Dimensional Inconsistency Measures and Postulates in Spatio-Temporal Databases (Extended Abstract)*

John Grant[1] , **Maria Vanina Martinez**[2] , **Cristian Molinaro**[3] , **Francesco Parisi**[3]
[1]Department of Computer Science and UMIACS, University of Maryland, USA
[2]Department of Computer Science, University of Buenos Aires, Argentina
[3] DIMES Department, University of Calabria, Italy
grant@cs.umd.edu, mvmartinez@dc.uba.ar, {cmolinaro, fparisi}@dimes.unical.it

Abstract

We define and investigate new inconsistency measures that are particularly suitable for dealing with inconsistent spatio-temporal information, as 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 Spatio-Temporal Databases

The representation and processing of spatio-temporal data has attracted much attention by the AI community [Cohn and Hazarika, 2001; Gabelaia *et al.*, 2005; Yaman *et al.*, 2005; Knapp *et al.*, 2006]. In this paper, we focus on Spatio-Temporal (ST) Databases (DBs), representing atomic statements of the form "object id is/was/will be inside region r at time t". Spatio-Temporal databases can be viewed as a special case of Spatial Probabilistic Temporal Databases [Parker *et al.*, 2009; Parisi *et al.*, 2010; Grant *et al.*, 2010; Grant *et al.*, 2013b; Grant *et al.*, 2013a; Parisi and Grant, 2016; Parisi and Grant, 2017; Grant *et al.*, 2017; Grant *et al.*, 2018].

Below we briefly introduce the syntax and semantics of ST databases. 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. We assume that an object can be in only one location at a time, but a single location may contain more than one object. A *region* is a nonempty subset of *Space*.

Definition 1. *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 *et al.*, 2014; Fazzinga *et al.*, 2016], but also information about the future, such as that derived from methods for predicting the destination of moving objects [Mittu and Ross, 2003; Hammel *et al.*, 2003; Southey *et al.*, 2007], or from querying predictive databases [Akdere *et al.*, 2011; Agarwal *et al.*, 2010; Parisi *et al.*, 2013]. The meaning of an ST database is given by the interpretations that satisfy it.

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

An interpretation specifies a trajectory for each $id \in$ *ID* by saying where in *Space* object id was/is/will be at each time $t \in T$. We now define satisfaction and ST models.

Definition 3. *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.,* $\mathcal{S}$*).*

Finally, we define when an ST database is consistent.

Definition 4. *An ST database* $\mathcal{S}$ *is* consistent *iff there exists at least one model for* $\mathcal{S}$*.*

The set of all minimal (under set-inclusion) inconsistent subsets of an ST database $\mathcal{S}$ is denoted as $\mathcal{M}(\mathcal{S})$.

2 Notion of Inconsistency Measure for ST Databases and (Classical) Postulates

An inconsistency measure (IM) is a function that assigns a nonnegative real value or infinity to every ST database. We use $\mathcal{R}^{\infty}_{\geq 0}$ for the set of nonnegative real numbers and the infinity symbol, and use $\mathcal{L}$ for the set of all ST databases.

Definition 5. *An* inconsistency measure $\mathcal{I} : \mathcal{L} \rightarrow \mathcal{R}^{\infty}_{\geq 0}$ *is a function such that, for every* $\mathcal{S}, \mathcal{S}' \in \mathcal{L}$*, the following two properties hold:*

*This paper is an extended abstract of [Grant *et al.*, 2021] published in the *Journal of Artificial Intelligence Research*.

1. (Consistency) $\mathcal{I}(\mathcal{S}) = 0$ *iff* $\mathcal{S}$ *is consistent.*
2. (Monotony) *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 w.r.t. set-inclusion.

We now define several properties that some inconsistency measures possess and that are intuitively desirable. Specifically, we present a list of eight additional properties that have counterparts for propositional knowledge bases. 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 $Problematic(\mathcal{S}) = \bigcup_{M \in \mathcal{M}(\mathcal{S})} M$ and $Free(\mathcal{S}) = \mathcal{S} \setminus Problematic(\mathcal{S})$.

Definition 6. *Let* $\mathcal{I}$ *be an IM, 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* $M, M' \in \mathcal{M}(\mathcal{S})$ *and* $|M| < |M'|$ *then* $\mathcal{I}(M) > \mathcal{I}(M')$.
6. (Equal Conflict) *If* $M, M' \in \mathcal{M}(\mathcal{S})$ *and* $|M| = |M'|$ *then* $\mathcal{I}(M) = \mathcal{I}(M')$.
7. (MI-Normalization) *If* $M \in \mathcal{M}(\mathcal{S})$ *then* $\mathcal{I}(M) = 1$.
8. (MI-Separability) *If* $\mathcal{M}(\mathcal{S} \cup \mathcal{S}') = \mathcal{M}(\mathcal{S}) \cup \mathcal{M}(\mathcal{S}')$ *and* $\mathcal{M}(\mathcal{S}) \cap \mathcal{M}(\mathcal{S}') = \emptyset$ *then* $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$.

We will find that in most cases these postulates are not satisfied, suggesting that they are not appropriate for the new IMs, because such 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 4), which are more suitable for ST databases and the new inconsistency measures.

3 Inconsistency Measures for ST Databases

In this section, we propose several inconsistency measures that are relevant for ST databases, as they explicitly take into account the dimensions characterizing such data.

3.1 Dimensional Inconsistency Measures

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 IMs 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 $\mathcal{I}_P(K) = |\{p \mid p \text{ is a problematic proposition}\}|$ for a propositional knowledge base K. Actually, we did not find this definition in the literature on IMs. However, it is the absolute version of a relative IM studied in [Xiao and Ma, 2012]. So $\mathcal{I}_P$ is our inspiration for measuring inconsistency along the three dimensions.

In some cases we may just be interested in how many objects or how many time values are involved in an inconsistency, which leads us to the two IMs defined below.

An IM based strictly on objects is:

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

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.

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

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

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 IM is thus defined looking at id, t pairs in minimal inconsistent subsets:

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

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

Let us now turn our attention to the space dimension. For an ST database $\mathcal{S}$ we define a region $R_{\mathcal{S}}$ as follows:

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

Then, we define $\mathcal{I}_S(\mathcal{S}) = |R_{\mathcal{S}}|$. Thus, $\mathcal{I}_S$ counts the number of points that are in regions involved in an inconsistency.

For the next measure involving space, we require a metric $d : \mathit{Space} \times \mathit{Space} \rightarrow [0, \infty)$. First, we define distance for minimal inconsistent subsets. Let M be a minimal inconsistent subset of an ST database $\mathcal{S}$: $M = \{(id, r_1, t), \ldots, (id, r_n, t)\}$. We start by defining n new regions, one for each i, $1 \leq i \leq n$, as $R_i = \bigcap_{j \neq i} r_j$. Since M is a minimal inconsistent set, for each i, $1 \leq i \leq n$, $R_i \neq \emptyset$, but $\cap_{i=1}^{n} r_i = \emptyset$. We define the value $d(M)$ as:

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

We can think of $d(M)$ as the minimal distance in *Space* between any regions involved in a minimal inconsistent subset. Then, we define: $\mathcal{I}_D(\mathcal{S}) = \sum_{M \in \mathcal{M}(\mathcal{S})} d(M)$. Thus $\mathcal{I}_D$ sums the minimal distances between regions involved in minimal inconsistent subsets.

3.2 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 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. *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 ways of changing an inconsistent ST database to a consistent one using updates.

Definition 8. *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 $Xcorr : \mathcal{S} \to \mathcal{S}'$ s.t. 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 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. In the following, we denote by $cost_{id}(a, a')$ the cost of changing atom a into an id-update 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. *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. *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 $\mathcal{S}$.*

The next repairing strategy we use is based on minimally updating the time values associated with the ST atoms. 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 the measure $\mathcal{I}_{time}$ are analogous to id-repair and $\mathcal{I}_{id}$ but for the time dimension.

Definition 11. *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 $\mathcal{S}$, $cost_{time}(\mathcal{S}, \mathcal{S}') \leq cost_{time}(\mathcal{S}, \mathcal{S}'')$.*

Definition 12. *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 $\mathcal{S}$.*

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 as follows.

Definition 13. *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 $\mathcal{S}$, $cost_{region}(\mathcal{S}, \mathcal{S}') \leq cost_{region}(\mathcal{S}, \mathcal{S}'')$.*

Definition 14. *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 $\mathcal{S}$.*

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 number of removed atoms be minimal.

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

Definition 16. *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 $\mathcal{S}$.*

4 Dimensional Postulates

In this section, we propose dimensional versions for some of the postulates—we will find in the next section that in several cases the dimensional postulate holds even though the original postulate does not hold.

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

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

	Dimensional IMs					Repair-based IMs			
	$\mathcal{I}_O$	$\mathcal{I}_T$	$\mathcal{I}_{OT}$	$\mathcal{I}_S$	$\mathcal{I}_D$	$\mathcal{I}_{id}$	$\mathcal{I}_{time}$	$\mathcal{I}_{region}$	$\mathcal{I}_{card}$
Free-Formula Independence	✓	✓	✓	✓	✓	✗	✗	✗	✓
Penalty	✗	✗	✗	✗	✓	✗	✗	✗	✗
Dominance	✓	✓	✓	✗	✓	✓	✓	✗	✓
Super-Additivity	✗	✗	✗	✗	✓	✓	✓	✓	✓
Attenuation	✗	✗	✗	✗	✗	✗	✗	✗	✗
Equal Conflict	✓	✓	✓	✗	✗	✗	✗	✗	✓
MI-Normalization	✓	✓	✓	✗	✗	✗	✗	✗	✓
MI-Separability	✗	✗	✗	✗	✓	✗	✗	✗	✗
Dimensional Penalty	✓	✓	–	✓	✓	✓*	✓*	✓	–
Dimensional Super-Additivity	✓	✓	–	✓	✓	✓	✓	✓	–
Dimensional MI-Separability	✓	✓	–	✓	✓	✗	✗	✓	–
Complexity	P	P	P	P *	P *	NP-c	NP-c	P †	P

Table 1: Classical/dimensional postulate satisfaction and complexity of dimensional and repair-based IMs (✓*: satisfied when measures return a finite value; *: under the restriction of isothetic rectangular regions; †: for the symmetric difference metric; –: not applicable).

In the following, for an ST database $\mathcal{S}$, we use the following notations: $ID(\mathcal{S}) = \{id \mid (id, r, t) \in \mathcal{S}\}$, $Time(\mathcal{S}) = \{t \mid (id, r, t) \in \mathcal{S}\}$, and $Region(\mathcal{S}) = \bigcup\{r \mid (id, r, t) \in \mathcal{S}\}$.

Definition 18 (**Dimensional Super-Additivity**). *Let $\mathcal{I}$ be an IM, 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 $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') \geq \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$.*

Definition 19 (**Dimensional MI-Separability**). *Let $\mathcal{I}$ be an IM, 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 $\mathcal{M}(\mathcal{S} \cup \mathcal{S}') = \mathcal{M}(\mathcal{S}) \cup \mathcal{M}(\mathcal{S}')$, then $\mathcal{I}(\mathcal{S} \cup \mathcal{S}') = \mathcal{I}(\mathcal{S}) + \mathcal{I}(\mathcal{S}')$.*

5 Postulate Satisfaction and Complexity

We analyze which postulates, both classical (cf. Definition 6) and dimensional (cf. Section 4) ones, are satisfied by the proposed IMs, considering dimensional as well as repair-based IMs (cf. Section 3). We also analyze the time complexity of our IMs. Results are summarized in Table 1. Each row in Table 1 but the last one refers to a postulate and shows which IMs satisfy it. The first 8 postulates are classical ones, while the last 3 are dimensional. The first 5 IMs (i.e., columns) are dimensional ones, while the remaining 4 are repair-based.

Classical Postulate Satisfaction. We first consider *dimensional IMs*. We note that $\mathcal{I}_O$, $\mathcal{I}_T$, and $\mathcal{I}_{OT}$ satisfy the same set of (classical) postulates. The two IMs based on the space dimension behave quite differently: $\mathcal{I}_S$ satisfies only Free-Formula Independence; $\mathcal{I}_D$ satisfies postulates (Penalty, Super-Additivity, and MI-Separability) that are not satisfied by any other dimensional IMs, but it does not satisfy postulates (Equal Conflict and MI-Normalization) that are satisfied by other dimensional IMs.

Let us consider now *repair-based IMs*. We first note that $\mathcal{I}_{id}$ and $\mathcal{I}_{time}$ behave in the same way, satisfying Dominance and Super-Additivity, while $\mathcal{I}_{region}$ satisfies only Super-Additivity. In contrast, the classical postulates fit quite well with $\mathcal{I}_{card}$, 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 IMs like $\mathcal{I}_{card}$. However, we will see below that the measures based on the dimensions satisfy the dimensional versions of some classical postulates they do not satisfy.

Dimensional Postulate Satisfaction. Notice that we do not include $\mathcal{I}_{OT}$ as it deals with two dimensions and we restricted the dimensional postulates to a single dimension. Also, we do not include $\mathcal{I}_{card}$, as it does not deal with dimensions.

We first consider *dimensional IMs.* All dimensional IMs satisfy all dimensional postulates. Considering the dimensional postulates in place of the classical counterpart, $\mathcal{I}_O$ and $\mathcal{I}_T$ satisfy all postulates but Attenuation, thus the highest possible number, since Attenuation is incompatible with MI-Normalization (cf. [Grant *et al.*, 2021]). The number of postulates satisfied by considering dimensional and classical postulates has doubled for all the measures except for $\mathcal{I}_D$.

We now consider *repair-based IMs.* $\mathcal{I}_{id}$ and $\mathcal{I}_{time}$ satisfy the dimensional postulates except Dimensional MI-Separability, with Dimensional Penalty being satisfied when the IM returns a finite value. Then, $\mathcal{I}_{region}$ satisfies all dimensional postulates unconditionally. Also, if the metric used for measuring region distance is the symmetric difference, $\mathcal{I}_{card}$ coincides with $\mathcal{I}_{region}$ (cf. [Grant *et al.*, 2021]), and 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.

Overall, our analysis shows that the dimensional postulates suit very well the dimensional inconsistency measures.

Complexity Analysis. Following [Thimm and Wallner, 2016; Thimm, 2018], we characterize the complexity of the following decision problem for an IM $\mathcal{I}$: given an ST database $\mathcal{S}$ and rational number k, decide whether $\mathcal{I}(\mathcal{S}) \leq k$.

The complexity of such problem for our IMs is reported in (the last row of) Table 1; more specifically: it is in P for $\mathcal{I}_O$, $\mathcal{I}_T$, $\mathcal{I}_{OT}$, and $\mathcal{I}_{card}$; it is in P for $\mathcal{I}_S$ and $\mathcal{I}_D$ under the restriction of *isothetic rectangular regions*, by which we mean rectangles oriented in the standard way, that is, with horizontal and vertical sides; it is in P for $\mathcal{I}_{region}$ when the symmetric difference is employed as the metric to measure distance between regions; finally, it is NP-complete for $\mathcal{I}_{id}$ and $\mathcal{I}_{time}$.

References

[Agarwal *et al.*, 2010] Deepak Agarwal, Datong Chen, Long ji Lin, Jayavel Shanmugasundaram, and Erik Vee. Forecasting high-dimensional data. In *Proc. SIGMOD*, pages 1003–1012, 2010.

[Akdere *et al.*, 2011] Mert Akdere, Ugur Cetintemel, Matteo Riondato, Eli Upfal, and Stanley B. Zdonik. The case for predictive database systems: Opportunities and challenges. In *Proc. CIDR*, pages 167–174, 2011.

[Cohn and Hazarika, 2001] Anthony G. Cohn and Shyamanta M. Hazarika. Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae*, 46(1-2):1–29, 2001.

[Fazzinga *et al.*, 2014] Bettina Fazzinga, Sergio Flesca, Filippo Furfaro, and Francesco Parisi. Offline cleaning of RFID trajectory data. In *Proc. SSDBM*, page 5, 2014.

[Fazzinga *et al.*, 2016] Bettina Fazzinga, Sergio Flesca, Filippo Furfaro, and Francesco Parisi. Exploiting integrity constraints for cleaning trajectories of RFID-monitored objects. *ACM Transactions on Database Systems*, 41(4), 2016.

[Gabelaia *et al.*, 2005] David Gabelaia, Roman Kontchakov, Ágnes Kurucz, Frank Wolter, and Michael Zakharyaschev. Combining spatial and temporal logics: Expressiveness vs. complexity. *Journal of Artificial Intelligence Research*, 23:167–243, 2005.

[Grant *et al.*, 2010] John Grant, Francesco Parisi, Austin Parker, and V. S. Subrahmanian. An agm-style belief revision mechanism for probabilistic spatio-temporal logics. *Artificial Intelligence*, 174(1):72–104, 2010.

[Grant *et al.*, 2013a] John Grant, Cristian Molinaro, and Francesco Parisi. Aggregate count queries in probabilistic spatio-temporal databases. In *Proc. SUM*, pages 255–268, 2013.

[Grant *et al.*, 2013b] John Grant, Francesco Parisi, and V. S. Subrahmanian. Research in probabilistic spatiotemporal databases: The spot framework. In *Advances in Probabilistic Databases for Uncertain Information Management*, volume 304 of *Studies in Fuzziness and Soft Computing*, pages 1–22. Springer, 2013.

[Grant *et al.*, 2017] John Grant, Cristian Molinaro, and Francesco Parisi. Count queries in probabilistic spatio-temporal knowledge bases with capacity constraints. In *Proc. ECSQARU*, volume 10369, pages 459–469, 2017.

[Grant *et al.*, 2018] John Grant, Cristian Molinaro, and Francesco Parisi. Probabilistic spatio-temporal knowledge bases: Capacity constraints, count queries, and consistency checking. *International Journal of Approximate Reasoning*, 100:1–28, 2018.

[Grant *et al.*, 2021] John Grant, Maria Vanina Martinez, Cristian Molinaro, and Francesco Parisi. Dimensional inconsistency measures and postulates in spatio-temporal databases. *Journal of Artificial Intelligence Research*, 71:733–780, 2021.

[Hammel *et al.*, 2003] Tom Hammel, Timothy J. Rogers, and Bernie Yetso. Fusing live sensor data into situational multimedia views. In *Proc. MIS*, pages 145–156, 2003.

[Knapp *et al.*, 2006] Alexander Knapp, Stephan Merz, Martin Wirsing, and Júlia Zappe. Specification and refinement of mobile systems in mtla and mobile uml. *Theoretical Computer Science*, 351(2):184–202, 2006.

[Mittu and Ross, 2003] Ranjeev Mittu and Robert Ross. Building upon the coalitions agent experiment (COAX) - integration of multimedia information in GCCS-M using IMPACT. In *Proc. MIS*, pages 35–44, 2003.

[Parisi and Grant, 2016] Francesco Parisi and John Grant. Knowledge representation in probabilistic spatio-temporal knowledge bases. *Journal of Artificial Intelligence Research*, 55:743–798, 2016.

[Parisi and Grant, 2017] Francesco Parisi and John Grant. On repairing and querying inconsistent probabilistic spatio-temporal databases. *International Journal of Approximate Reasoning*, 84:41–74, 2017.

[Parisi *et al.*, 2010] Francesco Parisi, Austin Parker, John Grant, and V. S. Subrahmanian. Scaling cautious selection in spatial probabilistic temporal databases. In *Methods for Handling Imperfect Spatial Information*, volume 256 of *Studies in Fuzziness and Soft Computing*, pages 307–340. Springer, 2010.

[Parisi *et al.*, 2013] Francesco Parisi, Amy Sliva, and V. S. Subrahmanian. A temporal database forecasting algebra. *International Journal of Approximate Reasoning*, 54(7):827–860, 2013.

[Parker *et al.*, 2009] Austin Parker, Guillaume Infantes, John Grant, and V. S. Subrahmanian. Spot databases: Efficient consistency checking and optimistic selection in probabilistic spatial databases. *IEEE Transactions on Knowledge and Data Engineering*, 21(1):92–107, 2009.

[Southey *et al.*, 2007] Finnegan Southey, Wesley Loh, and Dana F. Wilkinson. Inferring complex agent motions from partial trajectory observations. In *Proc. IJCAI*, pages 2631–2637, 2007.

[Thimm and Wallner, 2016] Matthias Thimm and Johannes Peter Wallner. Some complexity results on inconsistency measurement. In *Proc. KR*, pages 114–124, 2016.

[Thimm, 2018] Matthias Thimm. On the evaluation of inconsistency measures. In John Grant and Maria Vanina Martinez, editors, *Measuring Inconsistency in Information*, volume 73 of *Studies in Logic*, pages 19–60. College Publications, 2018.

[Xiao and Ma, 2012] Guohui Xiao and Yue Ma. Inconsistency measurement based on variables in minimal unsatisfiable subsets. In *Proc. ECAI*, pages 864–869, 2012.

[Yaman *et al.*, 2005] Fusun Yaman, Dana S. Nau, and V. S. Subrahmanian. Going far, logically. In *Proc. IJCAI*, pages 615–620, 2005.