}_O(\mathcal{S}) = \mathcal{I}_O(\mathcal{S}') = \mathcal{I}_O(\mathcal{S} \cup \mathcal{S}') < \mathcal{I}_O(\mathcal{S}) + \mathcal{I}_O(\mathcal{S}') = 2$.
5. Attenuation. Let $\mathcal{S}_2 = \{(id, r_1, t_1), (id, r_2, t_1), (id, r_3, t_2), (id, r_4, t_2), (id, r_5, t_2)\}$ where $r_1 \cap r_2 = \emptyset$, $r_3 \cap r_4 \neq \emptyset$, $r_3 \cap r_5 \neq \emptyset$, $r_4 \cap r_5 \neq \emptyset$, and $r_3 \cap r_4 \cap r_5 = \emptyset$ —an ST database of this form is illustrated in Figure 4.

Then, $\mathcal{MI}(\mathcal{S}_2) = \{MIS, MIS'\}$ where $MIS = \{(id, r_1, t_1), (id, r_2, t_1)\}$ and $MIS' = \{(id, r_3, t_2), (id, r_4, t_2), (id, r_5, t_2)\}$.

Here, $2 = |MIS| < |MIS'| = 3$, but $\mathcal{I}_O(MIS) = 1 = \mathcal{I}_O(MIS')$.

6. Equal Conflict. This follows because any minimal inconsistent subset can deal with only one object, hence if MIS is a minimal inconsistent subset then $\mathcal{I}_O(MIS) = 1$.
7. MI-Normalization. All atoms of any minimal inconsistent set MIS must have the same object id ; hence $\mathcal{I}_O(MIS) = 1$.

8. MI-Separability. Let $ID = \{id\}$, $T = \{t\}$, and $Space = \{p_1, p_2, p_3\}$. We will use the following two ST databases: $\mathcal{S}_1 = \{(id, \{p_1\}, t), (id, \{p_2, p_3\}, t)\}$ and $\mathcal{S}_2 = \{(id, \{p_1\}, t), (id, \{p_2\}, t)\}$. Then $\mathcal{MI}(\mathcal{S}_1) = \{\mathcal{S}_1\}$ and $\mathcal{MI}(\mathcal{S}_2) = \{\mathcal{S}_2\}$. Here $\mathcal{S}_1 \cup \mathcal{S}_2 = \{(id, \{p_1\}, t), (id, \{p_2, p_3\}, t), (id, \{p_2\}, t)\}$ and so $\mathcal{MI}(\mathcal{S}_1 \cup \mathcal{S}_2) = \{\mathcal{S}_1, \mathcal{S}_2\} = \mathcal{MI}(\mathcal{S}_1) \cup \mathcal{MI}(\mathcal{S}_2)$ and $\mathcal{MI}(\mathcal{S}_1) \cap \mathcal{MI}(\mathcal{S}_2) = \emptyset$. Thus the conditions for MI-Separability are satisfied. $\mathcal{I}_O(\mathcal{S}_1 \cup \mathcal{S}_2) = 1 \neq \mathcal{I}_O(\mathcal{S}_1) + \mathcal{I}_O(\mathcal{S}_2) = 2$.

□

4.1.2 MEASURING INCONSISTENCY ALONG THE TIME DIMENSION

Similar to the inconsistency measure along the object dimension, a natural inconsistency measure along the time dimension counts how many time values are involved in an inconsistency:

$$\mathcal{I}_T(\mathcal{S}) = |\{t \in T \mid (id, r, t) \in MIS \in \mathcal{MI}(\mathcal{S})\}|.$$

For instance, considering again the US Navy databases containing the ship location data used by Parker et al. (2009) and Parisi et al. (2010), \mathcal{I}_T could be used to get the number of time points at which an inconsistency occurs, that is, at least one ship is located in different regions having no common point.

Example 7. For \mathcal{S}_e we have $\mathcal{I}_T(\mathcal{S}_e) = 2$, as both time values 1 and 2 are involved in an inconsistency.

As stated in the following theorem, \mathcal{I}_T satisfies the same postulates as \mathcal{I}_O (in fact, the result for \mathcal{I}_T can be proved by reasoning analogously to the case of \mathcal{I}_O , often substituting “time” for “object”).

Theorem 2. *The \mathcal{I}_T inconsistency measure satisfies Free-Formula Independence, Dominance, Equal Conflict, and MI-Normalization, but it does not satisfy Penalty, Super-Additivity, Attenuation, and MI-Separability.*

Proof. Analogous to the proof of Theorem 1. □

4.2 Measuring Inconsistency along the Object and Time Dimensions Together

ST databases treat objects and time in a similar manner and differently from space, as seen from Definition 2. So it is natural to combine the object and time dimensions before dealing with the spatial dimension. This can be done by combining the two dimensions individually, that is, computing \mathcal{I}_O and \mathcal{I}_T and then applying some operation(s) to the two numbers. Instead of doing so, we observe that in many cases we are dealing with id, t pairs. Indeed, all the ST atoms of a minimal inconsistent set must have the same id and t values. The following inconsistency measure is thus defined looking at id, t pairs in minimal inconsistent subsets:

$$\mathcal{I}_{OT}(\mathcal{S}) = |\{(id, t) \mid (id, r, t) \in MIS \in \mathcal{MI}(\mathcal{S})\}|.$$

Thus, \mathcal{I}_{OT} counts how many object-time pairs are involved in an inconsistency.

Example 8. In our running example for ST database \mathcal{S}_e , we have $\mathcal{I}_{OT}(\mathcal{S}_e) = 2$, as two pairs of (id, t) values are involved in inconsistencies, namely $(id_1, 1)$ and $(id_1, 2)$.

It should not be surprising that \mathcal{I}_{OT} satisfies exactly the same properties as \mathcal{I}_O and \mathcal{I}_T .

Theorem 3. *The \mathcal{I}_{OT} inconsistency measure satisfies Free-Formula Independence, Dominance, Equal Conflict, and MI-Normalization, but it does not satisfy Penalty, Super-Additivity, Attenuation, and MI-Separability.*

Proof. Analogous to the proof of Theorem 1. □

4.3 Measuring Inconsistency along the Space Dimension

In this subsection we measure inconsistency by the amount of space that the inconsistencies involve. We give two different measures. The first one is similar to what we did for objects and time. The second involves distance.

4.3.1 THE \mathcal{I}_S INCONSISTENCY MEASURE

For an ST database \mathcal{S} we define a region $R_{\mathcal{S}}$ as follows:

$$R_{\mathcal{S}} = \bigcup \{r \mid (id, r, t) \in MIS \in \mathcal{MI}(\mathcal{S})\}.$$

Then, we define $\mathcal{I}_S(\mathcal{S}) = |R_{\mathcal{S}}|$.

Thus, \mathcal{I}_S counts the number of points that are in regions that are involved in an inconsistency.

Example 9. *In \mathcal{S}_e there are three regions involved in $\mathcal{MI}(\mathcal{S}_e)$, namely region a with 49 spatial points in it, region b with 49 but 15 of them also belong to a , and region c with 48 points. Therefore, among the three regions we have $\mathcal{I}_S(\mathcal{S}_e) = |R_{\mathcal{S}_e}| = 131$.*

The next result shows that only one postulate from Definition 6, Free-Formula Independence, is satisfied by \mathcal{I}_S ; we provide counterexamples for the rest of the postulates that show why they are not satisfied. The reason is similar to why such a postulate is satisfied by the inconsistency measures introduced so far, that is, a free ST atom (id, r, t) is not involved in any inconsistency, and thus the points in r cannot contribute to the number of points involved in an inconsistency.

Theorem 4. *The \mathcal{I}_S inconsistency measure satisfies Free-Formula Independence, but it does not satisfy any of the other postulates from Definition 6.*

Proof. We deal with each postulate individually.

1. Free-Formula Independence. A free formula does not add an inconsistency.
2. Penalty. Let $\mathcal{S}_6 = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t), (id_2, \{p_1\}, t), (id_2, \{p_2\}, t)\}$ where $id_1 \neq id_2$ and $p_1 \neq p_2$. All the formulas are problematic because of the two inconsistencies. However, $\mathcal{I}_S(\mathcal{S}_6) = 2 = \mathcal{I}_S(\mathcal{S}_6 \setminus \{(id_1, \{p_1\}, t)\})$.
3. Dominance. Let $\mathcal{S}_7 = \{(id, \{p_1\}, t)\}$, $r = \{p_2\}$, and $r' = \{p_2, p_3\}$, where p_1, p_2 , and p_3 are distinct points. Here $2 = \mathcal{I}_S(\mathcal{S}_7 \cup \{(id, r, t)\}) < \mathcal{I}_S(\mathcal{S}_7 \cup \{(id, r', t)\}) = 3$.
4. Super-Additivity. Consider the ST databases $\mathcal{S}_8 = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t)\}$ and $\mathcal{S}_9 = \{(id_2, \{p_1\}, t), (id_2, \{p_2\}, t)\}$, where $id_1 \neq id_2$ and $p_1 \neq p_2$. Then, $\mathcal{I}_S(\mathcal{S}_8) = \mathcal{I}_S(\mathcal{S}_9) = \mathcal{I}_S(\mathcal{S}_8 \cup \mathcal{S}_9) = 2$. Hence, $\mathcal{I}_S(\mathcal{S}_8 \cup \mathcal{S}_9) < \mathcal{I}_S(\mathcal{S}_8) + \mathcal{I}_S(\mathcal{S}_9)$.

5. Attenuation. Consider the ST database

$$\mathcal{S}_{10} = \{(id, \{p_1\}, t_1), (id, \{p_2\}, t_1), (id, \{p_1, p_2\}, t_2), (id, \{p_2, p_3\}, t_2), (id, \{p_1, p_3\}, t_2)\},$$

where $t_1 \neq t_2$, and p_1, p_2 , and p_3 are distinct points. Clearly, $MIS_4 = \{(id, \{p_1\}, t_1), (id, \{p_2\}, t_1)\}$ and $MIS'_4 = \{(id, \{p_1, p_2\}, t_2), (id, \{p_2, p_3\}, t_2), (id, \{p_1, p_3\}, t_2)\}$ are in $\mathcal{MI}(\mathcal{S}_{10})$, $2 = |MIS_4| < |MIS'_4| = 3$, but $2 = \mathcal{I}_S(MIS_4) < \mathcal{I}_S(MIS'_4) = 3$.

6. Equal Conflict. Let $\mathcal{S}_{11} = \{(id, \{p_1\}, t_1), (id, \{p_2\}, t_1), (id, \{p_1\}, t_2), (id, \{p_2, p_3\}, t_2)\}$ where $t_1 \neq t_2$ and p_1, p_2 , and p_3 are distinct points. Then, $MIS_5 = \{(id, \{p_1\}, t_1), (id, \{p_2\}, t_1)\}$ and $MIS'_5 = \{(id, \{p_1\}, t_2), (id, \{p_2, p_3\}, t_2)\}$ are in $\mathcal{MI}(\mathcal{S}_{11})$, $|MIS_5| = 2 = |MIS'_5|$, but $\mathcal{I}_S(MIS_5) = 2$ and $\mathcal{I}_S(MIS'_5) = 3$.

7. MI-Normalization. Consider $MIS_1 = \{(id, \{p_1\}, t), (id, \{p_2\}, t)\}$ where $p_1 \neq p_2$. Then, $\mathcal{I}_S(MIS_1) = 2$.

8. MI-Separability. Let $ID = \{id\}$, $T = \{t\}$, and $Space = \{p_1, p_2, p_3\}$. We will use the following two ST databases: $\mathcal{S}_1 = \{(id, \{p_1\}, t), (id, \{p_2, p_3\}, t)\}$ and $\mathcal{S}_2 = \{(id, \{p_1\}, t), (id, \{p_2\}, t)\}$. Then $\mathcal{MI}(\mathcal{S}_1) = \{\mathcal{S}_1\}$ and $\mathcal{MI}(\mathcal{S}_2) = \{\mathcal{S}_2\}$. Here $\mathcal{S}_1 \cup \mathcal{S}_2 = \{(id, \{p_1\}, t), (id, \{p_2, p_3\}, t), (id, \{p_2\}, t)\}$ and so $\mathcal{MI}(\mathcal{S}_1 \cup \mathcal{S}_2) = \{\mathcal{S}_1, \mathcal{S}_2\} = \mathcal{MI}(\mathcal{S}_1) \cup \mathcal{MI}(\mathcal{S}_2)$ and $\mathcal{MI}(\mathcal{S}_1) \cap \mathcal{MI}(\mathcal{S}_2) = \emptyset$. Thus the conditions for MI-Separability are satisfied. $\mathcal{I}_S(\mathcal{S}_1 \cup \mathcal{S}_2) = 3 \neq \mathcal{I}_T(\mathcal{S}_1) + \mathcal{I}_T(\mathcal{S}_2) = 5$.

□

Clearly, space is one important dimension for ST databases, however, the rationality postulates are not satisfied for the simple (and intuitive) inconsistency measure \mathcal{I}_S .

4.3.2 THE \mathcal{I}_D INCONSISTENCY MEASURE

For this measure we require a metric $d : Space \times Space \rightarrow [0, \infty)$ satisfying the following axioms, which are the standard properties a metric (or distance function) has to satisfy:

- $d(p, p') = 0$ if and only if $p = p'$ (identity of indiscernibles);
- $d(p, p') = d(p', p)$ (symmetry);
- $d(p, p') \leq d(p, p'') + d(p'', p')$ (triangle inequality).

This allows us to define a function on regions as follows:¹

$$d(r, r') = \min\{d(p, p') \mid p \in r, p' \in r'\}.$$

Next we define distance for minimal inconsistent subsets. Let MIS be a minimal inconsistent subset of an ST database \mathcal{S} : $MIS = \{(id, r_1, t), \dots, (id, r_n, t)\}$. We start by defining n new regions, one for each i , $1 \leq i \leq n$,

$$R_i = \bigcap_{j \neq i} r_j$$

1. Note that $d(r, r')$ is not a metric because the triangle inequality does not hold. For instance, consider $r_1 = \{(x, y) \mid 0 \leq x \leq 2, 0 \leq y \leq 2\}$, $r_2 = \{(x, y) \mid 0 \leq x \leq 7, 3 \leq y \leq 5\}$, and $r_3 = \{(x, y) \mid 5 \leq x \leq 7, 0 \leq y \leq 2\}$. Then, $d(r_1, r_3) = 3 \not\leq d(r_1, r_2) + d(r_2, r_3) = 1 + 1$.

Since MIS is a minimal inconsistent set, for each i , $1 \leq i \leq n$, $R_i \neq \emptyset$, but $\bigcap_{i=1}^n r_i = \emptyset$. We define the value $d(MIS)$ as follows:

$$d(MIS) = \min\{d(R_i, r_i) \mid 1 \leq i \leq n\}.$$

We can think of $d(MIS)$ as the minimal distance in $Space$ between any regions involved in a minimal inconsistent subset.

Then, we define:

$$\mathcal{I}_D(\mathcal{S}) = \sum_{MIS \in \mathcal{MI}(\mathcal{S})} d(MIS).$$

Thus \mathcal{I}_D sums the minimal distances between regions involved in minimal inconsistent subsets for the given ST database.

Example 10. Let $\mathcal{MI}(\mathcal{S}_e) = \{M_1, M_2\}$, where $M_1 = \{(id_1, b, 1), (id_1, c, 1)\}$ and $M_2 = \{(id_1, a, 2), (id_1, c, 2)\}$. Then for M_1 , $R_1 = c$ and $R_2 = b$, $d(M_1) = \min\{d(c, b), d(b, c)\}$, and assuming that d is the Euclidean distance, we obtain $d(M_1) = d(c, b) = d(b, c) = 1$, because a pair of closest points are $(8, 10) \in b$ and $(9, 10) \in c$. In the same way, $d(M_2) = d(a, c) = d(c, a) = 3$, because a pair of closest points are $(6, 12) \in a$ and $(9, 12) \in c$. Therefore, $\mathcal{I}_D(\mathcal{S}) = \sum_{M_i \in \mathcal{MI}(\mathcal{S}_e)} d(M_i) = 3 + 1 = 4$.

Recall that the definition of \mathcal{I}_D relies on an arbitrary metric d over $Space$. Thus, in a sense, \mathcal{I}_D is a *family* of inconsistency measures, one inconsistency measure for each specific metric. The next result shows the postulates from Definition 6 that are satisfied by \mathcal{I}_D and counterexamples for those that are not satisfied by the measure. Notice that when we say that \mathcal{I}_D satisfies a postulate, it means that the postulate is satisfied no matter which metric is employed, that is, \mathcal{I}_D satisfies the postulate *for every* metric. On the other hand, when we say that \mathcal{I}_D does not satisfy a postulate, it means that there is *at least one* metric for which \mathcal{I}_D does not satisfy the postulate—in such a case, we show the metric and the scenario that violates the postulate (notice that there might be specific metrics for which the postulate might be satisfied). Thus, in the following theorem, specific metrics are employed to show that \mathcal{I}_D does not satisfy Attenuation, Equal Conflict, and MI-Normalization (details are reported in the proof). For Free-Formula Independence, Penalty, Dominance, Super-Additivity, and MI-Separability, we show that \mathcal{I}_D satisfies such postulates whichever metric is employed.

Theorem 5. *The \mathcal{I}_D inconsistency measure satisfies Free-Formula Independence, Penalty, Dominance, Super-Additivity, and MI-Separability, but it does not satisfy Attenuation, Equal Conflict, and MI-Normalization.*

Proof. We deal with each postulate individually.

1. Free-Formula Independence. A free formula does not add an inconsistency.
2. Penalty. First of all, notice that, for any two ST databases \mathcal{S} and \mathcal{S}' , if $\mathcal{S} \subseteq \mathcal{S}'$, then $\mathcal{MI}(\mathcal{S}) \subseteq \mathcal{MI}(\mathcal{S}')$. Let \mathcal{S} be an ST database. If $(id, r, t) \in Problematic(\mathcal{S})$ then there exists $MIS \in \mathcal{MI}(\mathcal{S})$ such that $MIS \notin \mathcal{MI}(\mathcal{S} \setminus \{(id, r, t)\})$. Hence, $\mathcal{I}_D(\mathcal{S}) > \mathcal{I}_D(\mathcal{S} \setminus \{(id, r, t)\})$.

3. **Dominance.** First, suppose that $(id, r', t) \notin \text{Problematic}(\mathcal{S} \cup \{(id, r', t)\})$. Then $\mathcal{I}_D(\mathcal{S} \cup \{(id, r', t)\}) = \mathcal{I}_D(\mathcal{S}) \leq \mathcal{I}_D(\mathcal{S} \cup \{(id, r, t)\})$. The other case is where $(id, r', t) \in \text{MIS}' \in \mathcal{MI}(\mathcal{S} \cup \{(id, r', t)\})$. This means that there is a set of ST atoms $(id, r'_1, t), \dots, (id, r'_j, t)$ in \mathcal{S} such that writing R'_j for $\bigcap_{i=1}^j r'_i$, $R'_j \neq \emptyset$, but $R'_j \cap r' = \emptyset$. But then $r \subseteq r'$ implies that $R'_j \cap r = \emptyset$ also. Hence $(id, r, t) \in \text{MIS} \in \mathcal{MI}(\mathcal{S} \cup \{(id, r, t)\})$. Consider how $d(\text{MIS})$ is calculated as the distance between 2 points, say p and p' . One of those points must be in r , say p . Let $p' \in R'_j$ such that $d(\text{MIS}) = d(p, p')$. By $r \subseteq r'$, $p \in r'$, and there may be points in r' closer to R'_j than p . Hence $d(\text{MIS}') \leq d(\text{MIS})$. This is true for every MIS' such that $(id, r', t) \in \text{MIS}'$. Therefore $\mathcal{I}_D(\mathcal{S} \cup \{(id, r, t)\}) \geq \mathcal{I}_D(\mathcal{S} \cup \{(id, r', t)\})$.
4. **Super-Additivity.** Given two ST databases \mathcal{S} and \mathcal{S}' , $\mathcal{S} \cap \mathcal{S}' = \emptyset$ implies that $\mathcal{MI}(\mathcal{S}) \cap \mathcal{MI}(\mathcal{S}') = \emptyset$. Then $\mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}') \subseteq \mathcal{MI}(\mathcal{S} \cup \mathcal{S}')$. So $\mathcal{I}_D(\mathcal{S} \cup \mathcal{S}') = \sum_{\text{MIS} \in \mathcal{MI}(\mathcal{S} \cup \mathcal{S}')} d(\text{MIS}) \geq \mathcal{I}_D(\mathcal{S}) + \mathcal{I}_D(\mathcal{S}') = \sum_{\text{MIS} \in \mathcal{MI}(\mathcal{S})} d(\text{MIS}) + \sum_{\text{MIS} \in \mathcal{MI}(\mathcal{S}')} d(\text{MIS})$.
5. **Attenuation.** Consider the following ST database $\mathcal{S}_{12} = \{(id, \{p_1\}, t_1), (id, \{p_2\}, t_1), (id, \{p_3, p_4\}, t_2), (id, \{p_4, p_5\}, t_2), (id, \{p_3, p_5\}, t_2)\}$, where $t_1 \neq t_2$, all the points p_1, p_2, p_3, p_4, p_5 are distinct, $d(p_1, p_2) = 1$, $d(p_3, p_4) = 2$, $d(p_4, p_5) = 2$, and $d(p_3, p_5) = 2$. Then, we have that $\text{MIS} = \{(id, \{p_1\}, t_1), (id, \{p_2\}, t_1)\}$ and $\text{MIS}' = \{(id, \{p_3, p_4\}, t_2), (id, \{p_4, p_5\}, t_2), (id, \{p_3, p_5\}, t_2)\}$ are in $\mathcal{MI}(\mathcal{S})$, $2 = |\text{MIS}| < |\text{MIS}'| = 3$, but $1 = \mathcal{I}_D(\text{MIS}) < \mathcal{I}_D(\text{MIS}') = 2$.
6. **Equal Conflict.** Let $\mathcal{S}_{13} = \{(id, \{p_1\}, t), (id, \{p_2\}, t), (id, \{p_3\}, t)\}$, where p_1, p_2, p_3 are all distinct, $d(p_1, p_2) = 1$, and $d(p_1, p_3) = 2$. Then $\text{MIS} = \{(id, \{p_1\}, t), (id, \{p_2\}, t)\}$ and $\text{MIS}' = \{(id, \{p_1\}, t), (id, \{p_3\}, t)\}$ are in $\mathcal{MI}(\mathcal{S})$, $|\text{MIS}| = 2 = |\text{MIS}'|$, but $\mathcal{I}_D(\text{MIS}) = 1$ and $\mathcal{I}_D(\text{MIS}') = 2$.
7. **MI-Normalization.** Consider $\text{MIS}_1 = \{(id, \{p_1\}, t), (id, \{p_2\}, t)\}$ where $p_1 \neq p_2$. If $d(p_1, p_2) = 2$ then $\mathcal{I}_D(\text{MIS}_1) = 2$.
8. **MI-Separability.** Consider any two ST databases \mathcal{S} and \mathcal{S}' that satisfy the two properties for MI-Separability. Then,

$$\begin{aligned} \mathcal{I}_D(\mathcal{S} \cup \mathcal{S}') &= \sum_{\text{MIS} \in \mathcal{MI}(\mathcal{S} \cup \mathcal{S}')} d(\text{MIS}) = \sum_{\text{MIS} \in \mathcal{MI}(\mathcal{S})} d(\text{MIS}) + \sum_{\text{MIS} \in \mathcal{MI}(\mathcal{S}')} d(\text{MIS}) = \\ &= \mathcal{I}_D(\mathcal{S}) + \mathcal{I}_D(\mathcal{S}'). \end{aligned}$$

□

4.4 Discussion

Table 1 summarizes which properties are satisfied by the inconsistency measures introduced in this section.

Not surprisingly, \mathcal{I}_O , \mathcal{I}_T , and \mathcal{I}_{OT} satisfy the same set of postulates. As already observed before, the results for \mathcal{I}_T can be proved by reasoning analogously to the case of \mathcal{I}_O , and the same holds for \mathcal{I}_{OT} , which behaves similarly to \mathcal{I}_O and \mathcal{I}_T .

	Dimensional IMs				
	\mathcal{I}_O	\mathcal{I}_T	\mathcal{I}_{OT}	\mathcal{I}_S	\mathcal{I}_D
Free-Formula Ind.	✓	✓	✓	✓	✓
Penalty	✗	✗	✗	✗	✓
Dominance	✓	✓	✓	✗	✓
Super-Additivity	✗	✗	✗	✗	✓
Attenuation	✗	✗	✗	✗	✗
Equal Conflict	✓	✓	✓	✗	✗
MI-Normalization	✓	✓	✓	✗	✗
MI-Separability	✗	✗	✗	✗	✓

Table 1: Postulates satisfaction of dimensional inconsistency measures.

The two inconsistency measures based on the space dimension, namely \mathcal{I}_S and \mathcal{I}_D , behave quite differently. Specifically, \mathcal{I}_S satisfies only Free-Formula Independence, which is also satisfied by all other inconsistency measures. Regarding \mathcal{I}_D , on the one hand, it satisfies postulates that are not satisfied by any of the other inconsistency measures (namely, Penalty, Super-Additivity, and MI-Separability), on the other hand, it does not satisfy postulates (namely, Equal Conflict and MI-Normalization) that are satisfied by other inconsistency measures (namely, \mathcal{I}_O , \mathcal{I}_T , and \mathcal{I}_{OT}). Thus, from this point of view, \mathcal{I}_D is incomparable to \mathcal{I}_O , \mathcal{I}_T , and \mathcal{I}_{OT} . We also notice that Attenuation is not satisfied by any of the measures (recall that Attenuation and MI-Normalization cannot be jointly satisfied).

In Section 6, we will propose dimensional counterparts of some postulates that better fit the dimensional nature of the inconsistency measures, and, as we will show, such dimensional postulates are satisfied in most of the cases.

5. Repair-Based Inconsistency Measures

In this section, we define three new inconsistency measures, namely \mathcal{I}_{id} , \mathcal{I}_{time} , and \mathcal{I}_{region} , which are based on the cost of restoring consistency along the object, time, and spatial dimensions. We also introduce measure \mathcal{I}_{card} , which is not dimensional (it deals with whole ST atoms), but it is also based on the cost of restoring consistency in a minimal way as the other three measures introduced in this section.

In this work we rely on the notion of a *repair* commonly used when dealing with inconsistency (Bertossi, 2011), which in our case is a consistent ST database that minimally differs from the original one. In this regard, we will explain two aspects: what kind of modifications can be applied to the original ST database so as to obtain a repair, and how to measure the difference between two ST databases (in particular, the original and the “repaired” ones).

To formalize “repair-based” inconsistency measures, we first need to define repairs for ST databases, inspired by the approach for *probabilistic* spatio-temporal databases by Parisi and Grant (2017), and then define an inconsistency measure as the minimum cost of a repair—a similar idea has been recently considered by Bertossi (2018) to define inconsistency

measures for relational databases, where the issues concerning spatio-temporal data, as well as postulates satisfaction, are not considered.

Before defining the concept of repair for ST databases, we formalize the basic concepts of *update* and *correction*, which will be used to define repairs. Basically, updates are ways of modifying ST databases by changing some attribute values of ST atoms, while corrections consist of updates that restore consistency.

5.1 Updates and Corrections

In general, an update of an ST atom a is an ST atom a' derived from a by changing (at most) one of its dimensions. Hence, we will deal with 3 types of updates.

Definition 7 (Update). *Given an ST atom $a = (id, r, t)$,*

1. *an id-update of a is an ST atom $a' = (id', r, t)$;*
2. *a time-update of a is an ST atom $a' = (id, r, t')$;*
3. *a region-update of a is an ST atom $a' = (id, r', t)$.*

Thus, an id-update (resp., time-update, region-update) of a is either the result of changing the id value (resp., time value, region) of a or a itself.

Corrections are specific ways of changing an inconsistent ST database to one that is consistent using updates.

Definition 8 (Correction). *Let \mathcal{S} be an ST database. A consistent ST database \mathcal{S}' is called an X -correction of \mathcal{S} , where $X \in \{id, time, region\}$, if there is a surjective function $Xcorr : \mathcal{S} \rightarrow \mathcal{S}'$ such that for every $a \in \mathcal{S}$, $Xcorr(a)$ is an X -update of a .*

So an id-correction allows only id-updates, and the other types of corrections are defined analogously. We assure that every element in \mathcal{S}' is the result of an update of an element of \mathcal{S} . Thus, it is possible that $|\mathcal{S}'| < |\mathcal{S}|$ because, for instance, for an id-correction it may happen that $idcorr(a_1) = idcorr(a_2)$ where a_1, a_2 are two distinct ST atoms in \mathcal{S} (a similar behavior can occur also for the other types of corrections).

Example 11. *Consider the atoms $a = (id_1, b, 1)$ and $b = (id_1, a, 2)$ in our running example, where id-updates of a and b , respectively, could be $a' = (id_2, b, 1)$ and $b' = (id_2, a, 2)$. In this case, if we define $idcorr(id_1, b, 1) = (id_2, b, 1)$, $idcorr(id_1, a, 2) = (id_2, a, 2)$, and $idcorr(x) = x$ for all other ST atoms x in \mathcal{S}_e , then the resulting ST database induced by $idcorr$ is an id-correction of \mathcal{S}_e .*

5.2 Measuring Inconsistency by the Repairing Cost

Repairs are *minimal* corrections. In general, we would like to preserve as much of the information in the original database as possible. The idea for the three types of corrections in Definition 8 is to define a distance between two ST atoms for the type of update under consideration, take the sum of the distances between all atoms and their updated counterparts, and find the corrections whose sum is minimal. To measure the distance between two ST atoms when id- (resp., time-, region-) corrections are considered, we will rely on a

metric on the set of object identifiers (resp., time values, regions). Rather than committing to specific metrics, we will assume arbitrary ones, as done for the \mathcal{I}_D inconsistency measure (cf. Section 4.3.2). As discussed for the \mathcal{I}_D inconsistency measure, we point out that when one of the inconsistency measures \mathcal{I}_{id} , \mathcal{I}_{time} , \mathcal{I}_{region} satisfies a postulate, it does so *for every* metric employed, while when an inconsistency measure does not satisfy a postulate, it means that there is at least one metric for which the postulate is not satisfied. After the aforementioned inconsistency measures based on repairing along the object identifier, time, and space dimensions, we will consider also maximal consistent subsets of an ST database, which correspond to a repair strategy performing atom deletions.

5.2.1 THE MEASURE \mathcal{I}_{id}

Our first repairing strategy is based on minimally updating the object identifiers of the ST atoms to restore consistency. This means assuming that some object may be associated with a wrong identifier, not corresponding to the object from which an ST atom has been generated. As an example, license plate recognition systems may give wrong identifiers to vehicles passing in the proximity of a roadside sensor due to both bad external conditions (such as sun and headlights, dirty plates) and the limited level of the recognition software and vision hardware employed. Nevertheless, recognizing systems are generally designed to minimize the error rate and in particular the difference between the acquired and the true values. Thus, this is taken into account by id-repairs by choosing as good corrections those for which the difference between the acquired and the true values is minimized according to a metric that measures such a difference.

We assume a metric d_{ID} on the set ID of object identifiers for measuring the cost of updating id to id' for an ST atom. Thus, the cost of fixing the object identifier id of an ST atom a by setting it to id' is given by the distance $d_{ID}(id, id')$ between the two identifiers. For instance, if the objects are cars and the identifiers are the license plate numbers, then d_{ID} can be the edit-distance between the strings encoding these numbers. In the following, we denote by $cost_{id}(a, a')$ the cost of changing atom a into an id-update a' (here the only difference between a and a' is that the object identifier id in a was changed to id' in a'). Thus, $cost_{id}(a, a') = d_{ID}(id, id')$. Moreover, given an ST database \mathcal{S} and an id-correction \mathcal{S}' , let C_{id} be the set of all functions $idcorr$ as per Definition 8. We define $cost_{id}(\mathcal{S}, \mathcal{S}') = \min_{idcorr \in C_{id}} \{ \sum_{a \in \mathcal{S}} cost_{id}(a, idcorr(a)) \}$.

Definition 9 (id-repair). *An id-repair for an ST database \mathcal{S} is an id-correction \mathcal{S}' of \mathcal{S} such that for all id-corrections \mathcal{S}'' of \mathcal{S} , $cost_{id}(\mathcal{S}, \mathcal{S}') \leq cost_{id}(\mathcal{S}, \mathcal{S}'')$.*

We are now ready to define the measure \mathcal{I}_{id} . It is the cost of an id-repair, if one exists. It is possible that there is no id-correction at all, as all possible sets of id-updates result in an inconsistent database. Recall that ID is fixed.

Definition 10 (Measure \mathcal{I}_{id}). *Given an ST database \mathcal{S} , $\mathcal{I}_{id}(\mathcal{S}) = \infty$ if there is no id-repair; otherwise $\mathcal{I}_{id}(\mathcal{S}) = cost_{id}(\mathcal{S}, \mathcal{S}')$ where \mathcal{S}' is an id-repair for \mathcal{S} .*

Example 12. *For the ST database \mathcal{S}_e of our running example, let $idcorr(id_1, b, 1) = (id_2, b, 1)$, $idcorr(id_1, a, 2) = (id_2, a, 2)$, and $idcorr(x) = x$ for all other ST atoms x . Assuming a cost based on edit distance, $idcorr$ yields an id-repair, and we have $\mathcal{I}_{id}(\mathcal{S}_e) = 2$.*

The following result shows that \mathcal{I}_{id} satisfies Dominance and Super-Additivity, but in general it does not satisfy the other postulates, as shown in the counterexamples reported in the proof (recall that counterexamples employ specific metrics d_{ID}).

Theorem 6. *The \mathcal{I}_{id} inconsistency measure satisfies Dominance and Super-Additivity, but it does not satisfy Free-Formula Independence, Penalty, Attenuation, Equal Conflict, MI-Normalization, and MI-Separability.*

Proof. In all the counterexamples we use $d_{ID}(id_i, id_j) = |i - j|$.

1. Free-Formula Independence. Let $ID = \{id_1, id_2\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3\}$, and $\mathcal{S} = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t), (id_2, \{p_3\}, t)\}$. Then $\mathcal{I}_{id}(\mathcal{S}) = \infty \neq \mathcal{I}_{id}(\mathcal{S} \setminus \{(id_2, \{p_3\}, t)\}) = 1$.
2. Penalty. Let $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3\}$, and $\mathcal{S} = \{(id, \{p_1\}, t), (id, \{p_2\}, t), (id, \{p_3\}, t)\}$. Clearly, $(id, \{p_3\}, t) \in Problematic(\mathcal{S})$ but $\mathcal{I}_{id}(\mathcal{S}) = \infty \not\asymp \mathcal{I}_{id}(\mathcal{S} \setminus \{(id, \{p_3\}, t)\}) = \infty$.
3. Dominance. Let \mathcal{S}' be an id-repair for $\mathcal{S} \cup \{(id, r, t)\}$, and let $idcorr$ be an id-correction s.t. $\mathcal{I}_{id}(\mathcal{S} \cup \{(id, r, t)\}) = cost_{id}(\mathcal{S} \cup \{(id, r, t)\}, \mathcal{S}') = \sum_{a \in \mathcal{S} \cup \{(id, r, t)\}} cost_{id}(a, idcorr(a))$. Since $r \subseteq r'$, it is the case that $idcorr$ is an id-correction also for $\mathcal{S} \cup \{(id, r', t)\}$ (in fact, every inconsistency that is solved w.r.t. (id, r, t) is also solved w.r.t. (id, r', t) , where $r \subseteq r'$). By the definition of an id-repair, every id-repair \mathcal{S}'' for $\mathcal{S} \cup \{(id, r', t)\}$ is s.t. $\mathcal{I}_{id}(\mathcal{S} \cup \{(id, r', t)\}) = cost_{id}(\mathcal{S} \cup \{(id, r', t)\}, \mathcal{S}'') \leq \sum_{a \in \mathcal{S} \cup \{(id, r', t)\}} cost_{id}(a, idcorr(a))$, from which it follows that $\mathcal{I}_{id}(\mathcal{S} \cup \{(id, r', t)\}) \leq \mathcal{I}_{id}(\mathcal{S} \cup \{(id, r, t)\})$.
4. Super-Additivity. Let \mathcal{S}^* be an id-repair for $\mathcal{S} \cup \mathcal{S}'$, and let $idcorr$ be an id-correction such that $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') = cost_{id}(\mathcal{S} \cup \mathcal{S}', \mathcal{S}^*) = \sum_{a \in \mathcal{S} \cup \mathcal{S}'} cost_{id}(a, idcorr(a))$. Given this, we define restrictions of $idcorr$ to the atoms in \mathcal{S} and \mathcal{S}' , respectively, as follows:
 $idcorr_1 : \mathcal{S} \rightarrow \mathcal{S}_1^*$ such that $idcorr_1(a) = idcorr(a)$ and
 $idcorr_2 : \mathcal{S}' \rightarrow \mathcal{S}_2^*$ such that $idcorr_2(a) = idcorr(a)$,
where \mathcal{S}_1^* (resp., \mathcal{S}_2^*) is the image of $idcorr$ under \mathcal{S} (resp., \mathcal{S}').
We can easily check that $idcorr_1$ and $idcorr_2$ are corrections for \mathcal{S} and \mathcal{S}' , respectively, since $idcorr$ is a correction for $\mathcal{S} \cup \mathcal{S}'$. Also, since $\mathcal{S} \cap \mathcal{S}' = \emptyset$, it holds that $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') = \sum_{a \in \mathcal{S} \cup \mathcal{S}'} cost_{id}(a, idcorr(a)) = \sum_{a \in \mathcal{S}} cost_{id}(a, idcorr_1(a)) + \sum_{a \in \mathcal{S}'} cost_{id}(a, idcorr_2(a))$. Finally, by the definition of an id-repair, every id-repair \mathcal{S}_1 for \mathcal{S} (resp., \mathcal{S}_2 for \mathcal{S}') is such that $\mathcal{I}_{id}(\mathcal{S}) = cost_{id}(\mathcal{S}, \mathcal{S}_1) \leq \sum_{a \in \mathcal{S}} cost_{id}(a, idcorr_1(a))$ (resp., $\mathcal{I}_{id}(\mathcal{S}') = cost_{id}(\mathcal{S}', \mathcal{S}_2) \leq \sum_{a \in \mathcal{S}'} cost_{id}(a, idcorr_2(a))$), from which it follows that $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}_{id}(\mathcal{S}) + \mathcal{I}_{id}(\mathcal{S}')$.
5. Attenuation. Let $ID = \{id_1, id_2\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3\}$, $\mathcal{S} = \{(id_1, \{p_1, p_2\}, t), (id_1, \{p_2, p_3\}, t), (id_1, \{p_1, p_3\}, t), (id_2, \{p_1\}, t), (id_2, \{p_2\}, t)\}$. There are two minimal inconsistent subsets, $MIS = \{(id_2, \{p_1\}, t), (id_2, \{p_2\}, t)\}$ and $MIS' = \{(id_1, \{p_1, p_2\}, t), (id_1, \{p_2, p_3\}, t), (id_1, \{p_1, p_3\}, t)\}$. Also, let d_{ID} be such that: $d_{ID}(id_x, id_y) \geq 1$ for all $id_x \neq id_y \in ID$ and $d_{ID}(id_1, id_2) = 1$. Then, $|MIS| < |MIS'|$, but $\mathcal{I}_{id}(MIS) = 1 \not\asymp \mathcal{I}_{id}(MIS') = 1$.

6. Equal Conflict. Let $ID = \{id_1, id_3, id_4\}$, $T = \{t\}$, $Space = \{p_1, p_2\}$, and $\mathcal{S} = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t), (id_3, \{p_1\}, t), (id_3, \{p_2\}, t)\}$. Clearly, there are two minimal inconsistent subsets, namely $MIS = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t)\}$ and $MIS' = \{(id_3, \{p_1\}, t), (id_3, \{p_2\}, t)\}$. Then, $|MIS| = |MIS'|$, but $\mathcal{I}_{id}(MIS) = 2 \neq \mathcal{I}_{id}(MIS') = 1$.
7. MI-Normalization. Let $ID = \{id_1, id_3\}$, $T = \{t\}$, $Space = \{p_1, p_2\}$, and $\mathcal{S} = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t)\}$. Thus, there is $MIS = \mathcal{S} \in \mathcal{MI}(\mathcal{S})$ such that $\mathcal{I}_{id}(MIS) = 2 \neq 1$.
8. MI-Separability. Let us consider $ID = \{id_1, id_2\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3, p_4\}$, $\mathcal{S} = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t)\}$, and $\mathcal{S}' = \{(id_2, \{p_3\}, t), (id_2, \{p_4\}, t)\}$. Although $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \{\mathcal{S}, \mathcal{S}'\} = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$ and $\mathcal{MI}(\mathcal{S}) \cap \mathcal{MI}(\mathcal{S}') = \emptyset$, $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') = \infty \neq \mathcal{I}_{id}(\mathcal{S}) + \mathcal{I}_{id}(\mathcal{S}') = 1 + 1 = 2$.

□

5.2.2 THE MEASURE \mathcal{I}_{time}

The next repairing strategy we use is based on minimally updating the time values associated with the ST atoms. This means assuming that some error occurred when assigning time stamps. This may happen for example in a tracking system where sensors report the position and time that they sense to a server application that combines and stores this information. If the sensors are not synchronized, considering together the time values obtained from the several sources is likely to reveal some inconsistency. Again, in the following, to define a time-repair we rely on the assumption that the differences between the actual time values and those detected by sensors are generally low, and thus time values in a repair should be as close as possible to the original ones.

We need a metric d_T on the set of time values in T measuring the cost of updating t to t' . For this purpose we can use for instance $d_T(t, t') = |t' - t|$. Then, we denote by $cost_{time}(a, a')$ the cost of changing atom a to a time-update a' , where t in a was changed to t' in a' , and define $cost_{time}(a, a') = d_T(t, t')$. Finally, given an ST database \mathcal{S} and a time-correction \mathcal{S}' , let C_t be the set of all functions *timecorr* as per Definition 8. We define $cost_{time}(\mathcal{S}, \mathcal{S}') = \min_{timecorr \in C_t} \{\sum_{a \in \mathcal{S}} cost_{time}(a, timecorr(a))\}$.

The notions of time-repair and \mathcal{I}_{time} are defined analogously to id-repair and \mathcal{I}_{id} but for the time dimension.

Definition 11 (time-repair). *A time-repair for an ST database \mathcal{S} is a time-correction \mathcal{S}' of \mathcal{S} such that for all time-corrections \mathcal{S}'' of \mathcal{S} , $cost_{time}(\mathcal{S}, \mathcal{S}') \leq cost_{time}(\mathcal{S}, \mathcal{S}'')$.*

Hence, we define measure \mathcal{I}_{time} as the cost of a time-repair, if one exists.

Definition 12 (Measure \mathcal{I}_{time}). *Given an ST database \mathcal{S} , $\mathcal{I}_{time}(\mathcal{S}) = \infty$ if there is no time-repair, otherwise $\mathcal{I}_{time}(\mathcal{S}) = cost_{time}(\mathcal{S}, \mathcal{S}')$ where \mathcal{S}' is a time-repair for \mathcal{S} .*

Example 13. *For \mathcal{S}_e , let $timecorr(id_1, c, 1) = (id_1, c, 2)$, $timecorr(id_1, a, 2) = (id_1, a, 1)$, and $timecorr(x) = x$ for all other ST atoms x . Using $d_T(t, t') = |t - t'|$, *timecorr* yields a time-repair and $\mathcal{I}_{time}(\mathcal{S}_e) = 2$.*

As stated in the following theorem, \mathcal{I}_{time} satisfies that same postulates as \mathcal{I}_{id} . As for \mathcal{I}_{id} , when \mathcal{I}_{time} satisfies a postulate, it does so for every metric d_T , while when a postulate is not satisfied by \mathcal{I}_{time} , a counterexample employing a specific metric d_T is used.

Theorem 7. *The \mathcal{I}_{time} inconsistency measure satisfies Dominance and Super-Additivity, but it does not satisfy Free-Formula Independence, Penalty, Attenuation, Equal Conflict, MI-Normalization, and MI-Separability.*

Proof. The proofs are the same as for Theorem 6 with the id and t values switched. In all the counterexamples, the standard time distance is used, that is, $d_T(t, t') = |t - t'|$. \square

The counterexamples reported in the proof of Theorem 6, which apply also to Theorem 7, show that the same kind of issues prevent both \mathcal{I}_{id} and \mathcal{I}_{time} from satisfying other postulates besides Dominance and Super-Additivity.

5.2.3 THE MEASURE \mathcal{I}_{region}

The next repairing strategy we use is based on minimally updating regions in ST atoms. Similar to the previously introduced notions of repairs, we use a metric $d_R(r, r')$ on the set of regions. For instance, since a region is a set of point locations, we might measure the cost of updating a region r into a region r' as the cardinality of their symmetric difference, that is, we might define $d_R(r, r') = |(r \setminus r') \cup (r' \setminus r)|$.

We denote by $cost_{region}(a, a')$ the cost of changing atom a to a region-update a' , where r in a was changed to r' in a' , and define $cost_{region}(a, a') = d_R(r, r')$. Given an ST database \mathcal{S} and a region-correction \mathcal{S}' , let C_r be the set of all functions $regcorr$ as per Definition 8. We define $cost_{region}(\mathcal{S}, \mathcal{S}') = \min_{regcorr \in C_r} \{\sum_{a \in \mathcal{S}} cost_{region}(a, regcorr(a))\}$.

The notions of region-repair and \mathcal{I}_{region} are defined analogously to the previous ones.

Definition 13 (region-repair). *A region-repair for an ST database \mathcal{S} is a region-correction \mathcal{S}' of \mathcal{S} such that for all region-corrections \mathcal{S}'' of \mathcal{S} , $cost_{region}(\mathcal{S}, \mathcal{S}') \leq cost_{region}(\mathcal{S}, \mathcal{S}'')$.*

Measure \mathcal{I}_{region} is defined as the cost of a region-repair.

Definition 14 (Measure \mathcal{I}_{region}). *Given an ST database \mathcal{S} , $\mathcal{I}_{region}(\mathcal{S}) = cost_{region}(\mathcal{S}, \mathcal{S}')$ where \mathcal{S}' is a region-repair for \mathcal{S} .*

Example 14. *Assuming that regions must be rectangles and the cardinality of the symmetric difference between two regions is the cost function, a region-repair for \mathcal{S}_e is the region-correction given by the function $regcorr$ defined as follows: $regcorr(id_1, c, 1) = (id_1, c', 1)$, $regcorr(id_2, c, 2) = (id_2, c'', 2)$, and $regcorr(x) = x$ for all other ST atoms x , where $c' = ((8, 9), (16, 14))$ and $c'' = ((6, 9), (16, 14))$. Hence, $\mathcal{I}_{region}(\mathcal{S}_e) = 6 + 18 = 24$.*

The following theorem shows that \mathcal{I}_{region} satisfies Super-Additivity, but it does not satisfy all the other postulates—for the latter, we exhibit counterexamples (where specific d_R metrics are employed).

Theorem 8. *The \mathcal{I}_{region} inconsistency measure satisfies Super-Additivity, but it does not satisfy Free-Formula Independence, Penalty, Dominance, Attenuation, Equal Conflict, MI-Normalization, and MI-Separability.*

Proof. We deal with each postulate individually.

1. Free-Formula Independence. Let $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3, p_4, p_5\}$, and $\mathcal{S} = \{(id, \{p_1\}, t), (id, \{p_2\}, t), (id, \{p_1, p_2\}, t)\}$, where $d_R(\{p_1\}, \{p_3, p_4\}) = d_R(\{p_2\}, \{p_4, p_5\}) = 1$ and $d_R(r_x, r_y) = 3$ for all other pairs of distinct regions r_x and r_y . Then, $(id, \{p_1, p_2\}, t) \in Free(\mathcal{S})$, and $\mathcal{I}_{region}(\mathcal{S}) = 3 \neq \mathcal{I}_{region}(\mathcal{S} \setminus \{(id, \{p_1, p_2\}, t)\}) = 2$.
2. Penalty. Let $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3, p_4\}$, and $\mathcal{S} = \{(id, \{p_1, p_2\}, t), (id, \{p_2, p_3\}, t), (id, \{p_4\}, t)\}$, where $d_R(r_x, r_y) = 1$ for all pairs of distinct regions r_x and r_y . Notice that every ST atom in \mathcal{S} is problematic. A region-repair for \mathcal{S} is obtained by updating $\{p_4\}$ into $\{p_1, p_2\}$ (or $\{p_2, p_3\}$)—the so obtained ST database is consistent and has minimal cost as at least one region has to be updated and every update has cost 1. Thus, $\mathcal{I}_{region}(\mathcal{S}) = 1 \not\geq \mathcal{I}_{region}(\mathcal{S} \setminus \{(id, \{p_1, p_2\}, t)\}) = 1$.
3. Dominance. Let $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3, p_4\}$, and $\mathcal{S} = \{(id, \{p_1, p_2\}, t)\}$. Consider two regions $r = \{p_3\}$ and $r' = \{p_3, p_4\}$. Let d_R be defined as follows: $d_R(r, \{p_2\}) = 1$, and $d_R(r_x, r_y) = 2$ for all other pairs of distinct regions r_x and r_y . A region-repair for $\mathcal{S} \cup \{(id, r, t)\}$ is obtained by updating r into $\{p_2\}$ —the so obtained ST database is consistent and has minimal cost, namely 1, as at least one region has to be updated and every update has cost at least 1. A region-repair for $\mathcal{S} \cup \{(id, r', t)\}$ can be obtained by updating r' into a region that overlaps with $\{p_1, p_2\}$ (or updating $\{p_1, p_2\}$ into a region that overlaps with r')—the so obtained ST database is consistent and has minimal cost, namely 2, as at least one region has to be updated, the only update with cost 1 (i.e., updating r into $\{p_2\}$ or vice versa) cannot be applied to $\mathcal{S} \cup \{(id, r', t)\}$, and all other updates have cost 2. Thus, $\mathcal{I}(\mathcal{S} \cup \{(id, r, t)\}) < \mathcal{I}(\mathcal{S} \cup \{(id, r', t)\})$.
4. Super-Additivity. Let \mathcal{S}^* be a region-repair for $\mathcal{S} \cup \mathcal{S}'$, and $regcorr$ be a region-correction s.t. $\mathcal{I}_{region}(\mathcal{S} \cup \mathcal{S}') = cost_{region}(\mathcal{S} \cup \mathcal{S}', \mathcal{S}^*) = \sum_{a \in \mathcal{S} \cup \mathcal{S}'} cost_{region}(a, regcorr(a))$. Let $regcorr_1$ and $regcorr_2$ be restrictions of $regcorr$ to the ST atoms in \mathcal{S} and \mathcal{S}' , respectively, defined as follows:
 $regcorr_1 : \mathcal{S} \rightarrow \mathcal{S}_1^*$ such that $regcorr_1(a) = regcorr(a)$ for all $a \in \mathcal{S}$, and
 $regcorr_2 : \mathcal{S}' \rightarrow \mathcal{S}_2^*$ such that $regcorr_2(a) = regcorr(a)$ for all $a \in \mathcal{S}'$,
 where \mathcal{S}_1^* (resp., \mathcal{S}_2^*) is the image of $regcorr$ under \mathcal{S} (resp., \mathcal{S}').
 It is easy to check that $regcorr_1$ and $regcorr_2$ are region-corrections for \mathcal{S} and \mathcal{S}' , respectively, since $regcorr$ is a region-correction for $\mathcal{S} \cup \mathcal{S}'$. Also, since $\mathcal{S} \cap \mathcal{S}' = \emptyset$, we have $\mathcal{I}_{region}(\mathcal{S} \cup \mathcal{S}') = \sum_{a \in \mathcal{S} \cup \mathcal{S}'} cost_{region}(a, regcorr(a)) = \sum_{a \in \mathcal{S}} cost_{region}(a, regcorr_1(a)) + \sum_{a \in \mathcal{S}'} cost_{region}(a, regcorr_2(a))$. Finally, by definition of region-repair, we have that every region-repair \mathcal{S}_1 for \mathcal{S} (resp., \mathcal{S}_2 for \mathcal{S}') is such that $\mathcal{I}_{region}(\mathcal{S}) = cost_{region}(\mathcal{S}, \mathcal{S}_1) \leq \sum_{a \in \mathcal{S}} cost_{region}(a, regcorr_1(a))$ (resp., we have that $\mathcal{I}_{region}(\mathcal{S}') = cost_{region}(\mathcal{S}', \mathcal{S}_2) \leq \sum_{a \in \mathcal{S}'} cost_{region}(a, regcorr_2(a))$), from which it follows $\mathcal{I}_{region}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}_{region}(\mathcal{S}) + \mathcal{I}_{region}(\mathcal{S}')$.
5. Attenuation. Let $ID = \{id\}$, $T = \{1, 2\}$, $Space = \{p_1, p_2, p_3\}$, and $\mathcal{S} = \{(id, \{p_1, p_2\}, 1), (id, \{p_2, p_3\}, 1), (id, \{p_1, p_3\}, 1), (id, \{p_1\}, 2), (id, \{p_2\}, 2)\}$. where $d_R(r_x, r_y) = 1$ for all pairs of distinct regions r_x and r_y . Clearly, there are two minimal inconsistent subsets, namely $MIS = \{(id, \{p_1, p_2\}, 1), (id, \{p_2, p_3\}, 1), (id, \{p_1, p_3\}, 1)\}$ and $MIS' =$

$\{(id, \{p_1\}, 2), (id, \{p_2\}, 2)\}$. Then, $|MIS| = 3 > |MIS'| = 2$ but $\mathcal{I}_{region}(MIS) = 1 \not\prec \mathcal{I}_{region}(MIS') = 1$.

6. Equal Conflict. Let $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3\}$, and $\mathcal{S} = \{(id, \{p_1\}, t), (id, \{p_2\}, t), (id, \{p_3\}, t)\}$, where $d_R(\{p_1\}, \{p_2\}) = d_R(\{p_2\}, \{p_3\}) = 1$ and $d_R(r_x, r_y) = 2$ for all other pairs of distinct regions r_x and r_y . Clearly, there are three minimal inconsistent subsets, namely $MIS = \{(id, \{p_1\}, t), (id, \{p_2\}, t)\}$, $MIS' = \{(id, \{p_1\}, t), (id, \{p_3\}, t)\}$, and $MIS'' = \{(id, \{p_2\}, t), (id, \{p_3\}, t)\}$. Then, $|MIS| = |MIS'| = 2$, but $\mathcal{I}_{region}(MIS) = 1 \neq \mathcal{I}_{region}(MIS') = 2$.
7. MI-Normalization. The scenario above discussed for Equal Conflict shows a minimal inconsistent subset MIS' with $\mathcal{I}_{region}(MIS') = 2$.
8. MI-Separability. Let us consider $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3, p_4\}$, and $\mathcal{S} = \{(id, \{p_1\}, t), (id, \{p_2, p_3\}, t)\}$ and $\mathcal{S}' = \{(id, \{p_1\}, t), (id, \{p_3, p_4\}, t)\}$, where $d_R(r_x, r_y) = 1$ for all pairs of distinct regions r_x and r_y . Then $\mathcal{MI}(\mathcal{S}) = \{\mathcal{S}\}$, $\mathcal{MI}(\mathcal{S}') = \{\mathcal{S}'\}$, and $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \{\mathcal{S}, \mathcal{S}'\}$. Therefore $\mathcal{I}_{region}(\mathcal{S} \cup \mathcal{S}') = 1 \neq \mathcal{I}_{region}(\mathcal{S}) + \mathcal{I}_{region}(\mathcal{S}') = 1 + 1 = 2$.

□

Similar to \mathcal{I}_{id} and \mathcal{I}_{time} , \mathcal{I}_{region} satisfies Super-Additivity. As opposed to \mathcal{I}_{id} and \mathcal{I}_{time} , \mathcal{I}_{region} does not satisfy Dominance. All other postulates are not satisfied by \mathcal{I}_{id} , \mathcal{I}_{time} , and \mathcal{I}_{region} . Thus, \mathcal{I}_{id} and \mathcal{I}_{time} satisfy one postulate more than \mathcal{I}_{region} . Interestingly, we will show in Section 6 that \mathcal{I}_{region} satisfies one *dimensional* postulate more than \mathcal{I}_{id} and \mathcal{I}_{time} .

5.2.4 THE MEASURE \mathcal{I}_{card}

The last repairing strategy we consider relies on assuming that some ST atoms were wrongly generated and thus need to be removed to restore consistency. We require that the result be a maximal consistent subset, for otherwise we would be deleting atoms that can be safely left in. Moreover, we require that the number of removed atoms be minimal.

Definition 15 (card-repair). *A card-repair for an ST database \mathcal{S} is a consistent subset \mathcal{S}' of \mathcal{S} such that $|\mathcal{S}'| = \max \{|\mathcal{S}''| \text{ such that } \mathcal{S}'' \text{ is a consistent subset of } \mathcal{S}\}$. The cost of card-repair \mathcal{S}' for \mathcal{S} is $cost_{card}(\mathcal{S}, \mathcal{S}') = |\mathcal{S}| - |\mathcal{S}'|$.*

Card-repairs assume that the information encoded by each ST atom is either fully correct or incorrect, and thus they delete a minimal number of whole atoms. Measure \mathcal{I}_{card} is the cost of a card-repair.

Definition 16 (Measure \mathcal{I}_{card}). *Given an ST database \mathcal{S} , $\mathcal{I}_{card}(\mathcal{S}) = cost_{card}(\mathcal{S}, \mathcal{S}')$ where \mathcal{S}' is a card-repair for \mathcal{S} .*

Example 15. *For \mathcal{S}_e at least 2 ST atoms must be deleted to attain consistency. Hence, $\mathcal{I}_{card}(\mathcal{S}_e) = 2$.*

The following theorem shows the postulates satisfied by \mathcal{I}_{card} . Interestingly, \mathcal{I}_{card} behaves differently from the other repair-based inconsistency measures introduced so far in that more postulates are satisfied.

Theorem 9. *The \mathcal{I}_{card} inconsistency measure satisfies Free-Formula Independence, Dominance, Super-Additivity, Equal Conflict, and MI-Normalization, but it does not satisfy Penalty, Attenuation, and MI-Separability.*

- Proof.* 1. Free-Formula Independence. Let \mathcal{S}' be a card-repair for \mathcal{S} . Notice that if $(id, r, t) \in Free(\mathcal{S})$ then $(id, r, t) \in \mathcal{S}'$. Then, $\mathcal{S}'' = \mathcal{S}' \setminus \{(id, r, t)\}$ is a card-repair for $\mathcal{S} \setminus \{(id, r, t)\}$. In fact, \mathcal{S}'' is consistent, as it is a subset of \mathcal{S}' . Moreover, \mathcal{S}'' has to be of maximal cardinality. Suppose by contradiction that there exists a consistent subset \mathcal{S}^* of $\mathcal{S} \setminus \{(id, r, t)\}$ such that $|\mathcal{S}^*| > |\mathcal{S}''|$. Then, $\mathcal{S}^* \cup \{(id, r, t)\}$ is a consistent subset of \mathcal{S} and has strictly higher cardinality than \mathcal{S}' , which is a contradiction. Clearly, $|\mathcal{S}| - |\mathcal{S}'| = |\mathcal{S} \setminus \{(id, r, t)\}| - |\mathcal{S}' \setminus \{(id, r, t)\}|$. Hence, $\mathcal{I}_{card}(\mathcal{S}) = \mathcal{I}_{card}(\mathcal{S} \setminus \{(id, r, t)\})$.
2. Penalty. Let $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3, p_4\}$, and $\mathcal{S} = \{(id, \{p_1, p_2\}, t), (id, \{p_2, p_3\}, t), (id, \{p_4\}, t)\}$. Notice that every ST atom in \mathcal{S} is problematic and $\mathcal{I}_{card}(\mathcal{S}) = 1$, as the only card-repair for \mathcal{S} is $\mathcal{S} \setminus \{(id, \{p_4\}, t)\}$. Now $\mathcal{I}_{card}(\mathcal{S} \setminus \{(id, \{p_1, p_2\}, t)\}) = 1$ as there are two card-repairs for it obtained by discarding either of the two ST atoms in \mathcal{S}' . Hence $\mathcal{I}_{card}(\mathcal{S}) \not\geq \mathcal{I}_{card}(\mathcal{S} \setminus \{(id, \{p_1, p_2\}, t)\})$.
3. Dominance. Let \mathcal{S}' be a card-repair for $\mathcal{S} \cup \{(id, r, t)\}$. There are two cases. If $(id, r, t) \notin \mathcal{S}'$, then \mathcal{S}' is also a consistent subset of $\mathcal{S} \cup \{(id, r', t)\}$, but it is possible that $\mathcal{S} \cup \{(id, r', t)\}$ has a card-repair \mathcal{S}'' such that $\mathcal{S}'' = \mathcal{S}' \cup \{(id, r', t)\}$. So, $\mathcal{I}_{card}(\mathcal{S} \cup \{(id, r, t)\}) \geq \mathcal{I}_{card}(\mathcal{S} \cup \{(id, r', t)\})$. If $(id, r, t) \in \mathcal{S}'$, then $\mathcal{S}'' = \mathcal{S}' \setminus \{(id, r, t)\} \cup \{(id, r', t)\}$ is a consistent subset of $\mathcal{S} \cup \{(id, r', t)\}$, and thus every card-repair for $\mathcal{S} \cup \{(id, r', t)\}$ does not delete more tuples than the number deleted to get \mathcal{S}' . Again, $\mathcal{I}_{card}(\mathcal{S} \cup \{(id, r, t)\}) \geq \mathcal{I}_{card}(\mathcal{S} \cup \{(id, r', t)\})$.
4. Super-Additivity. Let $\mathcal{S}'' = \mathcal{S} \cup \mathcal{S}'$ and \mathcal{S}_r'' a card-repair for \mathcal{S}'' . Also, let $\mathcal{S}_r = \mathcal{S}_r'' \cap \mathcal{S}$ and $\mathcal{S}'_r = \mathcal{S}_r'' \cap \mathcal{S}'$. Since $\mathcal{S} \cap \mathcal{S}' = \emptyset$, $|\mathcal{S}_r''| = |\mathcal{S}_r| + |\mathcal{S}'_r|$. Since both \mathcal{S}_r and \mathcal{S}'_r are consistent, every card-repair of \mathcal{S} (resp., \mathcal{S}') has cardinality no less than $|\mathcal{S}_r|$ (resp., $|\mathcal{S}'_r|$). Thus, $\mathcal{I}_{card}(\mathcal{S}) \leq |\mathcal{S}| - |\mathcal{S}_r|$ and $\mathcal{I}_{card}(\mathcal{S}') \leq |\mathcal{S}'| - |\mathcal{S}'_r|$. Hence, $\mathcal{I}_{card}(\mathcal{S}) + \mathcal{I}_{card}(\mathcal{S}') \leq |\mathcal{S}| + |\mathcal{S}'| - |\mathcal{S}_r| - |\mathcal{S}'_r| = |\mathcal{S}''| - |\mathcal{S}_r''| = \mathcal{I}_{card}(\mathcal{S}'')$.
5. Attenuation and Equal Conflict. The fact that Attenuation does not hold and Equal Conflict holds follow from the fact that MI-Normalization holds. See below.
6. MI-Normalization. Let \mathcal{S} be an ST database and $MIS \in \mathcal{MI}(\mathcal{S})$. By the definition of minimal inconsistent subset, every proper subset of MIS is consistent and thus $MIS \setminus \{(id, r, t)\}$ is consistent for every (id, r, t) in MIS . Hence, $\mathcal{I}_{card}(MIS) = 1$.
7. MI-Separability. Let $ID = \{id\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3, p_4\}$, $\mathcal{S} = \{(id, \{p_1\}, t), (id, \{p_2, p_3\}, t)\}$, and $\mathcal{S}' = \{(id, \{p_1\}, t), (id, \{p_3, p_4\}, t)\}$. Then $\mathcal{MI}(\mathcal{S}) = \{\mathcal{S}\}$, $\mathcal{MI}(\mathcal{S}') = \{\mathcal{S}'\}$, and $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \{\mathcal{S}, \mathcal{S}'\}$. Hence $\mathcal{I}_{card}(\mathcal{S} \cup \mathcal{S}') = 1 \neq \mathcal{I}_{card}(\mathcal{S}) + \mathcal{I}_{card}(\mathcal{S}') = 1 + 1 = 2$.

□

Table 2 summarizes the behavior of the repair-based inconsistency measures with respect to the postulates. Observe that the classical postulates fit quite well with the measure \mathcal{I}_{card} ,

	Repair-based IMs			
	\mathcal{I}_{id}	\mathcal{I}_{time}	\mathcal{I}_{region}	\mathcal{I}_{card}
Free-Formula Ind.	✗	✗	✗	✓
Penalty	✗	✗	✗	✗
Dominance	✓	✓	✗	✓
Super-Additivity	✓	✓	✓	✓
Attenuation	✗	✗	✗	✗
Equal Conflict	✗	✗	✗	✓
MI-Normalization	✗	✗	✗	✓
MI-Separability	✗	✗	✗	✗

Table 2: Postulates satisfaction of repair-based inconsistency measures.

which satisfies all of them except Penalty, Attenuation, and MI-Separability. This should not be surprising since the classical postulates are designed for non-dimensional measures like \mathcal{I}_{card} . However, we will show in Section 6 that the measures based on the dimensions satisfy the dimensional versions of some classical postulates they do not satisfy.

5.2.5 RELATIONSHIPS BETWEEN \mathcal{I}_{card} AND THE OTHER MEASURES

In this section, we identify some conditions under which measure \mathcal{I}_{card} coincides with previously considered measures, which in turn implies also equivalences between measures different from \mathcal{I}_{card} .

The following theorem states the equivalence between \mathcal{I}_{card} and \mathcal{I}_{id} , as well as between \mathcal{I}_{card} and \mathcal{I}_{time} , under the restrictions that the distance between different objects (resp., time values) is 1 and there is a sufficient number of object ids (resp., time values).

Theorem 10. *Let \mathcal{S} be an ST database. If $d_{ID}(id, id') = 1$ for all distinct $id, id' \in ID$ and $|ID| \geq |\mathcal{S}|$, then $\mathcal{I}_{card}(\mathcal{S}) = \mathcal{I}_{id}(\mathcal{S})$. Similarly, if $d_T(t, t') = 1$ for all distinct $t, t' \in T$ and $|T| \geq |\mathcal{S}|$, then $\mathcal{I}_{card}(\mathcal{S}) = \mathcal{I}_{time}(\mathcal{S})$.*

Proof. We prove the statement for \mathcal{I}_{id} only as the proof for \mathcal{I}_{time} is similar: it suffices to use time points instead of object identifiers.

Given a card-repair \mathcal{S}' for \mathcal{S} with $cost_{card}(\mathcal{S}, \mathcal{S}') = |\mathcal{S}| - |\mathcal{S}'|$, we define an id-repair \mathcal{S}^* for \mathcal{S} as follows. Let $ID(\mathcal{S}) = \{id_1, \dots, id_n\}$ be the set of distinct ids in \mathcal{S} (clearly, $n \leq |\mathcal{S}|$). For each ST atom $a_i \in \mathcal{S} \setminus \mathcal{S}'$, let id_{n+i} be a fresh id in $ID \setminus ID(\mathcal{S})$. Let \mathcal{S}^* be the id-correction $\mathcal{S}' \cup \{(id_{n+i}, r, t) \mid a_i \in \mathcal{S} \setminus \mathcal{S}'\}$. Observe that, since $|ID| \geq |\mathcal{S}|$, there is a sufficient number of ids in ID so that each ST atom in $\mathcal{S} \setminus \mathcal{S}'$ can be updated to a different id. Clearly, $cost_{id}(\mathcal{S}, \mathcal{S}') = cost_{card}(\mathcal{S}, \mathcal{S}')$ as $d_{ID}(id, id') = 1$ for all distinct $id, id' \in ID$. Moreover, \mathcal{S}^* is a minimal id-correction as changing the id of any ST atom $(id_x, r, t) \in (\mathcal{S} \setminus \mathcal{S}')$ to any other id in $ID(\mathcal{S})$ would still have cost equal to 1 and cannot resolve any inconsistency other than that involving (id_x, r, t) , thus the number of id-updates needed to obtain \mathcal{S}^* is minimum. Therefore, \mathcal{S}^* is an id-repair for \mathcal{S} whose cost is equal to the cost of a card-repair, from which it follows that $\mathcal{I}_{card}(\mathcal{S}) = \mathcal{I}_{id}(\mathcal{S})$. \square

The following corollary follows.

Corollary 1. *Let \mathcal{S} be an ST database. If $d_{ID}(id, id') = 1$ for all distinct $id, id' \in ID$, $|ID| \geq |\mathcal{S}|$, $d_T(t, t') = 1$ for all distinct $t, t' \in T$, and $|T| \geq |\mathcal{S}|$, then $\mathcal{I}_{card}(\mathcal{S}) = \mathcal{I}_{id}(\mathcal{S}) = \mathcal{I}_{time}(\mathcal{S})$.*

The previous results will be particularly useful in Section 7, where we characterize the complexity of inconsistency measures. In fact, the results above will allow us to single out some tractable cases for \mathcal{I}_{id} and \mathcal{I}_{time} , which are NP-complete in the general case. More specifically, we will prove that \mathcal{I}_{card} has polynomial time complexity and, since \mathcal{I}_{card} coincides with \mathcal{I}_{id} and \mathcal{I}_{time} under the aforementioned conditions, this in turn allows us to get tractability for \mathcal{I}_{id} and \mathcal{I}_{time} under those conditions.

Finally, it turns out that the value of \mathcal{I}_{card} is equal to that of \mathcal{I}_{region} if the metric is the symmetric difference.

Theorem 11. *Given an ST databases \mathcal{S} , if $d_R(r, r') = |(r \setminus r') \cup (r' \setminus r)|$ for $r, r' \subseteq Space$, then $\mathcal{I}_{card}(\mathcal{S}) = \mathcal{I}_{region}(\mathcal{S})$.*

Proof. Given a card-repair \mathcal{S}' for \mathcal{S} with cost $cost_{card}(\mathcal{S}, \mathcal{S}') = |\mathcal{S}| - |\mathcal{S}'|$, we define a region-repair \mathcal{S}^* for \mathcal{S} as follows. Let \mathcal{S}^* be the region-correction defined as follows. For each id, t pair in \mathcal{S}' , let $p^{id,t}$ be a point in $Space$ such that $p^{id,t}$ belongs to the intersection of the regions of all the ST atoms in $\mathcal{S}'_{id,t}$ (clearly such a point exists as $\mathcal{S}'_{id,t}$ is consistent). Let \mathcal{S}^* be $\mathcal{S}' \cup \{(id, r \cup \{p^{id,t}\}), t) \mid (id, r, t) \in \mathcal{S} \setminus \mathcal{S}'\}$. Clearly, $cost_{region}(\mathcal{S}, \mathcal{S}') = |\mathcal{S}| - |\mathcal{S}'|$ since d_R is symmetric difference, and thus the addition of a point to a region has cost equal to 1. Moreover, as every ST atom in $\mathcal{S} \setminus \mathcal{S}'$ needs to be updated to restore consistency, this is the minimum cost that can be achieved. Therefore, \mathcal{S}^* is a region-repair for \mathcal{S} whose cost is equal to that of a card-repair, from which it follows that $\mathcal{I}_{card}(\mathcal{S}) = \mathcal{I}_{region}(\mathcal{S})$. \square

Besides its own interest, the previous theorem will be particularly valuable to characterize the complexity of \mathcal{I}_{card} in Section 7. In fact, we will show that $\mathcal{I}_{regcorr}$ has polynomial time complexity when the metric for regions is the symmetric difference, which, combined with Theorem 11 above, allows us to claim that the complexity of \mathcal{I}_{card} is polynomial time too.

6. Dimensional Postulates

In Section 3.2, we presented eight rationality postulates that are properties proposed as desirable in the literature for inconsistency measures. These postulates were formulated for propositional knowledge bases. In our case we deal with a database where the data has three dimensions: objects, time, and space. There seems to be a mismatch in some cases between the postulates that have nothing to do with dimensions and our inconsistency measures that are based on dimensions. In this section, we propose dimensional versions for some of the postulates and find that in several cases the dimensional postulate holds even though the original postulate does not hold.

We start with the fact that the atoms under consideration are tuples in n dimensions. For spatio-temporal databases there are three dimensions, but in other contexts the number of dimensions may be different. In some dimensions (objects and time in our case) the value is an element while in other dimensions (space in our case) the value is a set, and there is a consistency requirement (for every pair of id and t values the space regions must have a

nonempty intersection in our case). As we deal only with spatio-temporal databases, we will restrict our consideration to this case.

Next we consider the postulates we have been dealing with and explain how to define dimensional versions for them. We notice that for some of them there are no dimensional versions. Attenuation and Equal Conflict simply deal with the size of minimal inconsistent subsets. MI-Normalization sets the inconsistency measure of a minimal inconsistent set to 1. Free-Formula Independence and Dominance also do not appear to have dimensional versions. Hence, we restrict our attention to the three remaining postulates: Penalty, Super-Additivity, and MI-Separability. These postulates have been formulated in Definition 6 as a direct translation of their original definition for propositional knowledge bases to ST databases. Roughly speaking, while for propositional knowledge bases we look at propositional formulas, for ST databases we look at ST atoms. However, looking at whole ST atoms can be a too coarse-grained approach, since each ST atom is indeed characterized by an object id, a time value, and a point location. More fine-grained formulations of the postulates can be achieved by focusing on the object, time, and space dimensions individually, which is the main idea of the dimensional postulates introduced in the following. Thus, for each postulate, we introduce three variants, with each of them looking at a specific dimension (as opposed to looking at entire ST atoms).

We start with the dimensional version of the Penalty postulate. First of all, we recall the Penalty postulate as formulated in Definition 6: If $(id, r, t) \in \text{Problematic}(\mathcal{S})$ then $\mathcal{I}(\mathcal{S}) > \mathcal{I}(\mathcal{S} \setminus \{(id, r, t)\})$. Thus, Penalty states that problematic ST atoms increase the inconsistency (or, from another point of view, deleting a problematic ST atom from an ST database decreases its inconsistency). As already mentioned above, this was a coarse-grained, direct translation of the Penalty postulate for propositional knowledge bases to ST databases. Indeed, such a postulate is not satisfied by any of the inconsistency measures we introduced other than \mathcal{I}_D (cf. Tables 1 and 2).

A natural “dimensional” version of the Penalty postulate would focus on the object (resp., time, space) dimension and say that the deletion of a problematic *object* (resp., *time value*, *region*) decreases the inconsistency of an ST database. This leads us to the following definition.

Definition 17 (Dimensional Penalty). *Let \mathcal{I} be an inconsistency measure and \mathcal{S} be an ST database.*

1. (Object Penalty) *If $(id, r, t) \in \text{Problematic}(\mathcal{S})$ and $A = \{(id, r', t') \in \text{Problematic}(\mathcal{S})\}$ then $\mathcal{I}(\mathcal{S}) > \mathcal{I}(\mathcal{S} \setminus A)$.*
2. (Time Penalty) *If $(id, r, t) \in \text{Problematic}(\mathcal{S})$ and $A = \{(id', r', t) \in \text{Problematic}(\mathcal{S})\}$ then $\mathcal{I}(\mathcal{S}) > \mathcal{I}(\mathcal{S} \setminus A)$.*
3. (Space Penalty) *If $(id, r, t) \in \text{Problematic}(\mathcal{S})$ and $A = \{(id', r', t') \in \text{Problematic}(\mathcal{S}) \mid r \cap r' \neq \emptyset\}$ then $\mathcal{I}(\mathcal{S}) > \mathcal{I}(\mathcal{S} \setminus A)$.*

We now turn our attention to the Super-Additivity postulate, which was formulated in Definition 6 as follows: If $\mathcal{S} \cap \mathcal{S}' = \emptyset$ then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$. The rationale behind this postulate is that if two ST databases do not share any ST atoms, then the inconsistency of their union is at least as great as the sum of the inconsistency of the two

ST databases. However, once again, we are reasoning at the level of ST atoms, which can be too coarse-grained, given the different kinds of information encoded in ST atoms. As for the Penalty postulate above, dimensional versions of the Super-Additivity postulate can be achieved by focusing on the object, time, and space dimensions individually. A natural formulation is thus as follows: if two ST databases do not share any *object* (resp., *time value*, *point location*), then the inconsistency of their union is at least as great as the sum of the inconsistency of the two ST databases, which is formalized in the definition below.

In the following, for an ST database \mathcal{S} , we use the following notations:

$$\begin{aligned} ID(\mathcal{S}) &= \{id \mid (id, r, t) \in \mathcal{S}\}, \\ Time(\mathcal{S}) &= \{t \mid (id, r, t) \in \mathcal{S}\}, \\ Region(\mathcal{S}) &= \bigcup \{r \mid (id, r, t) \in \mathcal{S}\}. \end{aligned}$$

Definition 18 (Dimensional Super-Additivity). *Let \mathcal{I} be an inconsistency measure, and $\mathcal{S}, \mathcal{S}'$ be two ST databases. For $X \in \{ID, Time, Region\}$, if $X(\mathcal{S}) \cap X(\mathcal{S}') = \emptyset$ then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$.*

Finally, we get to MI-Separability, which was formulated in Definition 6 as follows: If $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$ and $\mathcal{MI}(\mathcal{S}) \cap \mathcal{MI}(\mathcal{S}') = \emptyset$ then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$. For this postulate, we can reason analogously to the Penalty and Super-Additivity postulates, that is, provide counterparts of the postulate that focus on each dimension individually. Here we find that because of the dimensional requirement the condition $\mathcal{MI}(\mathcal{S}) \cap \mathcal{MI}(\mathcal{S}') = \emptyset$ is implied by $X(\mathcal{S}) \cap X(\mathcal{S}') = \emptyset$ for $X \in \{ID, Time, Region\}$, thus we omit $\mathcal{MI}(\mathcal{S}) \cap \mathcal{MI}(\mathcal{S}') = \emptyset$ in the dimensional version of the MI-Separability postulate defined below.

Definition 19 (Dimensional MI-Separability). *Let \mathcal{I} be an inconsistency measure, and $\mathcal{S}, \mathcal{S}'$ be two ST databases. For $X \in \{ID, Time, Region\}$, if $X(\mathcal{S}) \cap X(\mathcal{S}') = \emptyset$ and $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$, then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$.*

As stated next, for the object and time dimensions, Dimensional MI-Separability implies Dimensional Super-Additivity.

Proposition 3. *If an inconsistency measure satisfies Object (resp., Time) Dimensional MI-Separability then it satisfies Object (resp., Time) Dimensional Super-Additivity.*

Proof. The claim follows from the fact that, given two ST databases \mathcal{S} and \mathcal{S}' , for $X \in \{ID, Time\}$, it is the case that $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$ whenever $X(\mathcal{S}) \cap X(\mathcal{S}') = \emptyset$. In fact, if \mathcal{S} and \mathcal{S}' do not share any object (resp., time value), then $\mathcal{S} \cap \mathcal{S}' = \emptyset$, from which it follows that $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$. Therefore, Object (resp., Time) Dimensional MI-Separability can be rewritten as follows: For $X \in \{ID, Time\}$, if $X(\mathcal{S}) \cap X(\mathcal{S}') = \emptyset$ then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$, from which the statement follows. \square

In contrast, Space Dimensional MI-Separability does not imply Space Dimensional Super-Additivity. To prove this claim, we need to show an *arbitrary* inconsistency measure that satisfies Space Dimensional MI-Separability but does not satisfy Space Dimensional Super-Additivity. Below we show that this is the case for the \mathcal{I}_{CC} measure defined for propositional knowledge bases (Thimm, 2018).

A set $\{\mathcal{S}_1, \dots, \mathcal{S}_n\}$ of pairwise disjoint subsets of \mathcal{S} is called a *conditional independent Minimal Unsatisfiable Subset (MUS) partition* of \mathcal{S} iff each \mathcal{S}_i is inconsistent and $\mathcal{MI}(\mathcal{S}_1 \cup \dots \cup \mathcal{S}_n)$ is the disjoint union of all $\mathcal{MI}(\mathcal{S}_i)$'s, that is, $\mathcal{MI}(\mathcal{S}_1 \cup \dots \cup \mathcal{S}_n)$ is the union of all $\mathcal{MI}(\mathcal{S}_i)$'s and the latter are pairwise disjoint. The \mathcal{I}_{CC} inconsistency measure is defined as

$$\mathcal{I}_{CC}(\mathcal{S}) = \max\{n \mid \{\mathcal{S}_1, \dots, \mathcal{S}_n\} \text{ is a conditional independent MUS partition of } \mathcal{S}\},$$

which is an adaptation of the \mathcal{I}_{CC} measure defined for propositional knowledge bases by Thimm (2018) to the case of *ST* databases.

To see why \mathcal{I}_{CC} satisfies Space Dimensional MI-Separability, observe that, given two *ST* databases \mathcal{S} and \mathcal{S}' , if $Region(\mathcal{S}) \cap Region(\mathcal{S}') = \emptyset$ and $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$, and thus as observed earlier $\mathcal{MI}(\mathcal{S}) \cap \mathcal{MI}(\mathcal{S}') = \emptyset$, then every maximal conditional independent MUS partition $\{\mathcal{S}_1, \dots, \mathcal{S}_n\}$ of $\mathcal{S} \cup \mathcal{S}'$ is such that it can be partitioned into two disjoint maximal conditional independent MUS partitions, one for \mathcal{S} and another for \mathcal{S}' , and thus $\mathcal{I}_{CC}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}_{CC}(\mathcal{S}) + \mathcal{I}_{CC}(\mathcal{S}')$.

To show that \mathcal{I}_{CC} does not satisfy Space Dimensional Super-Additivity, it suffices to consider the *ST* databases $\mathcal{S} = \{(id, \{p_1\}, t), (id, \{p_2\}, t)\}$ and $\mathcal{S}' = \{(id, \{p_3\}, t), (id, \{p_4\}, t)\}$, where all p_i 's are distinct. Although $Region(\mathcal{S}) = \{p_1, p_2\}$ does not intersect $Region(\mathcal{S}') = \{p_3, p_4\}$, it is the case that $\mathcal{I}_{CC}(\mathcal{S} \cup \mathcal{S}') = 1 \not\geq 2 = \mathcal{I}_{CC}(\mathcal{S}) + \mathcal{I}_{CC}(\mathcal{S}')$. Notice that $\mathcal{I}_{CC}(\mathcal{S}) = 1$, as the largest conditional independent MUS partition of \mathcal{S} is $\{\mathcal{S}\}$, and $\mathcal{I}_{CC}(\mathcal{S}') = 1$, as the largest conditional independent MUS partition of \mathcal{S}' is $\{\mathcal{S}'\}$. Also, $\mathcal{I}_{CC}(\mathcal{S} \cup \mathcal{S}') = 1$, as the largest conditional independent MUS partitions of $\mathcal{S} \cup \mathcal{S}'$ are of the form $\{(id, p_i, t), (id, p_j, t)\}$ with p_i, p_j distinct and belonging to $\{p_1, p_2, p_3, p_4\}$. Notice that there is no conditional independent MUS partition of $\mathcal{S} \cup \mathcal{S}'$ of cardinality 2. The reason is that for such a partition $\{\mathcal{S}_1, \mathcal{S}_2\}$, it has to be the case that $|\mathcal{S}_1| = |\mathcal{S}_2| = 2$ in order for \mathcal{S}_1 and \mathcal{S}_2 to be inconsistent and disjoint. Then, $\mathcal{MI}(\mathcal{S}_1 \cup \mathcal{S}_2)$ has six minimal inconsistent subsets and thus it is not the disjoint union of $\mathcal{MI}(\mathcal{S}_1)$ and $\mathcal{MI}(\mathcal{S}_2)$, as $|\mathcal{MI}(\mathcal{S}_1)| = |\mathcal{MI}(\mathcal{S}_2)| = 1$. Thus, $\{\mathcal{S}_1, \mathcal{S}_2\}$ cannot be a conditional independent MUS partition of $\mathcal{S} \cup \mathcal{S}'$.

Now we are ready to consider the dimensional postulates for the inconsistency measures we have dealt with. We do not include \mathcal{I}_{OT} as it deals with two dimensions and we restricted the dimensional postulates to a single dimension.

In the following, we first focus on the dimensional inconsistency measures introduced in Section 4 (namely \mathcal{I}_O , \mathcal{I}_T , \mathcal{I}_S , and \mathcal{I}_D), and then focus on the repair-based inconsistency measures introduced in Section 5 (namely \mathcal{I}_{id} , \mathcal{I}_{time} , \mathcal{I}_{region}).

The following four propositions show which dimensional postulates are satisfied by \mathcal{I}_O , \mathcal{I}_T , \mathcal{I}_S , and \mathcal{I}_D , respectively—indeed, all of them satisfy all the dimensional postulates.

We start by considering the object dimension and show that \mathcal{I}_O satisfies all the dimensional postulates in the following proposition.

Proposition 4. *The \mathcal{I}_O inconsistency measure satisfies Object Penalty, Object Super-Additivity, and Object MI-Separability.*

Proof. 1. Object Penalty. $\mathcal{I}_O(\mathcal{S} \setminus A) = \mathcal{I}_O(\mathcal{S}) - 1 < \mathcal{I}_O(\mathcal{S})$.

2. Object Super-Additivity. This follows from Proposition 3.

3. Object MI-Separability. $\mathcal{I}_O(\mathcal{S})$ and $\mathcal{I}_O(\mathcal{S}')$ count different objects. As there is no interaction between \mathcal{S} and \mathcal{S}' , $\mathcal{I}_O(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}_O(\mathcal{S}) + \mathcal{I}_O(\mathcal{S}')$. \square

We now focus on the time dimension and show that \mathcal{I}_T satisfies all the dimensional postulates.

Proposition 5. *The \mathcal{I}_T inconsistency measure satisfies Time Penalty, Time Super-Additivity, and Time MI-Separability.*

- Proof.* 1. Time Penalty. $\mathcal{I}_T(\mathcal{S} \setminus A) = \mathcal{I}_T(\mathcal{S}) - 1 < \mathcal{I}_T(\mathcal{S})$.
2. Time Super-Additivity. This follows from Proposition 3.
3. Time MI-Separability. $\mathcal{I}_T(\mathcal{S})$ and $\mathcal{I}_T(\mathcal{S}')$ count different time values. As there is no interaction between \mathcal{S} and \mathcal{S}' , $\mathcal{I}_T(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}_T(\mathcal{S}) + \mathcal{I}_T(\mathcal{S}')$. \square

Finally, we consider the space dimension, that is, we focus on the inconsistency measures \mathcal{I}_S (in Proposition 6) and \mathcal{I}_D (in Proposition 7) showing that both satisfy all the dimensional postulates.

Proposition 6. *The \mathcal{I}_S inconsistency measure satisfies Space Penalty, Space Super-Additivity, and Space MI-Separability.*

Proof. First of all, for an ST database \mathcal{S} we define a region $R_{\mathcal{S}}$ as follows:

$$R_{\mathcal{S}} = \bigcup \{r \mid (id, r, t) \in MIS \in \mathcal{MI}(\mathcal{S})\}$$

and $\mathcal{I}_S(\mathcal{S}) = |R_{\mathcal{S}}|$.

1. Space Penalty. $\mathcal{I}_S(\mathcal{S} \setminus A) \leq \mathcal{I}_S(\mathcal{S}) - |r| < \mathcal{I}_S(\mathcal{S})$.
2. Space Super-Additivity. By definition, $\mathcal{I}_S(\mathcal{S}) = |R_{\mathcal{S}}|$ and $\mathcal{I}_S(\mathcal{S}') = |R_{\mathcal{S}'}|$. Then $R_{\mathcal{S}} \cup R_{\mathcal{S}'} \subseteq R_{\mathcal{S} \cup \mathcal{S}'}$ and as $R_{\mathcal{S}}$ and $R_{\mathcal{S}'}$ are disjoint, $\mathcal{I}_S(\mathcal{S}) + \mathcal{I}_S(\mathcal{S}') = |R_{\mathcal{S}}| + |R_{\mathcal{S}'}| \leq |R_{\mathcal{S} \cup \mathcal{S}'}| = \mathcal{I}_S(\mathcal{S} \cup \mathcal{S}')$.
3. Space MI-Separability. Continuing with the proof of Space Super-Additivity, the additional condition assures that all minimal inconsistent sets in $\mathcal{S} \cup \mathcal{S}'$ must be either in \mathcal{S} or \mathcal{S}' . Hence $R_{\mathcal{S}} \cup R_{\mathcal{S}'} = R_{\mathcal{S} \cup \mathcal{S}'}$ and the result follows. \square

In the next proposition we show the relationship between the dimensional postulates and the inconsistency measure \mathcal{I}_D .

Proposition 7. *The \mathcal{I}_D inconsistency measure satisfies Space Penalty, Space Super-Additivity, and Space MI-Separability.*

- Proof.* 1. Space Penalty. $\mathcal{MI}(\mathcal{S} \setminus A) \subset \mathcal{MI}(\mathcal{S})$, hence $\mathcal{I}_D(\mathcal{S} \setminus A) < \mathcal{I}_D(\mathcal{S})$.

2. Space Super-Additivity. $\mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}') \subseteq \mathcal{MI}(\mathcal{S} \cup \mathcal{S}')$, hence $\mathcal{I}_D(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}_D(\mathcal{S}) + \mathcal{I}_D(\mathcal{S}')$.
3. Space MI-Separability. $\mathcal{MI}(\mathcal{S} \cup \mathcal{S}') = \mathcal{MI}(\mathcal{S}) \cup \mathcal{MI}(\mathcal{S}')$, hence $\mathcal{I}_D(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}_D(\mathcal{S}) + \mathcal{I}_D(\mathcal{S}')$.

□

We now consider the dimensional repair-based inconsistency measures introduced in Section 5, namely \mathcal{I}_{id} , \mathcal{I}_{time} , and \mathcal{I}_{region} . We show that all of them satisfy the Penalty postulate, even though \mathcal{I}_{id} and \mathcal{I}_{time} do so provided that they return finite values—recall that \mathcal{I}_{id} and \mathcal{I}_{time} return ∞ when there is no id- and time-repair, respectively. All the three inconsistency measures satisfy Super-Additivity, while only \mathcal{I}_{region} satisfies MI-Separability.

Proposition 8. *The \mathcal{I}_{id} inconsistency measure satisfies Object Penalty when the measure returns a finite value and Object Super-Additivity, but does not satisfy Object MI-Separability.*

Proof. 1. Object Penalty. Assume that $\mathcal{I}_{id}(\mathcal{S}) < \infty$. As $(id, r, t) \in \text{Problematic}(\mathcal{S})$, any id-repair must update at least one atom, say a , whose object id is id to an atom a' . Then $\mathcal{I}_{id}(\mathcal{S} \setminus A) \leq \mathcal{I}_{id}(\mathcal{S}) - \text{cost}_{id}(a, a') < \mathcal{I}_{id}(\mathcal{S})$.

2. Object Super-Additivity. If either \mathcal{S} or \mathcal{S}' does not have an id-repair, then $\mathcal{S} \cup \mathcal{S}'$ cannot have one either. In this case $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') = \infty$ and the result holds. So assume that both \mathcal{S} and \mathcal{S}' have an id-repair, say \mathcal{S}_1 and \mathcal{S}'_1 respectively. If $\mathcal{S} \cup \mathcal{S}'$ does not have an id-repair, then $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') = \infty$ and the result holds. Otherwise, any id-repair for $\mathcal{S} \cup \mathcal{S}'$ must do an id-correction for each atom that needed one in \mathcal{S} and \mathcal{S}' . Such an id-correction cannot have smaller cost than the corresponding one for \mathcal{S} and \mathcal{S}' . Hence the cost cannot decrease. Thus in all cases, $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}_{id}(\mathcal{S}) + \mathcal{I}_{id}(\mathcal{S}')$.

3. Object MI-Separability. Let $ID = \{id_1, id_2, id_3\}$, $T = \{t\}$, $Space = \{p_1, p_2, p_3\}$, $\mathcal{S} = \{(id_1, \{p_1\}, t), (id_1, \{p_2\}, t), (id_3, \{p_1\}, t)\}$, and also let $\mathcal{S}' = \{(id_2, \{p_3\}, t)\}$ where $d_{ID}(id_i, id_j) = |i - j|$. Then $\mathcal{I}_{id}(\mathcal{S}) = 1$ as an id-repair of \mathcal{S} must correct the id_1 in either the first or the second atom to id_2 and $\mathcal{I}_{id}(\mathcal{S}') = 0$ (\mathcal{S}' is consistent). But for $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}')$ an id-repair must change id_1 in the first atom to id_3 . Hence $\mathcal{I}_{id}(\mathcal{S} \cup \mathcal{S}') = 2 \neq \mathcal{I}_{id}(\mathcal{S}) + \mathcal{I}_{id}(\mathcal{S}') = 1 + 0 = 1$.

□

An analogous result is shown for \mathcal{I}_{time} in the next proposition.

Proposition 9. *The \mathcal{I}_{time} inconsistency measure satisfies Time Penalty when the measure returns a finite value and Time Super-Additivity, but does not satisfy Time MI-Separability.*

Proof. 1. Time Penalty. Assume that $\mathcal{I}_{time}(\mathcal{S}) < \infty$. As $(id, r, t) \in \text{Problematic}(\mathcal{S})$, any time-repair must update at least one atom, say a , whose time value is t to an atom a' . Then $\mathcal{I}_{time}(\mathcal{S} \setminus A) \leq \mathcal{I}_{time}(\mathcal{S}) - \text{cost}_{time}(a, a') < \mathcal{I}_{time}(\mathcal{S})$.

2. Time Super-Additivity. If either \mathcal{S} or \mathcal{S}' does not have a time-repair, then $\mathcal{S} \cup \mathcal{S}'$ cannot have one either. In this case $\mathcal{I}_{time}(\mathcal{S} \cup \mathcal{S}') = \infty$ and the result holds. So assume that both \mathcal{S} and \mathcal{S}' have a time-repair, say \mathcal{S}_2 and \mathcal{S}'_2 respectively. If $\mathcal{S} \cup \mathcal{S}'$ does not have a

time-repair, then $\mathcal{I}_{time}(\mathcal{S} \cup \mathcal{S}') = \infty$ and the result holds. Otherwise, any time-repair for $\mathcal{S} \cup \mathcal{S}'$ must do a time-correction for each atom that needed one in \mathcal{S} and \mathcal{S}' . Such a time-correction cannot have smaller cost than the corresponding one for \mathcal{S} and \mathcal{S}' . Hence the cost cannot decrease. Thus in all cases, $\mathcal{I}_{time}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}_{time}(\mathcal{S}) + \mathcal{I}_{time}(\mathcal{S}')$.

3. Time MI-Separability. Let $ID = \{id\}$, $T = \{1, 2, 3\}$, $Space = \{p_1, p_2, p_3\}$, $\mathcal{S} = \{(id, \{p_1\}, 1), (id, \{p_2\}, 1), (id, \{p_1\}, 3)\}$, and $\mathcal{S}' = \{(id, \{p_3\}, 2)\}$ where $d_T(i, j) = |i - j|$. Then $\mathcal{I}_{time}(\mathcal{S}) = 1$ as a time-repair of \mathcal{S} must correct the 1 in either the first or the second atom to 2 and $\mathcal{I}_{time}(\mathcal{S}') = 0$ (\mathcal{S}' is consistent). But for $\mathcal{I}_{time}(\mathcal{S} \cup \mathcal{S}')$ a time-repair must change 1 in the first atom to 3. Hence $\mathcal{I}_{time}(\mathcal{S} \cup \mathcal{S}') = 2 \neq \mathcal{I}_{time}(\mathcal{S}) + \mathcal{I}_{time}(\mathcal{S}') = 1 + 0 = 1$.

□

In contrast with the two previous results, Proposition 10 shows that all the dimensional postulates are satisfied by \mathcal{I}_{region} without further assumptions or conditions.

Proposition 10. *The \mathcal{I}_{region} inconsistency measure satisfies Space Penalty, Space Super-Additivity, and Space MI-Separability.*

Proof. 1. Space Penalty. Let $a = (id, r, t) \in Problematic(\mathcal{S})$ and suppose that a region-repair, among other changes, corrects a . Then the removal of a removes from $\mathcal{I}_{region}(\mathcal{S})$ the cost of such a change, and hence $\mathcal{I}_{region}(\mathcal{S} \setminus A) < \mathcal{I}_{region}(\mathcal{S})$. On the other hand, suppose that a region-repair does not correct a . Then it must correct some other problematic atom, say $a' = (id, r', t)$ with the same id and t values. So again, the value of \mathcal{I}_{region} is lessened giving $\mathcal{I}_{region}(\mathcal{S} \setminus A) < \mathcal{I}_{region}(\mathcal{S})$.

2. Space Super-Additivity. Every region-repair of $\mathcal{S} \cup \mathcal{S}'$ must repair both \mathcal{S} and \mathcal{S}' separately in the process. In fact, a region-repair of $\mathcal{S} \cup \mathcal{S}'$ may require additional region updates. In any case, under the condition of separate combined regions for \mathcal{S} and \mathcal{S}' , $\mathcal{I}_{region}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}_{region}(\mathcal{S}) + \mathcal{I}_{region}(\mathcal{S}')$.
3. Space MI-Separability. Continuing with the proof of Space Super-Additivity consider the additional condition about minimal inconsistent subsets. Combining the two conditions means that an atom in \mathcal{S} cannot have the same id and t values as an atom in \mathcal{S}' . Therefore a region-repair of $\mathcal{S} \cup \mathcal{S}'$ must be a region-correction whose restriction to \mathcal{S} (resp., \mathcal{S}') is a region-correction for \mathcal{S} (resp., \mathcal{S}'). Hence $\mathcal{I}_{region}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}_{region}(\mathcal{S}) + \mathcal{I}_{region}(\mathcal{S}')$.

□

Dimensional postulates satisfaction for the inconsistency measures are summarized in Table 3, which suggests that the dimensional postulates suit very well the dimensional inconsistency measures. Table 3 shows that the dimensional postulates are jointly satisfiable.

On the whole, Tables 1, 2, and 3 show that, considering the dimensional postulates in place of the classical ones, \mathcal{I}_O and \mathcal{I}_T satisfy all the postulates except Attenuation, that is, \mathcal{I}_O and \mathcal{I}_T satisfy the maximum number of postulates considered, since Attenuation is incompatible with MI-Normalization, as observed after Definition 6. In general, the number of postulates satisfied by considering dimensional and classical postulates has doubled for all

	\mathcal{I}_O	\mathcal{I}_T	\mathcal{I}_S	\mathcal{I}_D	\mathcal{I}_{id}	\mathcal{I}_{time}	\mathcal{I}_{region}
Dimensional Penalty	✓	✓	✓	✓	✓*	✓*	✓
Dimensional Super-Additivity	✓	✓	✓	✓	✓	✓	✓
Dimensional MI-Separability	✓	✓	✓	✓	✗	✗	✓

Table 3: Dimensional postulates satisfaction (✓* means satisfied when measures return a finite value).

the measures except for \mathcal{I}_D (which satisfies both the classical and the dimensional version of the postulates in Table 3).

As for the repair-based measures \mathcal{I}_{id} , \mathcal{I}_{time} , and \mathcal{I}_{region} , under the assumptions of Theorem 10, for which \mathcal{I}_{card} coincides with \mathcal{I}_{id} and \mathcal{I}_{time} , considering dimensional and classical postulates, \mathcal{I}_{id} and \mathcal{I}_{time} satisfy all the postulates except (Dimensional) MI-Separability and Attenuation (which is incompatible with MI-Normalization). This follows from the fact that, in addition to the postulates satisfied by \mathcal{I}_{card} , also Dimensional Penalty is satisfied in this case. Analogously, under the assumptions of Theorem 11, for which \mathcal{I}_{card} coincides with \mathcal{I}_{region} , considering dimensional and classical postulates, \mathcal{I}_{region} satisfies all the postulates except Attenuation, that is, the maximum number of postulates that can be jointly satisfied, as it happens also for \mathcal{I}_O and \mathcal{I}_T .

7. Complexity Analysis

In this section, we provide complexity results for all the inconsistency measures we introduced.

Following Thimm and Wallner (2016) and Thimm (2018), who have explored the complexity of problems related to computing classical inconsistency measures, we characterize the complexity of the following decision problem, named UPPER by Thimm (2018), for an inconsistency measure \mathcal{I} : given an ST database \mathcal{S} and rational number k , decide whether $\mathcal{I}(\mathcal{S}) \leq k$. Let \mathcal{C} be a complexity class. In the following, we will say that UPPER for \mathcal{I} is in \mathcal{C} if deciding whether $\mathcal{I}(\mathcal{S}) \leq k$ is in \mathcal{C} . We say that UPPER for \mathcal{I} is \mathcal{C} -hard if deciding whether $\mathcal{I}(\mathcal{S}) \leq k$ is \mathcal{C} -hard. If deciding whether $\mathcal{I}(\mathcal{S}) \leq k$ is in \mathcal{C} and it is \mathcal{C} -hard, we say that UPPER for \mathcal{I} is \mathcal{C} -complete. We assume the reader is familiar with standard complexity classes such as P and NP (Papadimitriou, 1994).

We start with the dimensional measures introduced in Section 4. It turns out that these measures are generally tractable in contrast to the measures for propositional knowledge bases considered by Thimm (2018). Then, we consider the repair-based measures introduced in Section 5, which in general have higher complexity.

7.1 Complexity Results for Dimensional Measures

We start with an algorithm that computes the \mathcal{I}_{OT} measure, which runs in polynomial time.

Proposition 11. *UPPER for the dimensional inconsistency measure \mathcal{I}_{OT} is in P.*

Proof. The calculation of $\mathcal{I}_{OT}(\mathcal{S})$ is accomplished by Algorithm 1, which is in P. It proceeds as follows. We initialize the inconsistency measure counter to 0 (line 1). For each id, t pair

Algorithm 1 Compute the \mathcal{I}_{OT} measure

Input: An ST database \mathcal{S} .

Output: $\mathcal{I}_{OT}(\mathcal{S})$.

```

1: let  $\mathcal{I}_{OT}(\mathcal{S}) = 0$ 
2: for each  $(id, r, t)$  in  $\mathcal{S}$  do
3:     let  $\mathcal{P}_{id,t} = \emptyset$ 
4: end for
5: for each  $(id, r, t)$  in  $\mathcal{S}$  do
6:     if  $\mathcal{P}_{id,t} \neq Space$  then
7:         let  $\mathcal{P}_{id,t} = \mathcal{P}_{id,t} \cup (Space \setminus r)$ 
8:         if  $\mathcal{P}_{id,t} = Space$  then
9:             let  $\mathcal{I}_{OT}(\mathcal{S}) = \mathcal{I}_{OT}(\mathcal{S}) + 1$ 
10:        end if
11:    end if
12: end for
13: return  $\mathcal{I}_{OT}(\mathcal{S})$ 
    
```

we use $\mathcal{P}_{id,t}$ to indicate the amount of space *not in the intersection* of all the regions for the id, t pair. If $\mathcal{P}_{id,t}$ becomes equal to $Space$ then the intersection of all the regions for that pair is empty, signifying an inconsistent id, t pair. The first loop (lines 2–4) traverses all the ST atoms in \mathcal{S} and initializes the amount of space not in the intersections of all the regions for all id, t pairs to the empty set. The next loop (lines 5–12) again traverses all the ST atoms and for each such atom (id, r, t) it first (line 6) checks to see if the amount of space already not in the intersection of all the regions for the id, t pair is not all of $Space$ —we assume this check can be accomplished in constant time. If it is all of $Space$ then we have already counted the id, t pair as inconsistent at a previous step. Otherwise, we add (line 7) the complement of r to the amount of space not in the intersection of all the regions for the id, t pair. If the result is all of $Space$ (line 8) then we have a new inconsistency for the id, t pair and add 1 to our inconsistency measure counter (line 9). At the end (line 13), the inconsistency measure counter, $\mathcal{I}_{OT}(\mathcal{S})$, is set correctly. Assuming that the basic union operation in line 7 can be done in constant time, we have that the running time of the algorithm is $O(|\mathcal{S}|)$; however, in the worst case this operation could take time $O(|Space|)$. Therefore, the general result is that the running time of the algorithm is $O(|\mathcal{S}| * C_{\cup})$, where C_{\cup} is the cost of computing the union of two regions. Notice that in the worst-case the algorithm has to keep a region $\mathcal{P}_{id,t}$ for each atom in \mathcal{S} (this case occurs when all atoms of \mathcal{S} refer to different id, t pairs), which requires polynomial space. \square

A small modification of this proposition results in obtaining the same complexity for the \mathcal{I}_O and \mathcal{I}_T measures, as stated in the following two propositions.

Proposition 12. UPPER for the dimensional inconsistency measure \mathcal{I}_O is in P.

Proof. We can use the algorithm given in the proof of Proposition 11 with a slight modification. The modification is that we also keep a list of objects that we initialize to the empty set. Then, as we obtain a new inconsistency in the algorithm, that is, $\mathcal{P}_{id,t} = Space$,

we check to see if the id is already in the list. We increment the inconsistency measure counter, $\mathcal{I}_O(\mathcal{S})$, and insert the object id into the list only if the object is not already in the list. \square

Proposition 13. *UPPER for the dimensional inconsistency measure \mathcal{I}_T is in P.*

Proof. Similar to the proof of Proposition 12, we can use the algorithm given in the proof of Proposition 11 with a slight modification. The modification is that we also keep a list of time values that we initialize to the empty set. Then, as we obtain a new inconsistency in the above algorithm, that is, $\mathcal{P}_{id,t} = \text{Space}$, we check to see if the time value is already in the list. We increment the inconsistency measure counter, $\mathcal{I}_T(\mathcal{S})$, and insert the time value t into the list only if the time value is not already in the list. \square

Complexity becomes more difficult to assess for the measures involving space. However, we have nice results for the case where the regions are rectangular. This is important because we expect many applications to deal with such a scenario. As stated in the following two propositions, both \mathcal{I}_S and \mathcal{I}_D are in P under the restriction of isothetic rectangular regions.

Proposition 14. *UPPER for the dimensional inconsistency measure \mathcal{I}_S is in P under the restriction of isothetic rectangular regions.*

Proof. Proposition 1 states that under the restriction of isothetic rectangular regions, every minimal inconsistent subset of an ST database has size 2. The algorithm proceeds as follows. Initialize $R_S = \emptyset$. We work with one $\mathcal{S}_{id,t}$ at a time. Let $\mathcal{S}_{id,t} = \{a_1, \dots, a_k\}$. If $k = 0$ or $k = 1$ we ignore $\mathcal{S}_{id,t}$. Otherwise, check each pair $\{a_i, a_j\}$, $1 \leq i < j \leq k$ for an inconsistency ($r_i \cap r_j = \emptyset$). If $r_i \cap r_j = \emptyset$, let $R_S = R_S \cup r_i \cup r_j$. After finishing all $\mathcal{S}_{id,t}$, compute $|R_S|$. Clearly, this algorithm is in P. \square

Proposition 15. *UPPER for the dimensional inconsistency measure \mathcal{I}_D is in P under the restriction of isothetic rectangular regions.*

Proof. Proposition 1 states that under the restriction of isothetic rectangular regions, every minimal inconsistent subset of an ST database has size 2. The algorithm proceeds as follows. Initialize $\mathcal{I}_D = 0$. We work with one $\mathcal{S}_{id,t}$ at a time. Let $\mathcal{S}_{id,t} = \{a_1, \dots, a_k\}$. If $k = 0$ or $k = 1$ we ignore $\mathcal{S}_{id,t}$. Otherwise, check each pair $\{a_i, a_j\}$, $1 \leq i < j \leq k$ for an inconsistency ($r_i \cap r_j = \emptyset$). If $r_i \cap r_j = \emptyset$, let $\mathcal{I}_D(\mathcal{S}) = \mathcal{I}_D(\mathcal{S}) + d(r_i, r_j)$. Assuming that d is computable in polynomial time, this algorithm is in P. \square

Table 4 summarizes the complexity results stated in this section. Notice that for \mathcal{I}_S and \mathcal{I}_D , membership in P holds under the restriction of isothetic rectangular regions. This restriction ensures that every minimal inconsistent subset has cardinality 2, while in the general case a minimal inconsistent subset can have arbitrary cardinality (e.g., $\mathcal{S} = \{(id, \{p_1, p_2\}, t), (id, \{p_1, p_3\}, t), (id, \{p_2, p_3\}, t)\}$ yields a ternary minimal inconsistent subset). For the remaining measures the problem is in P without restrictions.

Dimensional IMs					
	\mathcal{I}_O	\mathcal{I}_T	\mathcal{I}_{OT}	\mathcal{I}_S	\mathcal{I}_D
Complexity	P	P	P	P *	P *

Table 4: Complexity of dimensional inconsistency measures (*: under the restriction of isothetic rectangular regions).

7.2 Complexity Results for Repair-Based Measures

We deal now with the complexity of the repair-based measures. We start with the cases where the repair is done along the object or time dimension, for which we show NP-completeness. Membership in NP holds for all metrics d_{ID} and d_T that can be computed in polynomial time. Hardness is shown via reductions that employ specific metrics.

Theorem 12. *UPPER for the repair-based inconsistency measure \mathcal{I}_{id} is NP-complete.*

Proof. (Membership) First, it is important to observe that to decide whether $\mathcal{I}_{id}(\mathcal{S}) \leq k$, it suffices to decide whether there exists a function $idcorr$ conforming Definition 8 such that $\sum_{a \in \mathcal{S}} cost_{id}(a, idcorr(a)) \leq k$. In fact, if such a function $idcorr$ exists then there is an id-correction \mathcal{S}' such that $cost_{id}(\mathcal{S}, \mathcal{S}') = \sum_{a \in \mathcal{S}} cost_{id}(a, idcorr(a)) \leq k$, meaning that every id-repair \mathcal{S}^* for \mathcal{S} is such that $cost_{id}(\mathcal{S}, \mathcal{S}^*) \leq k$, which in turn implies that $\mathcal{I}_{id}(\mathcal{S}) \leq k$. Thus, to show membership in NP, it suffices to show a guess-and-check strategy for deciding whether such a function $idcorr$ exists, as follows: first guess $idcorr$ consisting of an id-update for each atom in \mathcal{S} , and then check in polynomial time that (i) the resulting ST database is consistent (i.e., for each id, t pair, it is the case that $\bigcap_{(id, r, t) \in \mathcal{S}_{id, t}} r \neq \emptyset$); and (ii) $\sum_{a \in \mathcal{S}} cost_{id}(a, idcorr(a)) \leq k$.

(Hardness) We show a reduction to our problem from the NP-hard Hitting Set problem (Papadimitriou, 1994), which is defined as follows: given a collection C of subsets of a set U , i.e., $C = \{S_1, \dots, S_n\}$ where $S_i \subseteq U$, and a positive integer K , decide whether U contains a hitting set for C of size K or less, that is, decide whether there is $H \subseteq U$ with $|H| \leq K$ and such that H contains at least one element from each subset in C .

Given U, C , and K , we construct an instance of our problem as follows. Let $Space = U \cup \{p_1^*, \dots, p_{K-1}^*\}$ where p_1^*, \dots, p_{K-1}^* are new elements not occurring in U , $T = \{t\}$, $ID = \{id\} \cup \{id_1, \dots, id_{K-1}\} \cup \{id'_1, \dots, id'_{K-1}\}$, and

$$\mathcal{S} = \{(id, S_i, t) \mid S_i \in C\} \cup \{(id_j, \{p_j^*\}, t) \mid 1 \leq j < K\}.$$

Hence, \mathcal{S} consists of $n + K - 1$ atoms, where $n = |C|$.

Assume d_{ID} is as follows:

- For all $i \in \{1, \dots, K - 1\}$, $d_{ID}(id, id_i) = 1$;
- For all $i, j \in \{1, \dots, K - 1\}$, $d_{ID}(id_i, id'_j) = n$;
- For all $j \in \{1, \dots, K - 1\}$, $d_{ID}(id, id'_j) = n + 1$;
- For all $i, j \in \{1, \dots, K - 1\}, i \neq j$, $d_{ID}(id_i, id_j) = d_{ID}(id'_i, id'_j) = 2$.

Finally, let $k = n(K - 1) + (n - 1)$.

We now show that there is $H \subseteq U$ such that $|H| \leq K$ and for all $i \in \{1, \dots, n\}$, $H \cap S_i \neq \emptyset$ iff $\mathcal{I}_{id}(\mathcal{S}) \leq k$.

(\Rightarrow) Let $H = \{p_0, \dots, p_{|H|-1}\} \subseteq U$ be a hitting set for C . W.l.o.g. assume that the elements in H are ordered by the number of sets of C with which they have non-empty intersection in descending order, that is, $\ell < j$ implies $|\{S_i \mid S_i \in C, S_i \cap \{p_\ell\} \neq \emptyset\}| \geq |\{S_i \mid S_i \in C, S_i \cap \{p_j\} \neq \emptyset\}|$. Let $m = |\{S_i \mid S_i \in C, S_i \cap \{p_0\} \neq \emptyset\}|$ be the maximum number of sets in C which intersect an element in U (this element is p_0 in our construction; clearly $m \geq 1$).

Consider the ST database \mathcal{S}' defined as follows:

$$\begin{aligned} \mathcal{S}' = & \{(id, S_i, t) \mid S_i \cap \{p_0\} \neq \emptyset\} \cup \\ & \{(id_j, S_i, t) \mid S_i \cap \{p_j\} \neq \emptyset, j \geq 1, \text{ there is no } \ell < j \text{ such that } S_i \cap \{p_\ell\} \neq \emptyset\} \cup \\ & \{(id'_j, \{p_j^*\}, t) \mid S_i \cap \{p_j\} \neq \emptyset, j \geq 1, \text{ there is no } \ell < j \text{ such that } S_i \cap \{p_\ell\} \neq \emptyset\} \cup \\ & \{(id'_j, \{p_j^*\}, t) \mid j \in \{|H|, \dots, K - 1\}\}. \end{aligned}$$

Thus, $cost(\mathcal{S}, \mathcal{S}') = n(|H| - 1) + (n - m)$ since \mathcal{S}' is obtained from \mathcal{S} by performing $(|H| - 1)$ id-updates changing id_j to id'_j with cost equal to n plus $(n - m)$ id-updates changing id to id_j with cost equal to 1.

It is easy to see that \mathcal{S}' is consistent. This is implied by the fact that H is a hitting set for $\{S_1, \dots, S_n\}$. In fact, a model M for \mathcal{S}' is as follows: $M(id, t) = p_0$, $M(id_j, t) = p_j$ and $M(id'_j, t) = p_j^*$ for all $j \in \{1, \dots, |H| - 1\}$, and $M(id_j, t) = p_j^*$ for all $j \in \{|H|, \dots, K - 1\}$. Hence, \mathcal{S}' is an id-correction for \mathcal{S} . Therefore, $\mathcal{I}_{id}(\mathcal{S}) \leq n(|H| - 1) + (n - m) \leq n(K - 1) + (n - 1) = k$, since $|H| \leq K$ and $m \geq 1$.

(\Leftarrow) Let \mathcal{S}' be an id-repair for \mathcal{S} , and $idcorr$ be an id-correction such that $\mathcal{I}_{id}(\mathcal{S}) = cost_{id}(\mathcal{S}, \mathcal{S}') = \sum_{a \in \mathcal{S}} cost_{id}(a, idcorr(a)) \leq k$, where $k = n(K - 1) + (n - 1)$.

We first observe that, for each atom $a_i = (id, S_i, t) \in \mathcal{S}$, it is the case that:

- 1) if $idcorr(a_i) = (id_j, S_i, t)$ then $idcorr(id_j, \{p_j^*\}, t) = (id'_j, \{p_j^*\}, t)$, and the total cost of these two id-updates is $1 + n$, as $d_{ID}(id, id_i) = 1$ and $d_{ID}(id_i, id'_j) = n$;
- 2) if $idcorr(a_i) = (id'_j, S_i, t)$ then the cost of this id-update is again $n + 1$, as $d_{ID}(id, id'_j) = n + 1$;
- 3) id-updates changing id_i into id_j , or id'_i into id'_j (with $i \neq j$), are not useful to define an id-correction whose cost is minimal, as the former leads to inconsistency which should be resolved with additional and unnecessary id-updates and the latter is never applied, since the id'_i 's do not appear in \mathcal{S} .

Therefore, if $idcorr$ updates a single atom $a_i = (id, S_i, t) \in \mathcal{S}$, the minimal cost of such an id-update is $n + 1$.

On the other hand, updating $x = |P_j|$ atoms in $P_j = \{a_{i_1} = (id, S_{i_1}, t), \dots, a_{i_x} = (id, S_{i_x}, t)\} \in \mathcal{S}$ for which there exists $p_j \in Space$ such that $p_j \in \bigcap_{\ell \in [1..x]} S_{i_\ell}$ has a minimal total cost equal to $|P_j| + n$, as this is the cost of $|P_j|$ id-updates of the form $idcorr(a_{i_\ell}) = (id_j, S_{i_\ell}, t)$ for all $\ell \in [1..|P_j|]$, each of them having cost $d_{ID}(id, id_i) = 1$, plus an additional id-update $idcorr(id_j, \{p_j^*\}, t) = (id'_j, \{p_j^*\}, t)$ whose cost is n .

From what was said above, it follows that $idcorr$ has the following properties:

- a) For each atom $a_i = (id, S_i, t) \in \mathcal{S}$ such that there exists $p_0 \in Space$, $S_i \cap \{p_0\} \neq \emptyset$, and $p_0 = \arg \max_{p \in Space} \{ |\{S_j \mid (id, S_j, t) \in \mathcal{S}, S_j \cap \{p\} \neq \emptyset\}| \}$, it is the case that $idcorr(a_i) = a_i$ (with cost $d_{ID}(id, id) = 0$). This way, $cost_{id}(\mathcal{S}, \mathcal{S}')$ is minimized because keeping unchanged any other set of atoms of \mathcal{S} would have a larger overall cost.
- b) To achieve consistency, the other atoms in \mathcal{S} need to be updated, and the minimal cost to do that is obtained by partitioning them into the minimal number B of sets P_1, \dots, P_B such that for each P_j , there exists $p_j \in Space, p_j \in \bigcap \{S_i \mid (id, S_i, t) \in P_j\}$. In fact, it turns out that $\sum_{a \in \mathcal{S}} cost_{id}(a, idcorr(a)) = \sum_{j=1}^B (n + |P_j|)$, as updating a set of $|P_j|$ atoms in \mathcal{S} has minimal cost equal to $|P_j| + n$, as discussed earlier.

Observe that $\sum_{j=1}^B (n + |P_j|) = nB + (n - m)$, where m (with $1 \leq m \leq n$) is the number of atoms $a_i \in \mathcal{S}$ such that $idcorr(a_i) = a_i$, i.e., those whose id is not changed. Therefore $cost_{id}(\mathcal{S}, \mathcal{S}') = nB + (n - m) \leq nB + (n - 1)$.

Now, as $\mathcal{I}_{id}(\mathcal{S}) \leq n(K - 1) + (n - 1)$, it must be the case that $B \leq K - 1$, meaning that the p_0 defined as in Item a) union the $\{p_j\}_{j \in [1..B]}$ defined as in Item b) above is a hitting set for $\{S_1, \dots, S_n\}$ of size K or less. \square

Theorem 13 below also shows an NP-complete result for \mathcal{I}_{time} .

Theorem 13. UPPER for the repair-based inconsistency measure \mathcal{I}_{time} is NP-complete.

Proof. Both the membership in NP and the hardness follow by reasoning analogously to the proof of the \mathcal{I}_{id} inconsistency measure (cf. Theorem 12). Specifically, hardness can be proved by showing a reduction from an instance $\langle U, C, K \rangle$ of the Hitting Set problem to an instance of our problem where $Space = U \cup \{p_1^* \dots, p_{K-1}^*\}$ where $p_1^* \dots, p_{K-1}^*$ are new elements not occurring in U , $ID = \{id\}$, $T = \{t\} \cup \{t_1, \dots, t_{K-1}\} \cup \{t'_1, \dots, t'_{K-1}\}$, $\mathcal{S} = \{(id, S_i, t) \mid S_i \in C\} \cup \{(id, \{p_j^*\}, t_j) \mid 1 \leq j < K\}$, and d_T is as follows:

- For all i such that $1 \leq i < K$, $d_T(t, t_i) = 1$;
- For all i, j such that $1 \leq i, j < K$, $d_T(t_i, t'_j) = n$;
- For all j such that $1 \leq j < K$, $d_T(t, t'_j) = n + 1$;
- For all distinct i, j such that $1 \leq i, j < K$, $d_T(t_i, t_j) = d_T(t'_i, t'_j) = 2$.

Given this, it can be shown that U contains a hitting set for C of size K or less iff $\mathcal{I}_{time}(\mathcal{S}) \leq k$, where $k = n(K - 1) + (n - 1)$. \square

The complexity is different for the repair-based measure dealing with regions. We get a nice result under the reasonable assumption that the metric among regions is the symmetric difference. In fact, in such a case, the next result shows that computing \mathcal{I}_{region} is tractable.

Theorem 14. UPPER for the repair-based inconsistency measure \mathcal{I}_{region} is in P under the restriction that the symmetric difference is used as the metric.

Proof. We work with one $\mathcal{S}_{id,t}$ at a time. Clearly, we only need to deal with inconsistent $\mathcal{S}_{id,t}$ s. The idea of the computation is to find a point that belongs to the maximum number of regions (of the atoms of $\mathcal{S}_{id,t}$) and add it to all the other regions as the repair. As d_R is symmetric difference, the addition of a point to a region has distance 1; so this is the best that can be achieved. For convenience we write more details below. Let $\mathcal{S}_{id,t} = \{a_i = (id, r_i, t) | 1 \leq i \leq n\}$. For each $p \in Space$ compute $count(p, \mathcal{S}_{id,t}) = |\{a_i \mid p \in r_i\}|$. Let p_{max} be a point with maximum $count$ number. Define $f : \mathcal{S}_{id,t} \rightarrow \mathcal{S}'_{id,t}$ where $f(a_i) = (id, r_i \cup \{p_{max}\}, t)$. The function f is a region-correction as $\mathcal{S}'_{id,t}$ is consistent (p_{max} is in every region) and is a region-repair as well. Hence $\mathcal{I}_{region}(\mathcal{S}_{id,t}) = cost_{region}(\mathcal{S}_{id,t}, \mathcal{S}'_{id,t}) = |\mathcal{S}_{id,t}| - count(p_{max}, \mathcal{S}_{id,t})$. Finally, $\mathcal{I}_{region}(\mathcal{S}) = \sum_{(id,t)} \text{there exists } (id,r,t) \in \mathcal{S} \text{ } cost(\mathcal{S}_{id,t})$. Clearly, the procedure described above can be accomplished in polynomial time, and hence the computation of $\mathcal{I}_{region}(\mathcal{S}) \leq k$ is also in P . \square

Our next result shows that computing the repair-based measure dealing with deletions is also feasible. This immediately follows from the tractability result for \mathcal{I}_{region} shown above and the fact that the value of \mathcal{I}_{card} coincides with that of \mathcal{I}_{region} when using the symmetric difference as metric (cf. Theorem 11).

Corollary 2. *UPPER for the repair-based inconsistency measure \mathcal{I}_{card} is in P .*

Finally, under the assumptions of Theorem 10, the tractability of \mathcal{I}_{card} entails that both \mathcal{I}_{id} and \mathcal{I}_{time} are tractable.

Corollary 3. *Given an ST database \mathcal{S} , if $d_{ID}(id, id') = 1$ for all distinct $id, id' \in ID$ and $|ID| \geq |\mathcal{S}|$, then UPPER for the repair-based inconsistency measure \mathcal{I}_{id} is in P .*

Corollary 4. *Given an ST database \mathcal{S} , if $d_{ID}(t, t') = 1$ for all distinct $t, t' \in T$ and $|T| \geq |\mathcal{S}|$, then UPPER for the repair-based inconsistency measure \mathcal{I}_{time} is in P .*

The complexity results for the repair-based inconsistency measures is summarized in Table 5. Although we believe that measuring inconsistency as the cost of restoring consistency is a natural approach, our complexity analysis shows that this comes at a cost for \mathcal{I}_{id} and \mathcal{I}_{time} for general ST databases, though we found some tractable cases for restricted ST databases.

On the whole, Tables 4 and 5 show that most of the inconsistency measures proposed can be evaluated in polynomial time. This suggests a nice computational behavior for our framework, compared with the case of general propositional knowledge bases where the complexity of inconsistency measurement is much higher, e.g., many measures belong to the second or third level of the polynomial hierarchy (Thimm & Wallner, 2016; Thimm, 2018).

8. Conclusion and Future Work

Though there has been an extensive body of work on measuring inconsistency in propositional knowledge bases (Hunter & Konieczny, 2008; Grant & Hunter, 2013, 2011; Thimm, 2016; Doder et al., 2010; Jabbour, Ma, & Raddaoui, 2014; Mu, Liu, Jin, & Bell, 2011), little work has been done on measuring inconsistency in spatio-temporal databases (Grant & Parisi, 2020, 2021).

Repair-based IMs				
	\mathcal{I}_{id}	\mathcal{I}_{time}	\mathcal{I}_{region}	\mathcal{I}_{card}
Complexity	NP-complete	NP-complete	P †	P

Table 5: Complexity of repair-based inconsistency measures (†: for the symmetric difference metric).

Some authors have already considered inconsistency measures separately in spatial and temporal databases, such as Condotta, Raddaoui, and Salhi (2016), who introduced inconsistency measures for qualitative constraint networks, and Brisaboa, Luaces, Rodríguez, and Seco (2014), who proposed inconsistency measures to evaluate how dirty a spatial dataset is with respect to a set of topological constraints. However, for spatio-temporal databases, to the best of our knowledge, this is new work that follows up on a very preliminary study done by Grant et al. (2018a). Besides considering more postulates than Grant et al. (2018a), we introduced the novel notions of *repair-based inconsistency measures* and analyzed them with respect to postulate satisfaction, we introduced the concept of *dimensional postulates*, and investigated the complexity of the inconsistency measures for spatio-temporal databases. Notably, this is the first work considering “dimensions” in inconsistency measurement and postulates—we believe this is an important aspect that should be taken into account in general non-propositional knowledge bases, not only in ST databases.

Our complexity characterization provides helpful results for addressing also the complexity of the (function) problem of computing the exact value of inconsistency measures. In fact, for the cases where the investigated decision problem is in P, the algorithms given in the proofs show that computing the value of the inconsistency measure is polynomial too. On the other hand, our NP-complete results entail that $FP^{NP[\log n]}$ is an upper bound on the complexity of the problem of computing the value of the repair-based inconsistency measures \mathcal{I}_{id} and \mathcal{I}_{time} whenever the distance between any pair of object identifiers (resp., time values) can be assumed to be bounded by n^c , where n is the number of ST atoms in the database and c is an arbitrary but fixed constant. In fact, under the assumption that $\forall id, id' \in ID, d_{ID}(id, id') \leq n^c$ (resp., $\forall t, t' \in T, d_T(t, t') \leq n^c$), the set V of the distinct values that \mathcal{I}_{id} (resp., \mathcal{I}_{time}) can have is $V = \{0, 1, 2, \dots, (n-1) \cdot n^c, \infty\}$. Thus, the cardinality of V is polynomial in the number of ST atoms. Since we showed that deciding whether $\mathcal{I}_{id}(\mathcal{S}) \leq k$ (resp., $\mathcal{I}_{time}(\mathcal{S}) \leq k$) is in NP, it follows that the value of these measures can be computed by performing a binary search on V , i.e., by asking a logarithmic number of queries to an NP-oracle deciding whether $\mathcal{I}_{id}(\mathcal{S}) \leq k$ (resp., $\mathcal{I}_{time}(\mathcal{S}) \leq k$), where $k \in V$ is chosen as usual in the binary search. This means that the value of \mathcal{I}_{id} (resp., \mathcal{I}_{time}) can be computed by a deterministic polynomial-time algorithm asking a logarithmic number of queries to an NP oracle solving UPPER for \mathcal{I}_{id} (resp., \mathcal{I}_{time}), from which the $FP^{NP[\log n]}$ upper bound follows. Providing also a lower bound on the complexity of the function problem for these measures, and thus obtaining a tight characterization, is left for future work.

Many other interesting issues concerning inconsistency measures in spatio-temporal databases remain to be investigated. We would like to extend our results to the case where an ST database is augmented with probabilities (Grant et al., 2010; Grant, Parisi, & Subrahmanian, 2013; Doder, Grant, & Ognjanović, 2013) and integrity constraints (Parisi

& Grant, 2016, 2020). To address the problem of measuring inconsistency in probabilistic spatio-temporal databases, we will start by considering the approaches proposed for measuring inconsistency in probabilistic logics (Thimm, 2013; Potyka, 2014; Bona & Finger, 2015; Bona, Finger, Potyka, & Thimm, 2018), where classical rationality postulates have been revisited for the probabilistic setting and the difficulty of dealing with probabilistic logics is discussed. Augmenting ST databases with integrity constraints (e.g., expressing restrictions on how objects can move over time) would yield a more general framework. Of course, in such a more expressive framework, an ST database can still be inconsistent because its ST atoms cannot be jointly satisfied as per Definition 3. However, there is now a new source of inconsistency: the ST atoms may violate the integrity constraints. Studying how constraints affect inconsistency measures, postulate satisfaction, and the complexity of inconsistency measures are intriguing issues to be explored. As another direction for future work, we plan to develop *relative* inconsistency measures (Grant, 2018) for ST databases—this is a quite unexplored topic also for propositional knowledge bases and has been recently addressed by Besnard and Grant (2020). Finally, we would like to consider practical applications involving GIS databases and data stored using RDF/ontologies.

Acknowledgments

The authors would like to thank the anonymous referees for many useful comments which resulted in a substantially improved paper.

References

- Agarwal, D., Chen, D., Ji Lin, L., Shanmugasundaram, J., & Vee, E. (2010). Forecasting high-dimensional data. In *Proceedings of the International Conference on Management of Data (SIGMOD)*, pp. 1003–1012.
- Agarwal, P. K., Arge, L., & Erickson, J. (2003). Indexing moving points. *Journal of Computer and System Sciences*, 66(1), 207–243.
- Akdere, M., Cetintemel, U., Riondato, M., Upfal, E., & Zdonik, S. B. (2011). The case for predictive database systems: Opportunities and challenges. In *Proceedings of the Biennial Conference on Innovative Data Systems Research (CIDR)*, pp. 167–174.
- An, L., Chen, X., Kafai, M., Yang, S., & Bhanu, B. (2013). Improving person re-identification by soft biometrics based reranking. In *Proceedings of the International Conference on Distributed Smart Cameras (ICDSC)*, pp. 1–6.
- Bayir, M. A., Demirbas, M., & Eagle, N. (2010). Mobility profiler: A framework for discovering mobility profiles of cell phone users. *Pervasive and Mobile Computing*, 6(4), 435–454.
- Bedagkar-Gala, A., & Shah, S. K. (2014). A survey of approaches and trends in person re-identification. *Image and Vision Computing*, 32(4), 270–286.
- Bertossi, L. E. (2011). *Database Repairing and Consistent Query Answering*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers.

- Bertossi, L. E. (2018). Measuring and computing database inconsistency via repairs. In *Proceedings of International Conference on Scalable Uncertainty Management (SUM)*, pp. 368–372.
- Besnard, P. (2014). Revisiting postulates for inconsistency measures. In *Proceedings of the European Conference Logics in Artificial Intelligence (JELIA)*, pp. 383–396.
- Besnard, P., & Grant, J. (2020). Relative inconsistency measures. *Artificial Intelligence*, 280, 103231.
- Bona, G. D., & Finger, M. (2015). Measuring inconsistency in probabilistic logic: rationality postulates and dutch book interpretation. *Artificial Intelligence*, 227, 140–164.
- Bona, G. D., Finger, M., Potyka, N., & Thimm, M. (2018). Inconsistency measurement in probabilistic logic. In Grant, & Martinez (Grant & Martinez, 2018), pp. 235–269.
- Bona, G. D., Grant, J., Hunter, A., & Konieczny, S. (2019). Classifying inconsistency measures using graphs. *Journal of Artificial Intelligence Research*, 66, 937–987.
- Brisaboa, N. R., Luaces, M. R., Rodríguez, M. A., & Seco, D. (2014). An inconsistency measure of spatial data sets with respect to topological constraints. *International Journal of Geographical Information Science*, 28(1), 56–82.
- Cohn, A. G., & Hazarika, S. M. (2001). Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae*, 46(1-2), 1–29.
- Condotta, J.-F., Raddaoui, B., & Salhi, Y. (2016). Quantifying conflicts for spatial and temporal information. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 443–452.
- Doder, D., Grant, J., & Ognjanović, Z. (2013). Probabilistic logics for objects located in space and time. *Journal of Logic and Computation*, 23(3), 487–515.
- Doder, D., Raskovic, M., Markovic, Z., & Ognjanovic, Z. (2010). Measures of inconsistency and defaults. *International Journal of Approximate Reasoning*, 51(7), 832–845.
- Fazzinga, B., Flesca, S., Furfaro, F., & Parisi, F. (2014). Offline cleaning of RFID trajectory data. In *Proceedings of the International Conference on Scientific and Statistical Database Management (SSDBM)*, p. 5.
- Fazzinga, B., Flesca, S., Furfaro, F., & Parisi, F. (2016). Exploiting integrity constraints for cleaning trajectories of RFID-monitored objects. *ACM Transactions on Database Systems*, 41(4).
- Gabelaia, D., Kontchakov, R., Kurucz, Á., Wolter, F., & Zakharyashev, M. (2005). Combining spatial and temporal logics: Expressiveness vs. complexity. *Journal of Artificial Intelligence Research*, 23, 167–243.
- Grant, J. (2018). From measuring infinities to measuring inconsistencies. In Grant, & Martinez (Grant & Martinez, 2018), pp. 1–18.
- Grant, J., & Hunter, A. (2011). Measuring consistency gain and information loss in stepwise inconsistency resolution. In *Proceedings of the European Conference Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, pp. 362–373.

- Grant, J., & Hunter, A. (2013). Distance-based measures of inconsistency. In *Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, pp. 230–241.
- Grant, J., & Martinez, M. V. (Eds.). (2018). *Measuring Inconsistency in Information*, Vol. 73 of *Studies in Logic*. College Publications.
- Grant, J., Martinez, M. V., Molinaro, C., & Parisi, F. (2018a). On measuring inconsistency in spatio-temporal databases. In Grant, & Martinez (Grant & Martinez, 2018), pp. 313–342.
- Grant, J., Molinaro, C., & Parisi, F. (2018b). Probabilistic spatio-temporal knowledge bases: Capacity constraints, count queries, and consistency checking. *International Journal of Approximate Reasoning*, 100, 1–28.
- Grant, J., & Parisi, F. (2020). Measuring inconsistency in a general information space. In *Proceedings of the International Symposium on Foundations of Information and Knowledge Systems (FoIKS)*, pp. 140–156.
- Grant, J., & Parisi, F. (2021). General information spaces: measuring inconsistency, rationality postulates, and complexity. *Annals of Mathematics and Artificial Intelligence*, 89, 1–35.
- Grant, J., Parisi, F., Parker, A., & Subrahmanian, V. S. (2010). An agm-style belief revision mechanism for probabilistic spatio-temporal logics. *Artificial Intelligence*, 174(1), 72–104.
- Grant, J., Parisi, F., & Subrahmanian, V. S. (2013). Research in probabilistic spatiotemporal databases: The spot framework. In *Advances in Probabilistic Databases for Uncertain Information Management*, Vol. 304 of *Studies in Fuzziness and Soft Computing*, pp. 1–22. Springer.
- Hadjieleftheriou, M., Kollios, G., Tsotras, V. J., & Gunopulos, D. (2002). Efficient indexing of spatiotemporal objects. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*, pp. 251–268.
- Hammel, T., Rogers, T. J., & Yetso, B. (2003). Fusing live sensor data into situational multimedia views. In *Multimedia Information Systems*, pp. 145–156.
- Hunter, A., & Konieczny, S. (2008). Measuring inconsistency through minimal inconsistent sets. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 358–366.
- Jabbour, S., Ma, Y., & Raddaoui, B. (2014). Inconsistency measurement thanks to MUS decomposition. In *Proceedings of the International conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 877–884.
- Karbassi, A., & Barth, M. (2003). Vehicle route prediction and time of arrival estimation techniques for improved transportation system management. In *Proceedings of IEEE Intelligent Vehicles Symposium (IV)*, pp. 511–516.
- Knapp, A., Merz, S., Wirsing, M., & Zappe, J. (2006). Specification and refinement of mobile systems in mta and mobile uml. *Theoretical Computer Science*, 351(2), 184–202.

- Kollios, G., Gunopulos, D., & Tsotras, V. J. (1999). On indexing mobile objects. In *Proceedings of the Symposium on Principles of Database Systems (PODS)*, pp. 261–272.
- Kurkovsky, S., & Harihar, K. (2006). Using ubiquitous computing in interactive mobile marketing. *Personal Ubiquitous Computing*, 10(4), 227–240.
- Mittu, R., & Ross, R. (2003). Building upon the coalitions agent experiment (COAX) - integration of multimedia information in GCCS-M using IMPACT. In *Multimedia Information Systems*, pp. 35–44.
- Mu, K., Liu, W., Jin, Z., & Bell, D. A. (2011). A syntax-based approach to measuring the degree of inconsistency for belief bases. *International Journal of Approximate Reasoning*, 52(7), 978–999.
- Papadimitriou, C. M. (1994). *Computational complexity*. Addison-Wesley, Reading, MA.
- Parisi, F., & Grant, J. (2016). Knowledge representation in probabilistic spatio-temporal knowledge bases. *Journal of Artificial Intelligence Research*, 55, 743–798.
- Parisi, F., & Grant, J. (2017). On repairing and querying inconsistent probabilistic spatio-temporal databases. *International Journal of Approximate Reasoning*, 84, 41–74.
- Parisi, F., & Grant, J. (2020). On measuring inconsistency in relational databases with denial constraints. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pp. 857–864.
- Parisi, F., Parker, A., Grant, J., & Subrahmanian, V. S. (2010). Scaling cautious selection in spatial probabilistic temporal databases. In *Methods for Handling Imperfect Spatial Information*, Vol. 256 of *Studies in Fuzziness and Soft Computing*, pp. 307–340. Springer.
- Parisi, F., Sliva, A., & Subrahmanian, V. S. (2013). A temporal database forecasting algebra. *International Journal of Approximate Reasoning*, 54(7), 827–860.
- Parker, A., Infantes, G., Grant, J., & Subrahmanian, V. S. (2009). Spot databases: Efficient consistency checking and optimistic selection in probabilistic spatial databases. *IEEE Transactions on Knowledge and Data Engineering*, 21(1), 92–107.
- Pelani, M., Saltenis, S., & Jensen, C. S. (2006). Indexing the past, present, and anticipated future positions of moving objects. *ACM Transactions on Database Systems*, 31(1), 255–298.
- Pfoser, D., & Jensen, C. S. (1999). Capturing the uncertainty of moving-object representations. In *Proceedings of the International Symposium on Advances in Spatial Databases (SSD)*, pp. 111–132.
- Pfoser, D., Jensen, C. S., & Theodoridis, Y. (2000). Novel approaches in query processing for moving object trajectories. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pp. 395–406.
- Potyka, N. (2014). Linear programs for measuring inconsistency in probabilistic logics. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*.

- Samet, H. (2006). *Foundations of multidimensional and metric data structures*. Morgan Kaufmann series in data management systems. Academic Press.
- Sedky, M. H., Moniri, M., & Chibelushi, C. C. (2005). Classification of smart video surveillance systems for commercial applications. In *Proceedings of the International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pp. 638–643.
- Southey, F., Loh, W., & Wilkinson, D. F. (2007). Inferring complex agent motions from partial trajectory observations. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2631–2637.
- Tao, Y., Papadias, D., & Sun, J. (2003). The TPR*-tree: an optimized spatio-temporal access method for predictive queries. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, pp. 790–801.
- Thimm, M. (2013). Inconsistency measures for probabilistic logics. *Artificial Intelligence*, *197*, 1–24.
- Thimm, M. (2016). Stream-based inconsistency measurement. *International Journal of Approximate Reasoning*, *68*, 68–87.
- Thimm, M. (2017). On the compliance of rationality postulates for inconsistency measures: A more or less complete picture. *Künstliche Intelligenz*, *31*(1), 31–39.
- Thimm, M. (2018). On the evaluation of inconsistency measures. In Grant, & Martinez (Grant & Martinez, 2018), pp. 19–60.
- Thimm, M., & Wallner, J. P. (2016). Some complexity results on inconsistency measurement. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 114–124.
- Xiao, G., & Ma, Y. (2012). Inconsistency measurement based on variables in minimal unsatisfiable subsets. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pp. 864–869.
- Yaman, F., Nau, D. S., & Subrahmanian, V. S. (2004). A logic of motion. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 85–94.
- Yaman, F., Nau, D. S., & Subrahmanian, V. S. (2005). Going far, logically. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 615–620.
- Zhang, B., & Trajcevski, G. (2014). The tale of (fusing) two uncertainties. In *Proceedings of the International Conference on Advances in Geographic Information Systems (GIS)*, pp. 521–524.
- Zhang, L., Kalashnikov, D. V., Mehrotra, S., & Vaisenberg, R. (2014). Context-based person identification framework for smart video surveillance. *Machine Vision and Applications*, *25*(7), 1711–1725.