

INFINITE PROBABILISTIC DATABASES *

MARTIN GROHE AND PETER LINDNER

RWTH Aachen University, Ahornstraße 55, 52074 Aachen, Germany
e-mail address: {grohe,lindner}@informatik.rwth-aachen.de

ABSTRACT. Probabilistic databases (PDBs) model uncertainty in data in a quantitative way. In the established formal framework, probabilistic (relational) databases are finite probability spaces over relational database instances. This finiteness can clash with intuitive query behavior (Ceylan et al., KR 2016), and with application scenarios that are better modeled by continuous probability distributions (Dalvi et al., CACM 2009).

We formally introduced infinite PDBs in (Grohe and Lindner, PODS 2019) with a primary focus on countably infinite spaces. However, an extension beyond countable probability spaces raises nontrivial foundational issues concerned with the measurability of events and queries and ultimately with the question whether queries have a well-defined semantics.

We argue that finite point processes are an appropriate model from probability theory for dealing with general probabilistic databases. This allows us to construct suitable (uncountable) probability spaces of database instances in a systematic way. Our main technical results are measurability statements for relational algebra queries as well as aggregate queries and Datalog queries.

1. INTRODUCTION

Probabilistic databases (PDBs) are used to model uncertainty in data. Such uncertainty can have various reasons like, for example, noisy sensor data, the presence of incomplete or inconsistent information, or information gathered from unreliable sources [Agg09, SORK11]. In the standard formal framework, probabilistic databases are finite probability spaces whose sample spaces consist of database instances in the usual sense, referred to as “possible worlds”. However, this framework has various shortcomings due to its inherent *closed-world assumption* [CDVdB16, CDVdB21], and the restriction to finite domains. In particular, any event outside of the finite scope of such probabilistic databases is treated as an impossible event. Yet, statistical models of uncertain data, say, for example, for temperature measurements as in Example 2.1, usually feature the use of continuous probability distributions in appropriate error models. This (continuous attribute-level uncertainty) is not expressible in the traditional PDB model. Finite PDBs also only have limited support for tuple-level uncertainty: in finite PDBs, all possible worlds have a fixed maximum number of tuples. Instead, in particular with respect to an open-world assumption, we would like

Key words and phrases: Probabilistic Databases, Possible Worlds Semantics, Query Measurability, Relational Algebra, Aggregation, Infinite Probabilistic Databases.

* Article version of the paper *Infinite Probabilistic Databases* [GL20] presented at ICDT 2020.

to be able to model probabilistic databases without an a priori bound on the number of tuples per instance. It is worth noting that there have been a number of approaches to PDB systems that are supporting continuous probability distributions, and hence going beyond finite probability spaces (see the related works section). These models, however, lack a general (unifying) formal basis in terms of a possible worlds semantics [DRS09]. While both open-world PDBs and continuous probability distributions in PDBs have received some attention in the literature, there is no systematic joint treatment of these issues with a sound theoretical foundation. In [GL19], we introduced an extended model of PDBs as arbitrary (possibly infinite) probability spaces over finite database instances. However, the focus there was on countably infinite PDBs. An extension to continuous PDBs, which is necessary to model probability distributions appearing in many applications that involve real-valued measurement data, raises new fundamental questions concerning the measurability of events and queries.

In this paper, we lay the foundations of a systematic and sound treatment of infinite, even uncountable, probabilistic databases. In particular, we prove that queries expressed in standard query languages have a well-defined semantics that is compatible with the existing theoretical point of view. Our model is based on the mathematical theory of finite point processes [Moy62, Mac75, DVJ03]. Adopting this theory to the context of relational databases, we give a suitable construction of measurable spaces over which our probabilistic databases can then be defined. The only (and mild) assumption we need is that the domains of all attributes are Polish spaces. Intuitively, this requires them to have nice topological properties. All typical domains one might encounter in database theory, for example integers, strings, and reals, satisfy this assumption.

For queries and views to have a well-defined open-world semantics, we need them to be measurable mappings between probabilistic databases. Our main technical result states that indeed all queries and views that can be expressed in the relational algebra, even equipped with arbitrary aggregate operators (satisfying some mild measurability conditions) are measurable mappings. The result holds for both a bag-based and set-based relational algebra and entails the measurability of Datalog queries.

Measurability of queries may seem like an obvious minimum requirement, but one needs to be very careful. We give an example of a simple, innocent looking “query” that is not measurable (see Example 3.12). The proofs of the measurability results are not trivial, and to our knowledge not immediately covered by standard results from point process theory. At their core, the proofs are based on finding suitable “countable approximations” of the queries. That such approximations can be obtained is guaranteed by our topological requirements.

In the last section of this paper, we briefly discuss queries for probabilistic databases that go beyond those that are just set-based versions of traditional database queries. This also casts other natural PDB queries into our framework. Examples of such a queries are probabilistic threshold queries and rank queries. Note that these examples refer not only to the facts in a database, but also to their probabilities, and hence are inherently probabilistic.

This article is an extended version of the paper with the same title, *Infinite Probabilistic Databases* that was presented at the 23rd International Conference on Database Theory (ICDT 2020) [GL20]. The presentation of various proofs and arguments has been extensively reworked, and the paper contains much more background information and new examples. The overall accessibility of the paper has been additionally enhanced by many notational and structural improvements.

Related Work. Early models for probabilistic databases date back to the 1980s [Won82, GH86, CP87] and 1990s [BGMP92, Pit94, DS96, FR97, Zim97]. Essentially, they are special cases or variations of the now-acclaimed formal model of probabilistic databases as a finite set of database instances (the “possible worlds”) together with a probability distribution [Agg09, SORK11].

The work [Koc08] presents a formal definition of the probabilistic semantics of relational algebra queries which is used in the MayBMS system [Koc09]. A probabilistic semantics for Datalog has already been proposed in the mid-90s [Fuh95]. More recently, a version of Datalog was considered in which rules may fire probabilistically [DKM10]. Aggregate queries in probabilistic databases were first treated systematically in [RSG05] and reappear in various works concerning particular PDB systems [MIW11, FHO12].

The models of possible worlds semantics mentioned above are the mathematical backbone of existing probabilistic database prototype systems such as MayBMS [Koc09], Trio [Wid09] and MystiQ [BDM⁺05]. Various subsequent prototypes feature uncountable domains as well, such as Orion [SMS⁺08], MCDB / SimSQL [JXW⁺11, CVP⁺13], Trio [Wid09, AW09] and PIP [KK10]. The MCDB system in particular allows programmers to specify probabilistic databases with infinitely many possible worlds, and with database instances that can grow arbitrarily large [JXW⁺11]. Its description does not feature a general formal, measure theoretic account of its semantics though. In a spirit similar to our work, the work [TPD⁺12] introduces a measure theoretic semantics for probabilistic data stream systems with probability measures composed from Gaussian mixture models but (to our knowledge) on a per tuple basis and without the possibility of inter-tuple correlations. Continuous probabilistic databases have already been considered earlier in the context of sensor networks [FGB02, CKP03, DGM⁺04]. The first work to formally introduce continuous possible worlds semantics (including aggregation) is [ACK⁺11] for probabilistic XML. However, the framework has an implicit restriction bounding the number of tuples in a PDB.

Models similar in expressivity to the one we present have also been suggested in the context of probabilistic modeling languages and probabilistic programming [MMR⁺05, Mil06, RD06, DRKNP16, Bel20, GHNR14]. In particular notable are the measure theoretic treatments of Bayesian Logic (BLOG) [MMR⁺05] in [WSH⁺18] and Markov Logic Networks (MLNs) [RD06] in [SD07]. While these data models are relational, their suitability for general database applications remains unclear. In particular, the investigation of typical database queries is beyond the scope of the corresponding works. Recently, point processes in particular have been abstractly investigated for probabilistic programming [DS20].

Problems caused by the closed-world assumption [Rei81] in probabilistic databases [ZP97] were initially discussed by Ceylan et al. in [CDVdB16] where they suggest the model of OpenPDBs. In [BCL18], the authors make a more fine-grained distinction between an *open-world* and *open-domain assumption*, the latter of which does not assume the attribute values of the database schema to come from a known finite domain. The semantics of OpenPDBs can be strengthened towards an open-domain assumption by the means of ontologies [BCL17, BCL18, BCL19].

There is a tight connection between the problem of computing the probability of Boolean queries in (finite) probabilistic databases and *weighted first-order model counting (WFOMC)* [VdBS17]. For this connection, the *lineage* of the query in question is considered. Then probabilistic query evaluation in finite PDBs boils down to weighted counting of the models of the lineage. While weighted model counting it itself is restricted to discrete weight functions, *weighted model integration (WMI)* is an extension that also supports uncountable

domains [BPVDB15, MPS19]. However, a direct connection to our infinite PDBs seems hard to obtain, as queries in infinite PDBs in general do not have finite lineage expressions. In particular, our infinite PDBs may have an unbounded number of random tuples, contrasting the setup of a fixed number of variables in a WMI problem. Also, in this paper we introduce a rather general notion of PDB queries, that covers, for example, parity tests and various fixpoint queries.

Our classification of views towards the end of this paper is similar to previous classifications of queries such as [CKP03, WvK15] in the sense that it distinguishes the level on which information is aggregated. The work [WvK15] suggests a distinction between “traditional” and “out-of-world aggregation” similar to the one we present.

2. PRELIMINARIES

Throughout this paper, \mathbb{N} , \mathbb{Q} , and \mathbb{R} denote the sets of non-negative integers, rationals, and real numbers, respectively. With \mathbb{N}_+ , \mathbb{Q}_+ , and \mathbb{R}_+ we denote their restrictions to positive numbers.

Sets and Bags. If S is a set and $k \in \mathbb{N}$, then $\mathcal{P}_k(S)$ denotes the set of all subsets of S that have cardinality exactly k . The set of all *finite* subsets of S is then given as $\mathcal{P}_{\text{fin}}(S) := \bigcup_{i=0}^{\infty} \mathcal{P}_i(S)$. The set of *all* subsets of S (that is, the *powerset* of S) is denoted by $\mathcal{P}(S)$.

A (finite) *bag* (or *multiset*) B over a set S is a function from S to \mathbb{N} , assigning a *multiplicity* to every element from S . We interpret bags as collections of elements that may contain duplicates. For $s \in S$ we let $|B|_s$ denote the multiplicity of s in B . If $\mathcal{S} \subseteq S$, we let $|B|_{\mathcal{S}} := \sum_{s \in \mathcal{S}} |B|_s$. The *cardinality* $|B|$ of the bag B is the sum of all multiplicities, i. e. $|B| := |B|_S$.

Similar to the set notation, we use $\mathcal{B}_k(S)$ to denote the set of all bags over S that have cardinality exactly k . The set of all *finite* bags over S is given as $\mathcal{B}_{\text{fin}}(S) := \bigcup_{i=0}^{\infty} \mathcal{B}_i(S)$. Occasionally, we explicitly denote bags by the elements they contain. Then $\{\{a_1, \dots, a_k\}\}$ denotes a bag of cardinality k with elements a_1, \dots, a_k (possibly including repetitions).

Relational Databases. For the remainder of this paper, we fix two countably infinite sets Rel and Att with $\text{Rel} \cap \text{Att} = \emptyset$. The elements of Rel are called *relation symbols*, and the elements of Att are called *attribute names*.

A *database schema* τ is a tuple $(\mathcal{A}, \mathcal{R}, \text{sort})$ such that \mathcal{A} is a finite subset of Att , \mathcal{R} is a finite subset of Rel , and $\text{sort}: \mathcal{R} \rightarrow \bigcup_{k=0}^{\infty} \mathcal{A}^k$ is a function that maps every relation symbol $R \in \mathcal{R}$ to a tuple (A_1, \dots, A_k) of pairwise distinct attribute names $A_1, \dots, A_k \in \mathcal{A}$ for some $k \in \mathbb{N}$. For $R \in \mathcal{R}$ we call $\text{sort}(R)$ the *sort* of R . If $\text{sort}(R) = (A_1, \dots, A_k)$, then k is called the *arity* of R , and denoted by $\text{ar}(R)$. We abuse notation and write $A \in \text{sort}(R)$ if A is an attribute name appearing in $\text{sort}(R)$. We also write $R \in \tau$ instead of $R \in \mathcal{R}$ for relation symbols.

Let \mathcal{A} be a subset of Att . A (*sorted*) *universe* with sorts \mathcal{A} is a pair (\mathbb{U}, dom) where \mathbb{U} is a non-empty set and $\text{dom}: \mathcal{A} \rightarrow \mathcal{P}(\mathbb{U})$. We abuse notation and refer a sorted universe (\mathbb{U}, dom) just by \mathbb{U} , and implicitly assume dom given. For $A \in \mathcal{A}$, we call $\text{dom}(A) =: \mathbb{A}$ the

(*attribute*) *domain* of A . If $\tau = (\mathcal{A}, \mathcal{R}, \text{sort})$ is a database schema and \mathbb{U} a universe with sorts \mathcal{A} and $R \in \mathcal{R}$ with $\text{sort}(R) = (A_1, \dots, A_k)$, then

$$\mathbb{T}_{[R, \mathbb{U}]} := \mathbb{A}_1 \times \dots \times \mathbb{A}_k \quad (2.1)$$

is called the *domain* of R . The elements of $\mathbb{T}_{[R, \mathbb{U}]}$ are called *R-tuples* (over \mathbb{U}). If R is a relation of sort $\text{sort}(R) = (A_1, \dots, A_k)$, and $t = (a_1, \dots, a_k)$ is an *R-tuple*, then

$$t[A_{i_1}, \dots, A_{i_\ell}] := (a_{i_1}, \dots, a_{i_\ell})$$

denotes the restriction of t to the attributes $A_{i_1}, \dots, A_{i_\ell}$ for all $1 \leq i_1 < \dots < i_\ell \leq k$ and all $\ell = 1, \dots, k$.

For $R \in \mathcal{R}$, an *R-fact* (over \mathbb{U}) is an expression of the shape $R(t)$ where t is an *R-tuple* over \mathbb{U} . If $\mathbf{T} \subseteq \mathbb{T}_{[R, \mathbb{U}]}$ is a set of *R-tuples* over \mathbb{U} , then

$$R(\mathbf{T}) := \{R(t) : t \in \mathbf{T}\}.$$

In particular, $R(\mathbb{T}_{[R, \mathbb{U}]})$ is the set of all *R-facts* over \mathbb{U} , which is formally given as

$$\mathbb{F}_{[R, \mathbb{U}]} := \{R\} \times \mathbb{T}_{[R, \mathbb{U}]} \quad (2.2)$$

The elements of

$$\mathbb{F}_{[\tau, \mathbb{U}]} := \bigcup_{R \in \tau} \mathbb{F}_{[R, \mathbb{U}]}$$

are called τ -*facts* (or just *facts*) over \mathbb{U} . We use (variants of) the letter f to denote facts.

A *database instance* D over τ and \mathbb{U} , or τ -*instance* over \mathbb{U} is a finite bag of τ -facts over \mathbb{U} . Thus,

$$\mathbb{DB}_{[\tau, \mathbb{U}]} := \mathcal{B}_{\text{fin}}(\mathbb{F}_{[\tau, \mathbb{U}]})$$

is the set of all τ -instances over \mathbb{U} . For $R \in \tau$ and $D \in \mathbb{DB}_{[\tau, \mathbb{U}]}$, we let $R(D)$ denote the restriction of D to *R-facts*. That is, the instance $R(D)$ is given by

$$|R(D)|_f := \begin{cases} |D|_f & \text{if } f \in \mathbb{F}_{[R, \mathbb{U}]} \text{ and} \\ 0 & \text{if } f \notin \mathbb{F}_{[R, \mathbb{U}]} \end{cases}$$

for all $f \in \mathbb{F}_{[\tau, \mathbb{U}]}$.

From all of the notation we defined above, we omit the explicit mention of τ and \mathbb{U} if they are clear from the context, and just write \mathbb{T} , \mathbb{T}_R , \mathbb{F} , \mathbb{F}_R and \mathbb{DB} instead of $\mathbb{T}_{[\tau, \mathbb{U}]}$, $\mathbb{T}_{[R, \mathbb{U}]}$, $\mathbb{F}_{[\tau, \mathbb{U}]}$, $\mathbb{F}_{[R, \mathbb{U}]}$ and $\mathbb{DB}_{[\tau, \mathbb{U}]}$. Table 1 recaps the notation we just defined.

TABLE 1. Notation used for relational databases.

Concept	Notation
Database schema	$\tau = (\mathcal{A}, \mathcal{R}, \text{sort})$ with $\mathcal{A} \subseteq \text{Att}$, $\mathcal{R} \subseteq \text{Rel}$
Attribute names	$A, A_1, A_2, \dots \in \mathcal{A}$
Relation names	$R, S, R_1, R_2, \dots \in \mathcal{R}$
Sort of a relation name $R \in \mathcal{R}$	$\text{sort}(R) \in \bigcup_{k \in \mathbb{N}} \mathcal{A}^k$
Arity of a relation name $R \in \mathcal{R}$	$\text{ar}(R) \in \mathbb{N}$
Underlying universe	\mathbb{U}
Space of all tuples over (τ, \mathbb{U})	$\mathbb{T}_{[\tau, \mathbb{U}]}, \mathbb{T}$
Space of all R -tuples over (τ, \mathbb{U}) , $R \in \mathcal{R}$	$\mathbb{T}_{[R, \mathbb{U}]}, \mathbb{T}_R$
Tuples and restricted tuples	t, t_1, t_2, \dots and $t[A_{i_1}, \dots, A_{i_k}]$
Space of all facts over (τ, \mathbb{U})	$\mathbb{F}_{[\tau, \mathbb{U}]}, \mathbb{F}$
Space of all R -facts over (τ, \mathbb{U}) , $R \in \mathcal{R}$	$\mathbb{F}_{[R, \mathbb{U}]}, \mathbb{F}_R$
Facts	f, f_1, f_2, \dots
Space of database instances over (τ, \mathbb{U})	$\mathbb{DB}_{[\tau, \mathbb{U}]}, \mathbb{DB}$

In [GL19], we used an example of temperature recordings to motivate infinite probabilistic databases, to illustrate the setup, and to highlight some observations. With some customization, we adopt this example to also serve as a running example for this paper.

Example 2.1. We consider a database that stores information about the rooms in an office building. An example database instance is shown in Figure 1.

OFFICE		TEMPREC		
RoomNo	Person	RoomNo	Date	Temp [°C]
4108	Alice	4108	2021-07-12	20.5
4108a	Bob	4108a	2021-07-12	21.0
4109	Charlie	4109	2021-07-12	21.1
		4109	2021-07-13	20.8

FIGURE 1. A database instance D , storing temperature recordings in an office building.

The database schema τ consists of

- relation symbols $\mathcal{R} = \{\text{OFFICE}, \text{TEMPREC}\}$,
- attribute names $\mathcal{A} = \{\text{RoomNo}, \text{Person}, \text{Date}, \text{Temp}\}$, and
- sorts $\text{sort}(\text{OFFICE}) = (\text{RoomNo}, \text{Person})$ and $\text{sort}(\text{TEMPREC}) = (\text{RoomNo}, \text{Date}, \text{Temp})$.

The sorted universe is given by (\mathbb{U}, dom) with $\mathbb{U} = \Sigma^* \cup \mathbb{R}$ where Σ is some alphabet, say, the set of ASCII symbols, and (for simplicity) $\text{dom}(\text{RoomNo}), \text{dom}(\text{Person}), \text{dom}(\text{Date}) \subseteq \Sigma^*$ and $\text{dom}(\text{TempRec}) = \mathbb{R}$.

For example facts $f_1 = \text{OFFICE}(4108, \text{Bob})$ and $f_2 = \text{TEMPREC}(4108, 2021-07-12, 20.5)$ are both facts from $\mathbb{F}_{[\tau, \mathbb{U}]}$, and it holds that $f_1 \notin D$ and $f_2 \in D$ for the database instance D from Figure 1.

Measurable Spaces and Functions. In the following we recap the relevant background and foundations from probability and measure theory that are needed throughout this paper. For a general introduction to measure theory we point the reader to [Kle14, Chapter 1] or [Kal02, Chapter 1].

Let $\mathbb{X} \neq \emptyset$. A σ -algebra on \mathbb{X} is a family \mathfrak{X} of subsets of \mathbb{X} such that $\mathbb{X} \in \mathfrak{X}$ and \mathfrak{X} is closed under complements and countable unions (an equivalent definition is obtained by replacing “countable unions” with “countable intersections”).

Notation 2.2. Throughout the paper, we use double struck letters ($\mathbb{X}, \mathbb{Y}, \mathbb{Z}, \mathbb{A}, \mathbb{B}, \dots$) to denote underlying spaces of interest, and fraktur letters ($\mathfrak{X}, \mathfrak{Y}, \mathfrak{Z}, \mathfrak{A}, \mathfrak{B}, \dots$) to denote set families and σ -algebras over such spaces, in particular. We usually denote subsets of the underlying spaces with bold italics letters ($\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{A}, \mathbf{B}, \dots$).

Let \mathfrak{G} (fraktur “ G ”) be a set of subsets of \mathbb{X} (that is, $\mathfrak{G} \subseteq \mathcal{P}(\mathbb{X})$). The σ -algebra generated by \mathfrak{G} is the coarsest (i.e. smallest with respect to set inclusion) σ -algebra \mathfrak{X} containing \mathfrak{G} . For any $\mathfrak{G} \subseteq \mathcal{P}(\mathbb{X})$, the σ -algebra generated by \mathfrak{G} is unique. A *measurable space* is a pair $(\mathbb{X}, \mathfrak{X})$ where \mathbb{X} is an arbitrary set and \mathfrak{X} is a σ -algebra on \mathbb{X} . The sets in \mathfrak{X} are called \mathfrak{X} -measurable (or *measurable*, if \mathfrak{X} is clear from the context). If $(\mathbb{Y}, \mathfrak{Y})$ is another measurable space, then a function $\varphi: \mathbb{X} \rightarrow \mathbb{Y}$ is called $(\mathfrak{X}, \mathfrak{Y})$ -measurable (or simply *measurable*, if \mathfrak{X} and \mathfrak{Y} are clear from the context), if for all $\mathbf{Y} \in \mathfrak{Y}$ it holds that

$$\varphi^{-1}(\mathbf{Y}) = \{X \in \mathbb{X}: \varphi(X) \in \mathbf{Y}\} \in \mathfrak{X}.$$

The following simple properties are needed throughout the paper.

Fact 2.3 (cf. [Kal02, Lemmas 1.4 and 1.7]). Let $(\mathbb{X}, \mathfrak{X})$, $(\mathbb{Y}, \mathfrak{Y})$, $(\mathbb{Z}, \mathfrak{Z})$ be measurable spaces.

- (1) Suppose \mathfrak{Y} is generated by some $\mathfrak{G} \subseteq \mathcal{P}(\mathbb{Y})$. If $\varphi: \mathbb{X} \rightarrow \mathbb{Y}$ satisfies $\varphi^{-1}(\mathbf{G}) \in \mathfrak{X}$ for all $\mathbf{G} \in \mathfrak{G}$, then φ is measurable.
- (2) Compositions of measurable functions are measurable. That is, if $\varphi: \mathbb{X} \rightarrow \mathbb{Y}$ and $\psi: \mathbb{Y} \rightarrow \mathbb{Z}$ are measurable, then $\psi \circ \varphi: \mathbb{X} \rightarrow \mathbb{Z}$ is $(\mathfrak{X}, \mathfrak{Z})$ -measurable. \square

Probability and Image Measures. A *probability measure* on a measurable space $(\mathbb{X}, \mathfrak{X})$ is a countably additive function $P: \mathfrak{X} \rightarrow [0, 1]$ with $P(\mathbb{X}) = 1$. Then $(\mathbb{X}, \mathfrak{X}, P)$ is called a *probability space*. Measurable sets in probability spaces are also called *events*. A measurable function φ from a probability space $(\mathbb{X}, \mathfrak{X}, P)$ into a measurable space $(\mathbb{Y}, \mathfrak{Y})$ introduces a new probability measure P' on $(\mathbb{Y}, \mathfrak{Y})$ by

$$P'(\mathbf{Y}) = P(\varphi^{-1}(\mathbf{Y})) = P(\{X \in \mathbb{X}: \varphi(X) \in \mathbf{Y}\})$$

for all $\mathbf{Y} \in \mathfrak{Y}$. Then $P' = P \circ \varphi^{-1}$ is called the *image* or *push-forward probability measure* of P under (or along) φ . The measurability of φ ensures that P' is well-defined.

Standard Constructions. Throughout this paper, we will encounter a variety of measurable spaces. The most basic of these are either well-known measurable spaces like countable spaces with the powerset σ -algebra or uncountable spaces like \mathbb{R} with its Borel σ -algebra. From these spaces, we then construct more complicated measurable spaces with the use of the following standard constructions.

- (1) **Product spaces.** For $i = 1, \dots, n$, let $(\mathbb{X}_i, \mathfrak{X}_i)$ be a measurable space. The *product σ -algebra*

$$\bigotimes_{i=1}^n \mathfrak{X}_i = \mathfrak{X}_1 \otimes \dots \otimes \mathfrak{X}_n$$

is the σ -algebra on $\prod_{i=1}^n \mathbb{X}_i$ that is generated by the sets $\{\text{proj}_j^{-1}(\mathbf{X}) : \mathbf{X} \in \mathfrak{X}_j\}$ with $j = 1, \dots, n$ where proj_j is the canonical projection $\text{proj}_j : \prod_{i=1}^n \mathbb{X}_i \rightarrow \mathbb{X}_j$. If $(\mathbb{X}_i, \mathfrak{X}_i) = (\mathbb{X}, \mathfrak{X})$ for all $i = 1, \dots, n$, we also write $\mathfrak{X}^{\otimes n}$ instead of $\bigotimes_{i=1}^n \mathfrak{X}$.

- (2) **Disjoint unions.** If the spaces \mathbb{X}_i are pairwise disjoint, then the *disjoint union σ -algebra* on $\bigcup_{i=1}^n \mathbb{X}_i$ is given by

$$\bigoplus_{i=1}^n \mathfrak{X}_i := \{ \mathbf{X} \subseteq \bigcup_{i=1}^n \mathbb{X}_i : \mathbf{X} \cap \mathbb{X}_i \in \mathfrak{X}_i \text{ for all } i = 1, \dots, n \}.$$

It may be easily verified that this is a σ -algebra on $\bigcup_{i=1}^n \mathbb{X}_i$.

- (3) **Subspaces.** Let $(\mathbb{X}, \mathfrak{X})$ be a measurable space and $\mathbf{X} \in \mathfrak{X}$. Then

$$\mathfrak{X}|_{\mathbf{X}} := \{ \mathbf{X}' \cap \mathbf{X} : \mathbf{X}' \in \mathfrak{X} \}$$

is a σ -algebra on \mathbf{X} .

Polish Topological Spaces. There are deep connections between measure theory and general topology. In fact, virtually everything we discuss and show in this paper relies on the presence of certain topological properties. The central notion from topology we need is that of *Polish spaces* [Kec95, Chapter 3]. In the following, we assume familiarity with the basic topological concepts. A small introduction to the basic terms can be found in Appendix A. For a thorough introduction, we refer to [Wil04].

A *Polish space* is a topological space that is separable (i. e. there exists a countable dense set), and completely metrizable (i. e. there is a complete metric on the space generating its topology). Polish and (the later introduced) standard Borel spaces are introduced and treated in detail in [Kec95, Chapter 1 & 2] and [Sri98]. In this paper, we heavily exploit the properties of Polish spaces. The existence of a countable dense set, together with the existence of a complete metric generating the topology allows us to approximate any point in the space by countable collections of open sets. More specifically, for every point in the space, there exists a sequence over a fixed *countable* set that converges to the point.

The following fact lists a few natural classes of spaces that are Polish.

Fact 2.4. All finite and countably infinite spaces (with the discrete topology) are Polish [Kec95, p. 13]. The Euclidean space \mathbb{R}^n is Polish for all $n \in \mathbb{N}_+$ [Kec95, *ibid.*].

Countable sums and products of sequences of Polish spaces are Polish (subject to the sum and product topologies) [Kec95, Proposition 3.3]. Moreover, subspaces of a Polish space that are countable intersections of open sets are Polish. This also implies that closed subspaces of Polish spaces are Polish, cf. [Kec95, Proposition 3.7 and Theorem 3.11].

The above properties arguably capture all typical spaces of interest for database theory. When we later work with Polish spaces, we always assume that we have a fixed Polish metric at hand, that is, a complete metric generating the Polish topology. With respect to said *compatible* metric, we write $B_\varepsilon(X)$ to denote the open ball of radius $\varepsilon > 0$ around the point $X \in \mathbb{X}$.

Standard Borel Spaces. If $(\mathbb{X}, \mathfrak{D})$ is a topological space, the *Borel σ -algebra* $\mathfrak{Bor}(\mathbb{X})$ on \mathbb{X} is the σ -algebra generated by the open sets \mathfrak{D} (the topology in use is usually clear from context; in our case, provided there is one, we always use a Polish topology on \mathbb{X}). Sets in the Borel σ -algebra are also called *Borel sets*. In the following, we state some nice basic properties of standard Borel spaces.

Fact 2.5 (see [Kal02, Lemmas 1.5 & 1.10]). Let $(\mathbb{X}, \mathfrak{D}_{\mathbb{X}})$ and $(\mathbb{Y}, \mathfrak{D}_{\mathbb{Y}})$ be topological spaces.

- (1) If $\varphi: \mathbb{X} \rightarrow \mathbb{Y}$ is continuous, then φ is $(\mathfrak{Bor}(\mathbb{X}), \mathfrak{Bor}(\mathbb{Y}))$ -measurable.
- (2) If $(\mathbb{Y}, \mathfrak{D}_{\mathbb{Y}})$ is a metric topological space, $(\varphi_n)_{n \geq 0}$ is a sequence of $(\mathfrak{Bor}(\mathbb{X}), \mathfrak{Bor}(\mathbb{Y}))$ -measurable functions $\varphi_n: \mathbb{X} \rightarrow \mathbb{Y}$ with $\lim_{n \rightarrow \infty} \varphi_n = \varphi$, then φ is measurable as well. \square

A measurable space $(\mathbb{X}, \mathfrak{X})$ is called a *standard Borel space* if there exists a Polish topology \mathfrak{D} on \mathbb{X} such that \mathfrak{X} is the (Borel) σ -algebra generated by \mathfrak{D} .

Fact 2.6 (Lusin-Souslin, see [Kec95, Theorem 15.1]). Let $(\mathbb{X}, \mathfrak{X})$ and $(\mathbb{Y}, \mathfrak{Y})$ be standard Borel spaces, let $\mathbf{X} \in \mathfrak{X}$, and let $\varphi: \mathbb{X} \rightarrow \mathbb{Y}$ be a continuous function such that $\varphi|_{\mathbf{X}}$ (the restriction of φ to the set \mathbf{X}) is injective. Then $\varphi(\mathbf{X}) \in \mathfrak{Y}$. \square

In other words, the above fact states that between standard Borel spaces, the image of a measurable set \mathbf{X} under a function is measurable itself, provided that the function is continuous, and injective on \mathbf{X} .

The following fact intuitively states that the property of measurable spaces being standard Borel is closed under using the constructions introduced before.

Fact 2.7 (see [Kec95, p. 75]). Let $(\mathbb{X}_i, \mathfrak{X}_i)$ be a sequence of standard Borel spaces.

- (1) The product space $(\prod_i \mathbb{X}_i, \otimes_i \mathfrak{X}_i)$ is standard Borel.
- (2) If the spaces \mathbb{X}_i are pairwise disjoint, then the disjoint sum $(\sum_i \mathbb{X}_i, \oplus_i \mathfrak{X}_i)$ is standard Borel.

Moreover, if $(\mathbb{X}, \mathfrak{X})$ is standard Borel, and $\mathbf{X} \in \mathfrak{X}$, then $(\mathbf{X}, \mathfrak{X}|_{\mathbf{X}})$ is standard Borel. \square

(Finite) Point Processes. A *point process* [DVJ03, DVJ08] is a probability space over countable sets of points in some abstract “state space” such as the Euclidean space \mathbb{R}^n . Point processes are a well-studied subject in probability theory and they appear in a variety of applications such as particle physics, ecology, geostatistics (cf. [DVJ03, Bad07] and target tracking [Deg17]). In computer science, they have applications in queuing theory [Fra82] and machine learning [KT12]. The individual outcomes (that is, point configurations) drawn from a point process are called a *realization of the process*. A point process is called *finite* if all of its realizations are finite. In the following, we recall the construction of a finite point process over a Polish state space, following the classic constructions of [Moy62, Mac75].

Let $(\mathbb{X}, \mathfrak{X})$ be a standard Borel space. Recall that $(\mathbb{X}^k, \mathfrak{X}^{\otimes k})$ is also standard Borel for all $k \in \mathbb{N}_+$. For all $(X_1, \dots, X_k), (X'_1, \dots, X'_k) \in \mathbb{X}^k$, we define

$$(X_1, \dots, X_k) \sim_k (X'_1, \dots, X'_k)$$

if and only if there exists a permutation π of $\{1, \dots, k\}$ such that

$$(X_1, \dots, X_k) = (X'_{\pi(1)}, \dots, X'_{\pi(k)}).$$

Then \sim_k is an equivalence relation on \mathbb{X}^k . Tuples in \mathbb{X}^k are equivalent under \sim_k if and only if they contain every individual element the exact same number of times. Thus, the quotient

space \mathbb{X}^k/\sim_k can be identified with the set $\mathcal{B}_k(\mathbb{X})$ of bags of cardinality exactly k over \mathbb{X} . Note that in the case $k = 0$, \mathbb{X}^0/\sim_0 (by convention) consists of a single distinguished point representing the empty realization. The space

$$\bigcup_{k=0}^{\infty} \mathbb{X}^k/\sim_k = \bigcup_{k=0}^{\infty} \mathcal{B}_k(\mathbb{X}) = \mathcal{B}_{\text{fin}}(\mathbb{X})$$

is one of the canonical and equivalent choices for the sample space for a finite point process on \mathbb{X} [Mac75]. The original construction of [Moy62] considers the *symmetrization* $\text{sym}: \bigcup_{k=0}^{\infty} \mathbb{X}^k \rightarrow \mathcal{B}_{\text{fin}}(\mathbb{X})$ where for all $(X_1, \dots, X_k) \in \mathbb{X}^k$ it holds that

$$\text{sym}(X_1, \dots, X_k) = [(X_1, \dots, X_k)]_{\sim_k} = \{\{X_1, \dots, X_k\}\}.$$

Then $\bigoplus_{k=1}^{\infty} \mathfrak{X}^{\otimes k}$ is the natural σ -algebra on $\bigcup_{k \in \mathbb{N}} \mathbb{X}^k$. The σ -algebra on $\mathcal{B}_{\text{fin}}(\mathbb{X})$ is then defined as

$$\{\mathbf{B} \subseteq \mathcal{B}_{\text{fin}}(\mathbb{X}) : \text{sym}^{-1}(\mathbf{B}) \in \bigoplus_{k=0}^{\infty} \mathfrak{X}^{\otimes k}\}, \quad (2.3)$$

that is, as the σ -algebra on $\mathcal{B}_{\text{fin}}(\mathbb{X})$ induced by $\text{sym}: \bigcup_{k=0}^{\infty} \mathbb{X}^k \rightarrow \mathcal{B}_{\text{fin}}(\mathbb{X})$. Note that this is indeed a σ -algebra on $\mathcal{B}_{\text{fin}}(\mathbb{X})$, cf. [Kal02, Lemma 1.3].

Remark 2.8 (cf. [Moy62, Section 2]). A set $\mathbf{X} \in \mathfrak{X}^{\otimes k}$ is called *symmetric*, if for all $(X_1, \dots, X_k) \in \mathbb{X}^k$ it holds that $(X_1, \dots, X_k) \in \mathbf{X}$ implies $(X_{\pi(1)}, \dots, X_{\pi(k)}) \in \mathbf{X}$ for all permutations π of $\{1, \dots, k\}$. The function $\mathbf{X} \mapsto \text{sym}(\mathbf{X})$ is a bijection between the symmetric sets \mathbf{X} in $\bigoplus_{k=0}^{\infty} \mathfrak{X}^{\otimes k}$ and the sets of (2.3).

An equivalent, but technically more convenient construction is motivated by interpreting point processes as *random counting measures* [DVJ08]. For $\mathbf{X} \in \mathfrak{X}$ and $n \in \mathbb{N}$, let

$$\#(\mathbf{X}, n) := \{B \in \mathcal{B}_{\text{fin}}(\mathbb{X}) : |B|_{\mathbf{X}} = \sum_{X \in \mathbf{X}} |B|_X = n\},$$

that is, $\#(\mathbf{X}, n)$ is the set of finite bags over \mathbb{X} containing exactly n elements from \mathbf{X} , counting multiplicities. The sets $\#(\mathbf{X}, n)$ with $\mathbf{X} \in \mathfrak{X}$ and $n \in \mathbb{N}$ are called *counting events*. The *counting σ -algebra* $\mathfrak{C}\text{ount}(\mathbb{X})$ on \mathbb{X} is the σ -algebra generated by all counting events.

Fact 2.9 ([Moy62, Theorem 3.2]). The σ -algebra from (2.3) coincides with $\mathfrak{C}\text{ount}(\mathbb{X})$. \square

Definition 2.10. A *finite point process* with standard Borel state space $(\mathbb{X}, \mathfrak{X})$ is a probability space $(\mathcal{B}_{\text{fin}}(\mathbb{X}), \mathfrak{C}\text{ount}(\mathbb{X}), P)$. A finite point process is *simple* if $P(\mathcal{P}_{\text{fin}}(\mathbb{X})) = 1$, that is, if its realizations are almost surely sets.

3. PROBABILISTIC DATABASES

In this section, we introduce our framework for infinite probabilistic databases and their query semantics. To begin with, we recall the conventional formal definition of probabilistic databases as it is found in textbooks on the subject [SORK11, VdBS17]:

Definition 3.1 (Finite Probabilistic Databases; adapted from [SORK11, Section 2.2]). Let τ be a database schema and let \mathbb{U} be a universe. A *probabilistic database (PDB)* \mathcal{D} over τ and \mathbb{U} is a probability space $\mathcal{D} = (\mathbb{DB}, \mathcal{P}(\mathbb{DB}), P)$ where \mathbb{DB} is a finite set of database instances over τ and \mathbb{U} .

From a probability theoretic point of view, Definition 3.1 is a severe limitation to the vast expressive power of stochastic models, only allowing probability distributions over finitely many alternative database instances. As an example, in a setting such as that of Example 2.1 (the database of temperature measurements), we would typically model noise or uncertainty in the sensor measurements by continuous distributions. This directly leads to (even uncountably) infinite probability spaces that are not covered by the typical textbook definition of probabilistic databases (Definition 3.1). In [GL19], we introduced the following general notion of probabilistic databases as probability spaces of database instances:

Definition 3.2 (see [GL19, Definition 3.1]). Let τ be database schema and let \mathbb{U} be a universe. Suppose that \mathfrak{F} is a σ -algebra on the space $\mathbb{F}_{[\tau, \mathbb{U}]}$ of all (τ, \mathbb{U}) -facts. A *probabilistic database (PDB)* \mathcal{D} over τ and \mathbb{U} is a probability space $\mathcal{D} = (\mathbb{DB}, \mathfrak{DB}, P)$ where $\mathbb{DB} \subseteq \mathbb{DB}_{[\tau, \mathbb{U}]}$ and \mathfrak{DB} is a σ -algebra on \mathbb{DB} satisfying

$$\{D \in \mathbb{DB} : |D|_{\mathbf{F}} > 0\} \in \mathfrak{DB}$$

for all $\mathbf{F} \in \mathfrak{F}$.

Therein, the sample space \mathbb{DB} may be infinite, even uncountable. While we only discussed set PDBs in [GL19], we broaden the definition to support bag instances here. In any case, it was left open in [GL19], how to construct such probability spaces, let alone how to obtain suitable measurable spaces $(\mathbb{DB}, \mathfrak{DB})$. This is no longer a trivial task once \mathbb{DB} is uncountable. In this section, we provide a general construction for such measurable spaces.

Remark 3.3. At this point, we want to stress again the meaning of our terminology. In general probabilistic databases there are lots of components, on different levels of abstraction that could, in principle, be infinite spaces. The term *infinite* probabilistic databases is derived from the notion of an *infinite* probability space, meaning that the sample space is of infinite size. Still, in the PDBs we consider in this paper, database instances themselves (the concrete outcomes, or realizations of a PDB) are always finite collections of facts. The framework we describe here is not suitable for discussing probability spaces over infinite database instances in the sense of [AHV95, Section 5.6], such as *constraint databases* [KLP00].

3.1. Probabilistic Databases are Finite Point Processes. Our construction of measurable spaces of database instances stems from interpreting a probabilistic database as a finite point process over the space of possible facts.

Convention 3.4. *From now on, we only consider database schemas τ in combination with sorted universes \mathbb{U} where all attribute domains are Polish.*

We first construct the measurable space of facts over τ and \mathbb{U} . For all $A \in \mathcal{A}$, by assumption, the domain $\mathbb{A} := \text{dom}(A)$ is Polish. Equipping it with its Borel σ -algebra \mathfrak{A} thus yields a standard Borel space $(\mathbb{A}, \mathfrak{A})$. Now let $R \in \tau$ be a relation symbol with $\text{sort}(R) = (A_1, \dots, A_k)$, so that the standard Borel spaces belonging to the attribute names A_1, \dots, A_k are $(\mathbb{A}_1, \mathfrak{A}_1), \dots, (\mathbb{A}_k, \mathfrak{A}_k)$. Recall from (2.1), that the set \mathbb{T}_R of R -tuples is the product of the attribute domains in R . Naturally, \mathbb{T}_R is equipped with its product σ -algebra

$$\mathfrak{T}_R := \bigotimes_{i=1}^k \mathfrak{A}_i. \quad (3.1)$$

Likewise, the set \mathbb{F}_R of R -facts is equipped with the σ -algebra

$$\mathfrak{F}_R := \{\emptyset, \{R\}\} \otimes \mathfrak{T}_R$$

(cf. (2.2)). As the set \mathbb{F} of all facts (over τ and \mathbb{U}) is the disjoint union of all the \mathbb{F}_R , it is naturally equipped with the disjoint union σ -algebra

$$\mathfrak{F} := \bigoplus_{R \in \tau} \mathfrak{F}_R.$$

Example 3.5. We reenact these definitions for our running example (Example 2.1). Recall that the attribute names $A \in \{\text{RoomNo}, \text{Person}, \text{Date}\}$, have domain Σ^* for some (finite) alphabet Σ . As Σ^* is countably infinite, we equip this space with its powerset σ -algebra, $\mathcal{P}(\Sigma^*)$. For the attribute TempRec, we assume the measurable space to be $(\mathbb{R}, \mathfrak{Bor}(\mathbb{R}))$, where $\mathfrak{Bor}(\mathbb{R})$ is the Borel σ -algebra on \mathbb{R} . Given the sorts of the relation names OFFICE and TEMPREC, we have the tuple spaces $\mathbb{T}_{\text{OFFICE}} = \Sigma^* \times \Sigma^*$ and $\mathbb{T}_{\text{TEMPREC}} = \Sigma^* \times \Sigma^* \times \mathbb{R}$. Following (3.1), the corresponding σ -algebras are

$$\mathfrak{T}_{\text{OFFICE}} = \mathcal{P}(\Sigma^*) \otimes \mathcal{P}(\Sigma^*)$$

and

$$\mathfrak{T}_{\text{TEMPREC}} = \mathcal{P}(\Sigma^*) \otimes \mathcal{P}(\Sigma^*) \otimes \mathfrak{Bor}(\mathbb{R}).$$

That is, $\mathfrak{T}_{\text{OFFICE}}$ is the σ -algebra generated by all events of the shape $L_1 \times L_2$ where $L_1, L_2 \subseteq \Sigma^*$, and $\mathfrak{T}_{\text{TEMPREC}}$ is the σ -algebra generated by all events of the shape $L_1 \times L_2 \times B$ where $L_1, L_2 \subseteq \Sigma^*$ and B is a Borel set in \mathbb{R} . The spaces of OFFICE- and TEMPREC-facts are given by

$$\mathbb{F}_{\text{OFFICE}} = \text{OFFICE}(\mathbb{T}_{\text{OFFICE}}) = \{\text{OFFICE}(r, s) : r, s \in \Sigma^*\}$$

and

$$\mathbb{F}_{\text{TEMPREC}} = \text{TEMPREC}(\mathbb{T}_{\text{TEMPREC}}) = \{\text{TEMPREC}(r, d, \theta) : r, d \in \Sigma^* \text{ and } \theta \in \mathbb{R}\}.$$

For example, (assuming that the alphabet Σ contains the respective symbols), it holds that $\text{OFFICE}(4108, \text{Bob}) \in \mathbb{F}_{\text{OFFICE}}$, and $\text{TEMPREC}(4108, 2021-07-12, 20.5) \in \mathbb{F}_{\text{TEMPREC}}$. The σ -algebras on these spaces of facts are directly obtained from the σ -algebras on the tuple spaces: $\mathfrak{F}_{\text{OFFICE}} = \mathcal{P}(\{\text{OFFICE}\}) \otimes \mathfrak{T}_{\text{OFFICE}}$ and $\mathfrak{F}_{\text{TEMPREC}} = \mathcal{P}(\{\text{TEMPREC}\}) \otimes \mathfrak{T}_{\text{TEMPREC}}$. Finally,

$$\mathbb{F} = \mathbb{F}_{\text{OFFICE}} \cup \mathbb{F}_{\text{TEMPREC}}$$

is the space of all facts, with σ -algebra

$$\mathfrak{F} = \mathfrak{F}_{\text{OFFICE}} \oplus \mathfrak{F}_{\text{TEMPREC}}.$$

For example, consider the sets

$$\mathbf{F}_1 = \{f \in \mathbb{F}_{\text{OFFICE}} : f = \text{OFFICE}(4108, s) \text{ for some } s \in \Sigma^*\}$$

and

$$\mathbf{F}_2 = \{f \in \mathbb{F}_{\text{TEMPREC}} : f = \text{TEMPREC}(4108, d, \theta) \text{ for some } d \in \Sigma^* \text{ and } \theta \in [21, 23]\}.$$

Then $\mathbf{F}_1 \in \mathfrak{F}_{\text{OFFICE}}$ and $\mathbf{F}_2 \in \mathfrak{F}_{\text{TEMPREC}}$. Note that due to the real interval, \mathbf{F}_2 contains uncountably many facts. As $\mathbf{F}_1 \in \mathfrak{F}_{\text{OFFICE}}$ and $\mathbf{F}_2 \in \mathfrak{F}_{\text{TEMPREC}}$, it holds that $\mathbf{F}_1 \cup \mathbf{F}_2 \in \mathfrak{F}$.

The constructions of the various measurable spaces above all started from standard Borel measurable spaces for the attribute domains, and then used the basic constructions of product and disjoint union measurable spaces. Therefore, Fact 2.7 immediately yields the following statement.

Lemma 3.6. *The spaces $(\mathbb{T}_R, \mathfrak{T}_R)_{R \in \tau}$, $(\mathbb{F}_R, \mathfrak{F}_R)_{R \in \tau}$, and $(\mathbb{F}, \mathfrak{F})$ are standard Borel spaces.*

Recall that $\mathbb{DB} = \mathbb{DB}_{[\tau, \mathbb{U}]} = \mathcal{B}_{\text{fin}}(\mathbb{F})$. We equip \mathbb{DB} with the σ -algebra $\mathfrak{DB} := \mathbf{Count}(\mathbb{F})$ from Section 2.

Definition 3.7. A *standard probabilistic database (standard PDB)* over τ and \mathbb{U} is a probability space $(\mathbb{DB}, \mathfrak{DB}, P)$ with $\mathbb{DB} = \mathcal{B}_{\text{fin}}(\mathbb{F})$ and $\mathfrak{DB} = \mathbf{Count}(\mathbb{F})$.

That is, a standard PDB is a finite point process over the state space $(\mathbb{F}, \mathfrak{F})$. Note that every standard PDB is a PDB in the sense of Definition 3.2. Standard PDBs that are simple are suitable for modeling PDBs with set semantics.

Example 3.8. We continue from Example 3.5. The space \mathbb{DB} of database instances is exactly the space of finite bags over $\mathbb{F} = \mathbb{F}_{\text{OFFICE}} \cup \mathbb{F}_{\text{TEMPREC}}$ with the counting σ -algebra $\mathfrak{DB} = \mathbf{Count}(\mathbb{F})$. This means, for example that the set

$$\#(\mathbf{F}_2, 3) = \{D \in \mathbb{DB} : |D|_{\mathbf{F}_2} = 3\}$$

is measurable (i. e., an event) in $(\mathbb{DB}, \mathfrak{DB})$. This is the set of database instances over \mathbb{F} that contain exactly three facts of the shape $\text{TEMPREC}(4108, d, \theta)$ where d is an arbitrary string and $21 \leq \theta \leq 23$. Note that these three facts need not be distinct, as facts are allowed to be present with multiplicities. A standard probabilistic database of the example schema is just a probability space with underlying measurable space $(\mathbb{DB}, \mathfrak{DB})$. In particular, in these PDBs, events such as $\#(\mathbf{F}_2, 3)$ carry a probability.

Lemma 3.9. *A standard PDB $(\mathbb{DB}, \mathfrak{DB}, P)$ is a standard Borel space itself.*

Proof. In [DVJ08, Proposition 9.1.IV], it is shown that the space of $\mathbb{N} \cup \{\infty\}$ -valued counting measures on a standard Borel space (with the property of being finite on bounded sets) is a Polish space, whose σ -algebra is generated by the functions that map a counting measure μ to $\mu(\mathbf{F}) \in \mathbb{N} \cup \{\infty\}$ for all $\mathbf{F} \in \mathfrak{F}$. Restricting this space to integer-valued counting measures, and equipping it with the corresponding subspace σ -algebra yields a standard Borel space. This space is isomorphic to $(\mathbb{DB}, \mathfrak{DB})$ via the function that maps a counting measure μ to the database instance whose multiplicity mapping is given by μ .¹ \square

Convention 3.10. *From now on, all PDBs we consider are standard PDBs. When we speak of PDBs, it is understood to refer to standard PDBs, unless explicitly stated otherwise.*

3.2. The Possible Worlds Semantics of Queries and Views. *Views* are mappings between database instances. That is, a view V is a function $V : \mathbb{DB}_{[\tau, \mathbb{U}]} \rightarrow \mathbb{DB}_{[\tau', \mathbb{U}]}$ for some database schemas τ and τ' and universes \mathbb{U} and \mathbb{U}' . We call τ the *input*, and τ' the *output schema* of the view V . If τ' consists of a single relation symbol only, we call V a *query*. Queries are typically denoted by Q . Usually, queries and views are given as syntactic expressions in some query language. As usual, we blur the distinction between a query or view, and its syntactic representation. In the following, we let $(\mathbb{DB}, \mathfrak{DB}) = (\mathbb{DB}_{[\tau, \mathbb{U}]}, \mathfrak{DB}_{[\tau, \mathbb{U}]})$, and $(\mathbb{DB}_V, \mathfrak{DB}_V) = (\mathbb{DB}_{[\tau', \mathbb{U}]}, \mathfrak{DB}_{[\tau', \mathbb{U}]})$.

¹Isomorphic here means that this function is bijective and measurable both ways; bijectivity is clear, and being measurable both ways stems from the fact that the generating events of the σ -algebra of the space of integer-valued counting measures on $(\mathbb{F}, \mathfrak{F})$ are identified with the counting events in $(\mathbb{DB}, \mathfrak{DB})$, i. e. the generating events of \mathfrak{DB} .

Let $\mathcal{D} = (\mathbb{DB}, \mathfrak{DB}, P)$ be a PDB and let $V: \mathbb{DB} \rightarrow \mathbb{DB}_V$ be a view. For $\mathbf{D} \subseteq \mathbb{DB}$, the image of \mathbf{D} under V is

$$V(\mathbf{D}) := \{V(D) : D \in \mathbf{D}\} \subseteq \mathbb{DB}_V.$$

If the function V is $(\mathfrak{DB}, \mathfrak{DB}_V)$ -measurable, then V introduces a push-forward measure P_V on $(\mathbb{DB}_V, \mathfrak{DB}_V)$ via

$$P_V(\mathbf{D}) := P(V^{-1}(\mathbf{D})) = P(\{D \in \mathbb{DB} : V(D) \in \mathbf{D}\}) \quad (3.2)$$

for all $\mathbf{D} \in \mathfrak{DB}_V$. In this case

$$V(\mathcal{D}) := (\mathbb{DB}_V, \mathfrak{DB}_V, P_V)$$

is a PDB.

Remark 3.11. The kind of semantics we introduce here is the natural generalization of a standard choice for semantics on probabilistic databases. Conceptually, we consider queries and views that have well-defined semantics on traditional database instances. That is, they get as input a database instance, and as output, produce a new database instance. Such a semantics is lifted to probabilistic databases by applying the query or view on every possible world, and weighting it according to the probability measure of the input PDB. This notion of semantics is commonly called the *possible answer sets semantics* of probabilistic databases [SORK11, Section 2.3.1]. It (or, to be more precise, its discrete version) has previously been also called *possible worlds semantics (of queries)* [DS07], which is the term we prefer to use in this work, as we deem it the natural choice on how to define query (or view) semantics on PDBs that are modelled as a collection of possible worlds with a probability distribution, matching the standard definition of output probabilities of a (measurable) function on a probability space (cf. [Gre09]). Note that strictly speaking, this overloads the term “possible worlds semantics”: in reference to PDBs, “possible worlds semantics” means the definition of PDBs as probability spaces over database instances, whereas in reference to queries or views, it means the definition of the output of a query with respect to an application per possible world, as in (3.2).

For queries, another semantics has been discussed in literature, which was later dubbed the *possible answers semantics* [SORK11, Section 2.3.2]. Under this semantics, the output of a query is the collection of tuples that may appear as an answer to the query (i. e. the tuples that appear in the output possible worlds under the possible worlds semantics), together with their marginal probability. For finite PDBs, this notion makes sense, because the result will be much smaller than a description of the whole output probability under possible worlds semantics. For uncountable infinite PDBs, however, this is not of much use. As soon as continuous probability distributions are involved, we naturally encounter PDBs where the marginal probability of every particular fact (or tuple in the output) may be zero.

We note that to the best of our knowledge, there has been no formal description of these semantics when duplicates are allowed. Note that for *Boolean queries* with set semantics, i. e., queries whose output is either $\{()\}$ (true) or \emptyset (false), both of the above semantics are essentially equivalent: the only possible answer tuple is the empty tuple $()$, and the only possible worlds of the answer are \emptyset and $\{()\}$.

Note that if V fails to be measurable, then (3.2) is not well-defined. In this case, V has no meaningful semantics on probabilistic databases! Thus, discussing the measurability of views and queries is an issue that requires attention. The following example shows that there are inconspicuous, seemingly simple queries that are not measurable.

Example 3.12. Consider $\mathbb{U} = \mathbb{U}' = \mathbb{R}$, together with the database schemas $\tau = \tau'$ consisting of the single, unary relation symbol R with domain $\mathbb{T}_R = \mathbb{R}$ (equipped with the Borel σ -algebra). Let $\mathbf{B} \in \mathfrak{Bot}(\mathbb{R}^2)$. We define a function $Q_{\mathbf{B}}: \mathbb{DB} \rightarrow \mathbb{DB}'$ (where $\mathbb{DB}' = \mathbb{DB}_{[\tau', \mathbb{U}']}$) as follows.

$$Q_{\mathbf{B}}(D) := \begin{cases} D & \text{if } D \text{ is a singleton } \{\{R(x)\}\} \text{ and there exists a } y \in \mathbb{R} \text{ s. t. } (x, y) \in \mathbf{B} \text{ and} \\ \emptyset & \text{otherwise.} \end{cases}$$

Observe that $Q_{\mathbf{B}}^{-1}(\mathbb{DB}') = \{\{\{R(x)\}\}: x \in \text{proj}_1(\mathbf{B})\}$, where

$$\text{proj}_1(\mathbf{B}) = \{x \in \mathbb{R}: \text{there is } y \in \mathbb{R} \text{ s. t. } (x, y) \in \mathbf{B}\}.$$

It is well known that there are Borel sets $\mathbf{B} \subseteq \mathbb{R}^2$ with the property that $\text{proj}_1(\mathbf{B})$ is *not* Borel in \mathbb{R} [Kec95, Theorem 14.2 (Souslin) and Exercise 14.3]. In this case, $Q_{\mathbf{B}}$ is not measurable even though \mathbf{B} is Borel.

3.3. Assembling Views from Queries. A view V with output relations R_1, \dots, R_k can be identified with a set of queries Q_1, \dots, Q_k , one per output relation symbol, such that for all $D \in \mathbb{DB}$ it holds that

$$V(D) = \bigcup_{i=1}^k Q_i(D).$$

Let $(\mathbb{F}_V, \mathfrak{F}_V)$ be the measurable space of facts belonging to $(\mathbb{DB}_V, \mathfrak{DB}_V)$ and let $(\mathbb{F}_{Q_i}, \mathfrak{F}_{Q_i})$ be the measurable space of facts belonging to $(\mathbb{DB}_{Q_i}, \mathfrak{DB}_{Q_i})$. From the above equality, we get that for all $D \in \mathbb{DB}$, all $\mathbf{F} \in \mathfrak{F}_V$, and all $n \in \mathbb{N}$ it holds that $|V(D)|_{\mathbf{F}} = n$ if and only if there exist non-negative integers n_1, \dots, n_k with $n_1 + \dots + n_k = n$ and the property that $|Q_i(D)|_{\mathbf{F} \cap \mathbb{F}_{Q_i}} = n_i$ for all $i = 1, \dots, k$. Note that this condition corresponds to an event given by a countable union of counting events. Thus, we obtain the following.

Lemma 3.13. *The view V is measurable if and only if Q_i is measurable for all $i = 1, \dots, k$.*

By the merit of Lemma 3.13, we only need to discuss the measurability of queries.

4. GENERAL MEASURABILITY CRITERIA

In the remainder of the paper we establish measurability results for various types of queries as they typically appear in database applications. In this section, we set out the general setup of said investigation and introduce some general measurability results that are not yet tailored to specific query languages.

4.1. Setup. Henceforth, we adhere to the following notational conventions when discussing the measurability of a query Q .

Convention 4.1 (Inputs). The input schema of Q is $\tau = (\mathcal{A}, \mathcal{R}, \text{sort})$. We consider input instances over τ , and the sorted universe \mathbb{U} (with all attribute domains Polish).

The associated fact space is denoted as $(\mathbb{F}, \mathfrak{F})$, with subspaces $(\mathbb{F}_R, \mathfrak{F}_R)$ for all $R \in \tau$. The space of R -tuples is given as $(\mathbb{T}_R, \mathfrak{T}_R)$. For all $R \in \tau$, we fix a compatible Polish metric d_R on \mathbb{T}_R and let \mathbb{T}_R^* be a countable, dense set in \mathbb{T}_R . With abuse of notation, we denote the corresponding metric on \mathbb{F}_R by d_R as well. Note that $\mathbb{F}_R^* := \mathbb{T}_R^*$ is a countable dense set in \mathbb{F}_R .

We denote the input (standard) PDB under consideration by $\mathcal{D} = (\mathbb{DB}, \mathfrak{B}, P)$, where $\mathbb{DB} = \mathcal{B}_{\text{fin}}(\mathbb{F})$ and $\mathfrak{B} = \mathbf{Count}(\mathbb{F})$.

Convention 4.2 (Outputs). The output schema of Q is $\tau_Q = (\mathcal{A}_Q, \mathcal{R}_Q, \text{sort}_Q)$ where R_Q is the only relation symbol in \mathcal{R}_Q . We consider output instances over τ_Q , and sorted universe \mathbb{U}_Q (with all attribute domains Polish). The associated fact space is denoted as $(\mathbb{F}_Q, \mathfrak{F}_Q)$. The space of R_Q -tuples is given as $(\mathbb{T}_Q, \mathfrak{T}_Q)$. We fix a compatible Polish metric d_Q on \mathbb{T}_Q , and a countable dense set \mathbb{T}_Q^* in \mathbb{T}_Q . Again d_Q will also denote the corresponding metric on \mathbb{F}_Q , and the set \mathbb{F}_Q^* is countable and dense in \mathbb{F}_Q .

The output measurable space is denoted by $(\mathbb{DB}_Q, \mathfrak{B}_Q)$, where $\mathbb{DB}_Q = \mathcal{B}_{\text{fin}}(\mathbb{F}_Q)$ and $\mathfrak{B}_Q = \mathbf{Count}(\mathbb{F}_Q)$.

Thus, our goal is to show that a given function $Q: \mathbb{DB} \rightarrow \mathbb{DB}_Q$ is $(\mathfrak{B}, \mathfrak{B}_Q)$ -measurable.

Remark 4.3. We have some straightforward measurability criteria from the general properties of measurable functions and the used σ -algebras.

(1) By Fact 2.3(1), to show the measurability of a query Q , it suffices to show that

$$\{D \in \mathbb{DB}: |Q(D)|_{\mathbf{F}} = n\} \in \mathfrak{B}$$

for all $\mathbf{F} \in \mathfrak{F}_Q$ and all $n \in \mathbb{N}$, as the counting events generate the σ -algebra \mathfrak{B}_Q . This remains true, if we replace “ $= n$ ” with “ $\geq n$ ” or only require $n \in \mathbb{N}_+$.

(2) By Fact 2.3(2), compositions of measurable functions are measurable. That is, for query languages whose semantics are defined inductively over the structure of their syntactic expression, it suffices to show measurability for the basic building blocks.

(3) By Fact 2.5(2), limits of measurable queries are measurable.

4.2. The Mapping Theorem. The following is a partial restatement of what is known as the *mapping theorem* of point processes. The original theorem from point process theory also involves the transfer of certain properties to the image space [LP17, Theorem 5.1] which is, however, of less importance for the remainder of the paper. Moreover, we allow partial transformations as long as their domain is measurable (and this does not invalidate the measurability statement from the mapping theorem).

Theorem 4.4 (cf. [LP17, Theorem 5.1]). *Let $\mathbf{F}_q \in \mathfrak{F}$. If $q: \mathbf{F}_q \rightarrow \mathbb{F}_Q$ is $(\mathfrak{F}, \mathfrak{F}_Q)$ -measurable, then the function $Q: \mathbb{DB} \rightarrow \mathbb{DB}_Q$ with*

$$|Q(D)|_f := \begin{cases} |D|_{q^{-1}(f)} & \text{if } q^{-1}(f) \neq \emptyset \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

for all $D \in \mathbb{DB}$ is a measurable query. □

In this (restricted) form, the theorem is straightforward to verify.

Proof. Let $\mathbf{F}_q \in \mathfrak{F}$ and $q: \mathbf{F}_q \rightarrow \mathbb{F}_Q$ be measurable. Now let $\mathbf{F} \in \mathfrak{F}_Q$ and $n \in \mathbb{N}$. Then

$$|Q(D)|_{\mathbf{F}} = n \iff |D|_{q^{-1}(\mathbf{F})} = n.$$

Since $q^{-1}(\mathbf{F}) \in \mathfrak{F}$, the claim follows. □

Intuitively, the theorem states that whenever we have a measurable transformation of the fact space of a PDB, then we obtain a measurable query when we just apply this transformation to all facts in the database instances.²

Example 4.5. We continue our running example. Recall that $\mathbb{F}_{\text{TEMPREC}}$ is the space of facts $\text{TEMPREC}(r, d, \theta)$ where r and d are strings and θ is a real number. Consider the function $q: \mathbb{F}_{\text{TEMPREC}} \rightarrow \mathbb{F}_{\text{TEMPREC}}$ that increases the temperature by 2 °C, i. e. with

$$q(\text{TEMPREC}(r, d, \theta)) = \text{TEMPREC}(r, d, \theta + 2).$$

Then q is $(\mathfrak{F}_{\text{TEMPREC}}, \mathfrak{F}_{\text{TEMPREC}})$ -measurable. This follows, since the addition of 2 is a continuous function on \mathbb{R} and by the construction of the measurable spaces.

Thus, by Theorem 4.4, the query that, given an instance D , applies q to every fact in D (i. e. increasing temperatures by 2) is measurable.

While Theorem 4.4 is a nice statement, it fails to cover most queries of interest, as database queries often consider and manipulate multiple tuples at once. Such transformations are not captured by Theorem 4.4. We therefore need measurability statements beyond Theorem 4.4.

4.3. Continuous One-to-One Decompositions. In this subsection we introduce a new criterion for query measurability that overcomes the aforementioned limitation of Theorem 4.4.

Lemma 4.6. *Let $Q: \mathbb{DB} \rightarrow \mathbb{DB}_Q$. Then Q is measurable if there exists some $k \in \mathbb{N}_+$, pairwise distinct $R_1, \dots, R_k \in \tau$ and functions $q_i: \mathbb{F}_Q \rightarrow \mathbb{F}_{R_i}$ for $i = 1, \dots, k$ with the following properties:*

(1) *For all $n \in \mathbb{N}_+$ there is a set $N_Q(n) \in \mathbb{N}^k \setminus \{(0, \dots, 0)\}$ with*

$$|Q(D)|_f = n \iff (|D|_{q_1(f)}, \dots, |D|_{q_k(f)}) \in N_Q(n)$$

for all $D \in \mathbb{DB}$ and all $f \in \mathbb{F}_Q$.

(2) *For all $i = 1, \dots, k$, the function q_i is injective and continuous.*

Think of the functions q_i as providing a decomposition of R_Q -facts into the R_i -facts of the input instance they originated from under the query. The set $N_Q(n)$ provides the recipe, how the number of occurrences of a fact in the output is determined by the counts of its decompositions in the input. That is, requirement (1) stipulates the query semantics. The topological requirement (2) ensures measurability.

A simple example application that we ask the reader to have in mind is that of a difference operator on bag instances (this is in fact, one of our later applications, where it is made precise). Intuitively, if Q is the difference of relations R and S , then the number of times a tuple t occurs in the output is given by the number of times it appears in R , minus the number of times it appears in S . This fits the pattern of Lemma 4.6 with functions q_1, q_2 being the identity $t \mapsto t$, and $N_Q(n)$ being the set of pairs (n_1, n_2) with $\max(0, n_1 - n_2) = n$. Lemma 4.6 provides a generalization of this setup, allowing for much more general functions q_i , and sets $N_Q(n)$.

²Functions of the shape of Q in Theorem 4.4 are a special case of “mapping constructs” (applying a function to every element of a bag) that can be found in previously considered bag query languages [GM96, LW97].

Remark 4.7. This is vaguely related to the notion of (*why- and how-*)*provenance* of a tuple in the output of a view [GKT07], with the functions q_i providing the “why-information”, and the set $N_Q(n)$ providing the “how”. Lemma 4.6 now only applies to queries inducing a very particular provenance structure (as governed by the q_i and $N_Q(n)$).

Proof of Lemma 4.6. Let $Q: \mathbb{DB} \rightarrow \mathbb{DB}_Q$, let $k \in \mathbb{N}_+$. Let R_1, \dots, R_k be distinct relation symbols in τ . In the following, we write $(\mathbb{F}_i, \mathfrak{F}_i)$ instead of $(\mathbb{F}_{R_i}, \mathfrak{F}_{R_i})$ for all $i = 1, \dots, k$. As required, let $q_i: \mathbb{F}_Q \rightarrow \mathbb{F}_i$ for all $i = 1, \dots, k$ and assume that (1) and (2) hold. We have to show that Q is measurable, which is settled, in particular, by showing that $\{D \in \mathbb{DB}: |Q(D)|_{\mathbf{F}} \geq n\} \in \mathfrak{DB}$ for all $\mathbf{F} \in \mathfrak{F}_Q$ and all $n \in \mathbb{N}_+$.

Let d be a fixed Polish metric on \mathbb{F}_Q generating \mathfrak{F}_Q and let $\mathbb{F}_Q^* \subseteq \mathbb{F}_Q$ be countable and dense in \mathbb{F}_Q . Let $\mathbf{F} \in \mathfrak{F}_Q$ and $n \in \mathbb{N}_+$. We show that for all $D \in \mathbb{DB}$, it holds that $|Q(D)|_{\mathbf{F}} \geq n$ is equivalent to the following condition.

Condition 4.8. For some $\ell \in \mathbb{N}_+$ there exist $n_1, \dots, n_\ell \in \mathbb{N}$ with $n_1 + n_2 + \dots + n_\ell \geq n$, and $(n_{ij})_{i=1, \dots, k; j=1, \dots, \ell}$ with

$$(n_{1j}, \dots, n_{kj}) \in N_Q(n_j) \quad (4.1)$$

for all $j = 1, \dots, \ell$; and there exists $\varepsilon_0 > 0$ such that for all positive $\varepsilon < \frac{\varepsilon_0}{4}$ there are $f_{1,\varepsilon}^*, \dots, f_{\ell,\varepsilon}^* \in \mathbb{F}_Q^*$ with the following properties:

- (a) For all $j, j' = 1, \dots, \ell$ with $j \neq j'$ it holds that $d_Q(f_{j,\varepsilon}^*, f_{j',\varepsilon}^*) > \frac{\varepsilon_0}{2}$.
- (b) For all $i = 1, \dots, k$ and all $j = 1, \dots, \ell$ it holds that $|D|_{q_i(\mathbf{F}) \cap q_i(B_\varepsilon(f_{j,\varepsilon}^*))} = n_{ij}$

An illustration of the situation of Condition 4.8 can be found in Figure 2. We now prove both directions of the claimed equivalence.

\Rightarrow : First assume that $|Q(D)|_{\mathbf{F}} \geq n > 0$ and suppose that f_1, \dots, f_ℓ are (pairwise distinct) facts from \mathbf{F} with the property $|D|_{q_i(f_j)} > 0$ for at least one $i = 1, \dots, k$ for all $j = 1, \dots, \ell$. Such f_1, \dots, f_ℓ , with $\ell \in \mathbb{N}_+$, exist by requirement (1). Let $n_j := |Q(D)|_{f_j}$ for all $j = 1, \dots, \ell$. Then $\sum_{j=1}^{\ell} n_j \geq n$, and, according to (1), $(n_{1j}, \dots, n_{kj}) \in N_Q(n_j)$, where $n_{ij} := |D|_{q_i(f_j)}$ for all $i = 1, \dots, k$ and $j = 1, \dots, \ell$. We choose $\varepsilon_0 > 0$ such that

$$\varepsilon_0 < \min\{d_Q(f_j, f): j = 1, \dots, \ell, f \neq f_j, f \in Q(D)\}$$

and let $0 < \varepsilon < \frac{\varepsilon_0}{4}$. Because \mathbb{F}_Q^* is dense in \mathbb{F}_Q , we can choose $f_{1,\varepsilon}^*, \dots, f_{\ell,\varepsilon}^* \in \mathbb{F}_Q^*$ such that $d_Q(f_j, f_{j,\varepsilon}^*) < \varepsilon < \frac{\varepsilon_0}{4}$ for all $j = 1, \dots, \ell$. Then from the triangle inequality it follows that

$$d_Q(f_{j,\varepsilon}^*, f_{j',\varepsilon}^*) \geq d_Q(f_j, f_{j'}) - d_Q(f_{j,\varepsilon}^*, f_j) - d_Q(f_{j',\varepsilon}^*, f_{j'}) > \varepsilon_0 - 2\frac{\varepsilon_0}{4} = \frac{\varepsilon_0}{2}.$$

In particular, since $\varepsilon < \frac{\varepsilon_0}{4}$, the balls $B_\varepsilon(f_{j,\varepsilon}^*)$ are pairwise disjoint and $f_j \in B_\varepsilon(f_{j,\varepsilon}^*)$ for all $j = 1, \dots, \ell$. Moreover, by the choice of ε_0 , no fact f from $Q(D)$ other than f_1, \dots, f_ℓ appears in the balls $B_\varepsilon(f_{1,\varepsilon}^*), \dots, B_\varepsilon(f_{\ell,\varepsilon}^*)$. Because q_i is injective, it follows that $q_i(f) \notin q_i(B_\varepsilon(f_{j,\varepsilon}^*))$ for all $f \neq f_j$ appearing in $Q(D)$, and all $i = 1, \dots, k$. This means that for all $i = 1, \dots, k$ and $j = 1, \dots, \ell$ it holds that

$$|D|_{q_i(\mathbf{F}) \cap q_i(B_\varepsilon(f_{j,\varepsilon}^*))} = |D|_{q_i(f_j)} = n_{ij}.$$

Thus, Condition 4.8 holds.

\Leftarrow : Suppose that Condition 4.8 holds and let $\ell \in \mathbb{N}_+$, $n_1 + \dots + n_\ell \geq n$ and n_{ij} for $i = 1, \dots, k$ and $j = 1, \dots, \ell$, and ε_0 be given accordingly. Note that for all $\varepsilon < \frac{\varepsilon_0}{4}$, the balls $B_\varepsilon(f_{j,\varepsilon}^*)$ are pairwise disjoint by condition (a). Since q_i is injective, the sets

$q_i(B_\varepsilon(f_{j,\varepsilon}^*))$ are disjoint as well. By (b), for all $i = 1, \dots, k$, and all $j = 1, \dots, \ell$ it holds that

$$|D|_{q_i(\mathbf{F}) \cap q_i(B_\varepsilon(f_{j,\varepsilon}^*))} = n_{ij}.$$

If ε is small enough, then for all n_{ij} with $n_{ij} > 0$ it holds that there exists a single fact $f_{i,j,\varepsilon} \in q_i(\mathbf{F}) \cap q_i(B_\varepsilon(f_{j,\varepsilon}^*))$ with the property that $|D|_{f_{i,j,\varepsilon}} = n_{ij}$. To see this, note that if the fact count n_{ij} were distributed over multiple facts, then by the injectivity of the functions q_i , these facts have pairwise different preimages under q_i . For small enough ε , however, $B_\varepsilon(f_{j,\varepsilon}^*)$ can only contain one of these facts, yielding a contradiction with (b). We now fix ε small enough such that the above holds, as well as an index $j = 1, \dots, \ell$. Moreover, we pick some $i(j) \in \{1, \dots, k\}$ with $n_{i(j),j} > 0$ and define $f_j := q_{i(j)}^{-1}(f_{i(j),j,\varepsilon})$. As $q_{i(j)}$ is injective, it follows that $f_j \in \mathbf{F} \cap B_\varepsilon(f_{j,\varepsilon}^*)$. Without loss of generality, we assume that $f_j \in B_{\varepsilon'}(f_{j,\varepsilon'}^*)$ for all $0 < \varepsilon' < \varepsilon$.³ This means that $\bigcap_{0 < \varepsilon' < \varepsilon} B_{\varepsilon'}(f_{j,\varepsilon'}^*) = \{f_j\}$.

We claim that for all $i = 1, \dots, k$ it then holds that

$$|D|_{q_i(f_j)} = n_{ij}.$$

By choice, this already holds for $i = i(j)$. Suppose i is an index with $n_{ij} = 0$. As $|D|_{q_i(\mathbf{F}) \cap q_i(B_\varepsilon(f_{j,\varepsilon}^*))} = n_{ij} = 0$ and $f_j \in B_\varepsilon(f_{j,\varepsilon}^*)$, it follows that $|D|_{q_i(f_j)} = 0$. Now let i be an index different from $i(j)$ such that $n_{ij} > 0$. We have argued before that then there exists a single fact $f_{i,j,\varepsilon} \in q_i(\mathbf{F}) \cap q_i(B_\varepsilon(f_{j,\varepsilon}^*)) = n_{ij}$. Assume that $f_{i,j,\varepsilon} \neq q_i(f_j)$. Then $q_i^{-1}(f_{i,j,\varepsilon}) \neq f_j$ and therefore, there is $\varepsilon' < \varepsilon$ such that $q_i^{-1}(f_{i,j,\varepsilon'}) \notin B_\varepsilon(f_{j,\varepsilon'}^*)$, leading to a contradiction with (b). Thus, $f_{i,j,\varepsilon} = q_i(f_j)$ and $|D|_{q_i(f_j)} = |D|_{f_{i,j,\varepsilon}} = n_{ij}$. From property (1), it follows that $|Q(D)|_{f_j} = n_j$. Thus, since $n_1 + \dots + n_\ell \geq n$ and since the facts $f_j \in \mathbf{F}$ are pairwise distinct, it follows that $|Q(D)|_{\mathbf{F}} \geq n$.

To conclude the proof, we argue that Condition 4.8 can be used to express $|Q(D)|_{\mathbf{F}} \geq n$ as a countable combination of counting events. Note that because \mathbb{Q}_+ is dense in \mathbb{R}_+ , the equivalence of $|Q(D)|_{\mathbf{F}} \geq n$ and Condition 4.8 still holds, if the numbers ε_0 and ε are additionally required to be rational. Also, the fact sets in (b) are measurable in \mathfrak{F}_i : the open balls are certainly measurable, and as q_i is injective and continuous, q_i maps measurable sets to measurable sets by Fact 2.6. Then the set of database instances $D \in \mathbb{DB}$ with Condition 4.8 is of the shape

$$\bigcup_{\ell, (n_i), (n_{ij})} \bigcup_{\varepsilon_0} \bigcap_{\varepsilon} \bigcup_{(f_{j,\varepsilon}^*)} \bigcap_{i,j} \{D \in \mathbb{DB} : |D|_{q_i(\mathbf{F}) \cap q_i(B_\varepsilon(f_{j,\varepsilon}^*))} = n_{ij}\},$$

with the indices ranging as in Condition 4.8 and $\varepsilon, \varepsilon_0$ additionally restricted to \mathbb{Q} . □

³Generally, it could happen that j switches roles with some j' where $(n_{1j'}, \dots, n_{kj'}) = (n_{1j}, \dots, n_{kj})$. If there are exactly m indices j for which the sequences (n_{1j}, \dots, n_{kj}) coincide, the corresponding balls $B_\varepsilon(f_{j,\varepsilon}^*)$ for these indices have (at least) m distinct accumulation points.




FIGURE 2. Illustration of Condition 4.8 in the case $k = 2$. Note that $f_{j,\epsilon}^*$ itself need not necessarily lie in \mathbf{F} . The dashed arrows indicate the intended convergence of $f_{j,\epsilon}^*$ to the fact f_j as $\epsilon \rightarrow 0$.

Incidentally, Theorem 4.4 is a special case of Lemma 4.6 for $k = 1$ and with $N_Q(n) := \{n\}$ for all $n \in \mathbb{N}_+$.

4.4. Coarse Preimages. For $D \in \mathbb{DB}$, and every relation symbol $R \in \tau$, we let

$$\mathbf{T}_R(D) := \{t \in \mathbf{T}_R : |D|_{R(t)} > 0\}$$

denote the set of R -tuples in D . Moreover, we let

$$d_R(D) := \begin{cases} \min\{d_R(t, t') : t, t' \in \mathbf{T}_R(D) \text{ with } t \neq t'\} & \text{if } |\mathbf{T}_R(D)| \geq 2 \text{ and} \\ \infty & \text{otherwise.} \end{cases}$$

That is, $d_R(D)$ is the smallest distance between any two R -tuples of D (or ∞ , if D contains at most one R -fact).

Definition 4.9. Let $\epsilon > 0$. An instance $D \in \mathbb{DB}$ is called ϵ -coarse if for all $R \in \tau$ it holds that $d_R(D) > \epsilon$. We denote the set of ϵ -coarse instances by $\mathbb{DB}|_\epsilon$.

Unfolding the definition, all instances in $\mathbb{DB}|_\epsilon$ have the property that their tuples (per relation) are sufficiently far apart with respect to the metric on the respective space of tuples.

Example 4.10. Consider our example of temperature recordings, but for simplicity (in order for not having to discuss the metric on the product space), assume that there is only one relation, with a single attribute for (real-valued) temperature recordings. Then a database instance is ϵ -coarse precisely if the temperatures occurring in its instances differ by more than ϵ between distinct facts.

The following lemma states that the set of ϵ -coarse instances is measurable for any $\epsilon > 0$.

Lemma 4.11. *It holds that $\mathbb{DB}|_\epsilon \in \mathfrak{DB}$ for all $\epsilon > 0$.*

Proof. An instance $D \in \mathbb{DB}$ is *not* ϵ -coarse, if for some $R \in \tau$, there exist $t_1, t_2 \in \mathbf{T}_R(D)$ such that $t_1 \neq t_2$ and $d_R(t_1, t_2) \leq \epsilon$. We claim that this is the case if and only if D satisfies the following condition.

Condition 4.12. *There are $k_1, k_2 \in \mathbb{N}_+$ and $\varepsilon_0 > 0$ such that for all $r \in (0, \varepsilon_0)$ there are $t_{1,r}^*, t_{2,r}^* \in \mathbb{T}_R^*$ with $\varepsilon_0 < d_R(t_{1,r}^*, t_{2,r}^*) < \varepsilon + 2r$ such that*

- (1) $|D|_{R(B_r(t_{1,r}^*))} = k_1$ and
- (2) $|D|_{R(B_r(t_{2,r}^*))} = k_2$.

The condition basically states that we can find approximations $t_{1,r}^*$ and $t_{2,r}^*$ of two tuples t_1 and t_2 in $\mathbf{T}_R(D)$ that witness that $d_R(t_1, t_2)$ is too small. In particular, the role of ε_0 in Condition 4.12 is to guarantee that $t_{1,r}^*$ and $t_{2,r}^*$ do not end up approximating the same tuple. We proceed to show the following equivalence:

there are $t_1, t_2 \in \mathbf{T}_R(D), t_1 \neq t_2$ with $d_R(t_1, t_2) \leq \varepsilon \iff D$ satisfies Condition 4.12.

To simplify notation, we let $d := d_R$ denote our metric on \mathbb{T}_R , and $d(D) := d_R(D)$.

\Rightarrow : Let t_1, t_2 be two distinct tuples in $\mathbf{T}_R(D)$ of minimal distance, in particular satisfying $0 < d(t_1, t_2) \leq \varepsilon$. Let $k_1 := |D|_{R(t_1)}$ and $k_2 := |D|_{R(t_2)}$. We fix some $\varepsilon_0 > 0$ that satisfies $\varepsilon_0 < \frac{1}{3}d(t_1, t_2)$. Since \mathbb{T}_R^* is dense in \mathbb{T}_R , for every $r > 0$ (in particular for $r < \varepsilon_0$) there exist $t_{1,r}^*, t_{2,r}^* \in \mathbb{T}_R^*$ with $d(t_{1,r}^*, t_1), d(t_{2,r}^*, t_2) < r$. Since $d(t_1, t_2) > 3\varepsilon_0$ and $r < \varepsilon_0, r > 0$, using the triangle inequality, it follows that

$$d(t_{1,r}^*, t_2) \geq d(t_1, t_2) - d(t_1, t_{1,r}^*) > 3\varepsilon_0 - r > \varepsilon_0 > r$$

and similarly that $d(t_{2,r}^*, t_1) > \varepsilon_0 > r$. Because t_1 and t_2 were chosen with minimal distance, this entails $|D|_{R(B_r(t_{1,r}^*))} = |D|_{R(t_1)} = k_1$ and $|D|_{R(B_r(t_{2,r}^*))} = |D|_{R(t_2)} = k_2$. Note that, again using the triangle inequality, it holds that

$$d(t_{1,r}^*, t_{2,r}^*) \leq d(t_{1,r}^*, t_1) + d(t_1, t_2) + d(t_2, t_{2,r}^*) < \varepsilon + 2r$$

and, moreover,

$$d(t_{1,r}^*, t_{2,r}^*) \geq d(t_1, t_2) - d(t_1, t_{1,r}^*) - d(t_2, t_{2,r}^*) > \varepsilon_0.$$

Together, D satisfies Condition 4.12.

\Leftarrow : Now suppose Condition 4.12 holds for some $k_1, k_2 \in \mathbb{N}_+$, and some $\varepsilon_0 > 0$. By Condition 4.12, for all positive $r < \varepsilon_0$ there exist $t_{1,r}^*, t_{2,r}^* \in \mathbb{T}_R^*$ with $\varepsilon_0 < d(t_{1,r}^*, t_{2,r}^*) < \varepsilon + 2r$ such that $|D|_{R(B_r(t_{1,r}^*))} = k_1 > 0$ and $|D|_{R(B_r(t_{2,r}^*))} = k_2 > 0$. Note that this implies $\varepsilon_0 < \varepsilon$.

Now for all $r < \frac{1}{3} \min\{d_R(D), \varepsilon_0\}$, the balls $B_r(t_{1,r}^*)$ and $B_r(t_{2,r}^*)$ are disjoint, and contain at exactly one tuple from $\mathbf{T}_R(D)$ each, say $\mathbf{T}_R(D) \cap B_r(t_{1,r}^*) = \{t_1\}$ and $\mathbf{T}_R(D) \cap B_r(t_{2,r}^*) = \{t_2\}$. Thus, $|D|_{R(t_1)} = |D|_{R(B_r(t_{1,r}^*))} = k_1 > 0$ and $|D|_{R(t_2)} = |D|_{R(B_r(t_{2,r}^*))} = k_2 > 0$. Moreover,

$$d(t_1, t_2) \leq d(t_1, t_{1,r}^*) + d(t_{1,r}^*, t_{2,r}^*) + d(t_{2,r}^*, t_2) < \varepsilon + 4r. \quad (4.2)$$

In particular, Condition 4.12 implies that for all $r > 0$ there exist distinct $t_1, t_2 \in \mathbf{T}_R(D)$ with $d(t_1, t_2) < \varepsilon + 4r$. As D is finite, letting $r \rightarrow 0$ implies that there exist distinct $t_1, t_2 \in \mathbf{T}_R(D)$ with $d(t_1, t_2) \leq \varepsilon$.

Because \mathbb{Q} is dense in \mathbb{R} , the equivalence still holds, if the numbers ε_0 and r are restricted to \mathbb{Q} . Then,

$$(\mathbb{DB}|_\varepsilon)^c = \bigcup_R \bigcup_{\varepsilon_0} \bigcap_r \bigcup_{t_{1,r}^*, t_{2,r}^*} \{D \in \mathbb{DB} : |D|_{R(B_r(t_{1,r}^*))} = k_1 \text{ and } |D|_{R(B_r(t_{2,r}^*))} = k_2\} \in \mathfrak{DB}$$

(with the indices ranging as in our equivalence, numbers being restricted to rationals). Thus, $\mathbb{DB}|_\varepsilon \in \mathfrak{DB}$. \square

Note that for every database instance D , there exists an $\varepsilon > 0$, small enough, such that $D \in \mathbb{DB}|_\varepsilon$, because D is finite. This means that $\bigcup_{\varepsilon>0} \mathbb{DB}|_\varepsilon$ covers \mathbb{DB} (even if the union is taken only over rational ε).

Corollary 4.13. *If for all $\varepsilon > 0$, all $\mathbf{F} \in \mathfrak{F}_Q$, and all $n \in \mathbb{N}$ it holds that*

$$\{D \in \mathbb{DB}|_\varepsilon : |Q(D)|_{\mathbf{F}} = n\} \in \mathfrak{DB},$$

then Q is measurable.

Proof. This follows directly from

$$\{D \in \mathbb{DB} : |Q(D)|_{\mathbf{F}} = n\} = \bigcup_{\varepsilon \in \mathbb{Q}_+} \{D \in \mathbb{DB}|_\varepsilon : |Q(D)|_{\mathbf{F}} = n\}. \quad \square$$

Corollary 4.13 can be used to leverage the finiteness of our instances. In an ε -coarse instance, we can approximate sets of facts by simpler sets of facts as long as these approximations are sufficiently fine. For example, in the context of Example 4.10, in order to prove measurability of a query, it suffices to prove the measurability with respect to preimages where the temperature recordings are “far apart”.

5. RELATIONAL ALGEBRA

As motivated in Section 3.2, we investigate the measurability of relational algebra queries in our model. The concrete relational algebra for bags that we use here is basically the (unnested version of the) algebra that was introduced in [DGK82] and investigated, respectively extended, and surveyed in [Alb91, GM96, GLMW96]. It is called BALG¹ (with superscript 1) in [GM96].

We do not introduce nesting as it would yield yet another layer of abstraction and complexity to the spaces we investigate, although by the properties that such spaces exhibit, we have strong reason to believe that there is no technical obstruction in allowing spaces of finite bags as attribute domains. It is unclear however, whether this extends to PDBs with unbounded nesting depth.

The operations we consider are shown in the Table 2 below. As seen in [Alb91, GM96, GLMW96], there is some redundancy within this set of operations that will be addressed later. A particular motivation for choosing this particular algebra is that possible worlds semantics are usually built on top of set semantics and these operations naturally extend the common behavior of relation algebra queries to bags. This is quite similar to the original motivation of [DGK82] and [Alb91] regarding their choice of operations.

TABLE 2. The operators of BALG^1 considered in this paper.

<i>Base Queries</i>	Constructors	$Q = \{\{\}\}$ and $Q = \{R(t_0)\}$
	Extractors	$Q = R$
	Renaming	$Q = \rho_{A \rightarrow B}(R)$
<i>Basic Bag Operations</i>	Additive Union	$Q = R \uplus S$
	Max Union	$Q = R \cup S$
	Intersection	$Q = R \cap S$
	Difference	$Q = R - S$
	Deduplication	$Q = \delta(R)$
<i>SPJ-Operations</i>	Selection	$Q = \sigma_{(A_1, \dots, A_k) \in B}(R)$
	Projection	$Q = \pi_{A_1, \dots, A_k}(R)$
	Cross Product	$Q = R \times S$

Since compositions of measurable functions are measurable, it suffices to show the measurability of the operators from Table 2, and the measurability of compound queries follows by structural induction.

Therefore, by investigating the measurability of the operators from Table 2 we will show the following main result of this section.

Theorem 5.1. *All queries expressible in the bag algebra BALG^1 are measurable.*

5.1. **Base Queries.** The base queries (Table 3) are easily seen to be measurable.

TABLE 3. Base Queries.

Query	Semantics (for all $D \in \mathbb{DB}$, $t \in \mathbb{T}_Q$)
$Q = \{\{\}\}$	$ Q(D) _{R_Q(t)} := 0$
$Q = \{R_Q(t_0)\}$	$ Q(D) _{R_Q(t)} := 1$ if $t = t_0$, and 0 otherwise
$Q = R$	$ Q(D) _{R_Q(t)} := D _{R(t)}$

Lemma 5.2. *The following queries are measurable:*

- (1) $Q = \{\{\}\}$,
- (2) $Q = \{f\}$ for all $f \in \mathbb{F}_Q$, and
- (3) $Q = R$ for all $R \in \tau$.

Proof. (1) Let $D \in \mathfrak{DB}_Q$. If $\{\{\}\} \in D$, then $Q^{-1}(D) = \mathbb{DB} \in \mathfrak{DB}$. Otherwise, if $\{\{\}\} \notin D$, then $Q^{-1}(D) = \emptyset \in \mathfrak{DB}$.

(2) This can be shown the same way.

(3) Let $Q = R$, let $R_Q(\mathbf{T}) \in \mathfrak{F}_Q$ and let $n \in \mathbb{N}$. Then $R(\mathbf{T}) \in \mathfrak{F}_R \subseteq \mathfrak{F}$. For all $D \in \mathbb{DB}$, it holds that

$$|Q(D)|_{R_Q(\mathbf{T})} = n \iff |D|_{R(\mathbf{T})} = n,$$

and the latter is measurable in \mathfrak{DB} . Thus, Q is measurable. \square

Note that there is nothing to show for the renaming query because it leaves all tuples themselves untouched. In the subsequent subsections, we deal with the remainder of Table 2.

5.2. Basic Bag Operations. We now treat the basic bag operations (Table 4). Assume that $R, S \in \tau$ are relation symbols of the same sort. From now on, we will only consider the case where R and S are distinct, as for the case $R = S$, measurability is trivial.

TABLE 4. Basic Bag Operations.

Query	Semantics (for all $D \in \mathbb{DB}$, $t \in \mathbb{T}_Q$)
$Q = R \uplus S$	$ Q(D) _{R_Q(t)} := D _{R(t)} + D _{S(t)}$
$Q = R - S$	$ Q(D) _{R_Q(t)} := \max\{0, D _{R(t)} - D _{S(t)}\}$
$Q = R \cup S$	$ Q(D) _{R_Q(t)} := \max\{ D _{R(t)}, D _{S(t)}\}$
$Q = R \cap S$	$ Q(D) _{R_Q(t)} := \min\{ D _{R(t)}, D _{S(t)}\}$

Lemma 5.3. *The following queries are measurable:*

- (1) $Q = R \uplus S$.
- (2) $Q = R - S$.
- (3) $Q = R \cup S$.
- (4) $Q = R \cap S$.

Proof. As \cup and \cap can be expressed by compositions of \uplus and $-$ [Alb91], it suffices to show (1) and (2).

- (1) Consider the functions $g: \mathbb{F}_Q \rightarrow \mathbb{F}_R$ and $h: \mathbb{F}_Q \rightarrow \mathbb{F}_S$ with

$$g(R_Q(t)) = R(t) \quad \text{and} \quad h(R_Q(t)) = S(t)$$

for all $t \in \mathbb{T}_Q = \mathbb{T}_R = \mathbb{T}_S$. Clearly, both functions are injective and continuous. For all $k \in \mathbb{N}$ we define

$$N_Q(n) := \{(n_1, n_2) \in \mathbb{N}^2 : n_1 + n_2 = n\}.$$

Then $|Q(D)|_f = n$ if and only if $(|D|_{g(f)}, |D|_{h(f)}) \in N_Q(n)$. Thus, Q is measurable by Lemma 4.6.

- (2) This works analogously to part (1), with

$$N_Q(n) := \{(n_1, n_2) \in \mathbb{N}^2 : \max\{n_1 - n_2, 0\} = n\}. \quad \square$$

5.3. Set Semantics. The deduplication operator (Table 5) maps bag instances to their underlying set instances.

TABLE 5. Deduplication.

Query	Semantics (for all $D \in \mathbb{DB}$, $t \in \mathbb{T}_Q$)
$Q = \delta(R)$	$ Q(D) _{R_Q(t)} := 1$ if $ D _{R(t)} > 0$, and 0 otherwise

Lemma 5.4. *The deduplication query $Q = \delta(R)$ is measurable for all $R \in \tau$.*

Proof. We apply Lemma 4.6 for $r = 1$ and the function $q_1 = q: \mathbb{F}_Q \rightarrow \mathbb{F}_R$ defined by $q(R_Q(x)) = R(x)$. Then for all $D \in \mathbb{DB}$ and all $f \in \mathbb{F}_Q$, it holds that

$$|Q(D)|_f = |D|_{q(f)}.$$

Thus, we let $N_Q(k) = \{0\}$ if $k = 0$ and $N_Q(k) = \mathbb{N} \setminus \{0\}$ otherwise. Then r , q and N_Q satisfy the requirements of Lemma 4.6, so Q is measurable. \square

Having the deduplication query measurable means that standard PDBs support set semantics.

Remark 5.5. The function associated with the deduplication query is countable-to-one (preimage of a single instance in the result is a countable collection of instances) and measurable by the lemma above. This can be used to infer that the space of set instances is standard Borel using [Sri98, Theorem 4.12.4]. This means that we could also completely restrict our setting to set instances without introducing new measurability problems. In general however, it can still be mathematically more convenient to use the full measurable space that was defined in Section 3.1, even in a set semantics setting. The point processes defining PDBs should then just be “simple” in the sense of Definition 2.10, i. e. have the probability 0 of containing duplicate facts.

5.4. Selection and Projection. In this section, we investigate the selection and projection operators Table 6. First, we note that reordering the attributes in the sort of a relation yields a measurable query.

Lemma 5.6. *Let $\text{sort}(R) = (A_1, \dots, A_r)$ and let β be a permutation of $\{1, \dots, r\}$. For $f = R(a_1, \dots, a_r)$ let $\beta(f) = R_Q(a_{\beta(1)}, \dots, a_{\beta(r)})$, and let $Q(D) := \{\beta(f) : f \in D\}$. Then Q is measurable.*

Proof. This directly follows from the mapping theorem (Theorem 4.4), because $\beta: \mathbb{F}_R \rightarrow \mathbb{F}_Q$ is measurable. \square

Lemma 5.6 is helpful for restructuring relations into a more convenient shape to work with later. Semantically, it is a special case of a projection query.

Let $R \in \tau$ be a relation symbol with $\text{ar}(R) = r$, and let A_1, \dots, A_k be pairwise distinct attributes appearing in $\text{sort}(R)$ where $0 < k \leq r$. By Lemma 5.6, wlog. we assume that $\text{sort}(R) = (A_1, \dots, A_r)$.

TABLE 6. Selection and Projection.

Query	Semantics (for all $D \in \mathbb{DB}$, $t \in \mathbb{T}_Q$)
$Q = \sigma_{(A_1, \dots, A_k) \in \mathbf{B}}(R)$	$ Q(D) _{R_Q(t)} := D _{R(t)}$ if $t[A_1, \dots, A_k] \in \mathbf{B}$, and 0 otherwise
$Q = \pi_{A_1, \dots, A_k}(R)$	$ Q(D) _{R_Q(t)} := \sum_{t': t'[A_1, \dots, A_k] = t} D _{R(t')}$

Lemma 5.7. *Let $\mathbf{B} \in \bigotimes_{i=1}^k \mathfrak{A}_i$. Then the query $Q = \sigma_{(A_1, \dots, A_k) \in \mathbf{B}}(R)$ is measurable.*

Proof. Fix $\mathbf{F} \in \mathfrak{F}_Q$, say $\mathbf{F} = R_Q(\mathbf{T})$ with $\mathbf{T} \in \bigotimes_{i=1}^k \mathfrak{A}_i$, and let $n \in \mathbb{N}$. Define

$$\mathbf{F}^{-1} := R(\mathbf{T}) \cap R(\mathbf{B} \times \mathbb{A}_{k+1} \times \dots \times \mathbb{A}_r) \quad (5.1)$$

Then $\mathbf{F}^{-1} \in \mathfrak{F}_R \subseteq \mathfrak{F}$. It holds that

$$|Q(D)|_{\mathbf{F}} = n \iff |D|_{\mathbf{F}^{-1}} = n.$$

Thus, Q is measurable. □

Example 5.8. Suppose $A_1, A_2 \in \text{sort}(R)$ with $\mathbb{A}_1 = \mathbb{A}_2 = \mathbb{R}$. The sets

$$\mathbf{T}_= := \{(x, y) \in \mathbb{R}^2 : x = y\} \quad \text{and} \quad \mathbf{T}_< := \{(x, y) \in \mathbb{R}^2 : x < y\}$$

are Borel in \mathbb{R}^2 . Thus, the queries

$$\sigma_{A_1=A_2}(R) := \sigma_{(A_1, A_2) \in \mathbf{T}_=}(R) \quad \text{and} \quad \sigma_{A_1 < A_2}(R) := \sigma_{(A_1, A_2) \in \mathbf{T}_<}(R)$$

are measurable by Lemma 5.7. In particular, in the context of our running example, we can do selections based on comparing temperatures.

Lemma 5.9. *The query $Q = \pi_{A_1, \dots, A_k}(R)$ is measurable.*

Proof. Fix $\mathbf{F} \in \mathfrak{F}_Q$, say $\mathbf{F} = R_Q(\mathbf{T})$ with $\mathbf{T} \in \bigotimes_{i=1}^k \mathfrak{A}_i$, and let $n \in \mathbb{N}$. Again, define

$$\mathbf{F}^{-1} := R(\mathbf{T} \times \mathbb{A}_{k+1} \times \dots \times \mathbb{A}_r). \quad (5.2)$$

Then $\mathbf{F}^{-1} \in \mathfrak{F}_R$ and it holds that

$$|Q(D)|_{\mathbf{F}} = n \iff |D|_{\mathbf{F}^{-1}} = n,$$

and, hence, Q is measurable. □

Remark 5.10. Above, we provided direct proofs of Lemma 5.7 and Lemma 5.9. Alternatively, they follow from Theorem 4.4 using the function

$$q: \{R(t) \in \mathbb{F}_R : t[A_1, \dots, A_k] \in \mathbf{B}\} \rightarrow \mathbb{F}_Q \quad \text{with} \quad q(R(t)) = R_Q(t)$$

for selection and the function

$$q: \mathbb{F}_R \rightarrow \mathbb{F}_Q \quad \text{with} \quad q(R(t)) = R_Q(t[A_1, \dots, A_k])$$

for projection. A closer look reveals that the sets \mathbf{F}^{-1} of (5.1) and (5.2) are really the preimages of the respective function q that we know from Theorem 4.4.

5.5. Products. Let $R, S \in \tau$ be relation symbols. In this section, we treat the cross product $Q = R \times S$ (Table 7). (For the discussions here, it does not matter whether $R = S$.)

TABLE 7. Cross Products.

Query	Semantics (for all $D \in \mathbb{DB}$, $t \in \mathbb{T}_R$, $t' \in \mathbb{T}_S$)
$Q = R \times S$	$ Q(D) _{R_Q(t, t')} := D _{R(t)} \cdot D _{S(t')}$

Lemma 5.11. *Let $\mathbf{F} \in \mathfrak{F}_Q$ such that $\mathbf{F} = R_Q(\mathbf{T}_R \times \mathbf{T}_S)$ with $\mathbf{T}_R \in \mathfrak{T}_R$ and $\mathbf{T}_S \in \mathfrak{T}_S$. Then*

$$\{D \in \mathbb{DB} : |Q(D)|_{\mathbf{F}} = n\} \in \mathfrak{DB}$$

for all $n \in \mathbb{N}$.

Proof. Let $\mathbf{F} \in \mathfrak{F}_Q$, say $\mathbf{F} = R_Q(\mathbf{T})$ with $\mathbf{T} \in \mathbb{T}_Q$. If $\mathbf{T} = \mathbf{T}_R \times \mathbf{T}_S$ for some $\mathbf{T}_R \in \mathbb{T}_R$ and $\mathbf{T}_S \in \mathbb{T}_S$, then for all $D \in \mathbb{D}$, it holds that

$$|Q(D)|_{\mathbf{F}} = |D|_{\mathbf{F}_R} \cdot |D|_{\mathbf{F}_S}, \quad (5.3)$$

where $\mathbf{F}_R = R(\mathbf{T}_R)$ and $\mathbf{F}_S = S(\mathbf{T}_S)$. Thus,

$$\{D \in \mathbb{D}\mathbb{B} : |Q(D)|_{\mathbf{F}} = n\} = \bigcup_{n_R \cdot n_S = n} \{D \in \mathbb{D}\mathbb{B} : |D|_{\mathbf{F}_R} = n_R \text{ and } |D|_{\mathbf{F}_S} = n_S\} \in \mathfrak{B}. \quad \square$$

Remark 5.12. Note that we proceeded similar to Lemma 4.6. Unfortunately, not every set $\mathbf{T} \in \mathbb{T}_Q$ can be decomposed into a product of measurable sets \mathbf{T}_R and \mathbf{T}_S as used above. The sets \mathbf{T}_R and \mathbf{T}_S are the projections of \mathbb{T}_Q to \mathbb{T}_R and \mathbb{T}_S . Recall that we already encountered a similar situation in Example 3.12. In general, such projections of measurable sets from products of standard Borel spaces to their factors yield the *analytic* sets (cf. [Kec95, Exercise 14.3]). A fundamental theorem in descriptive set theory states that every uncountable standard Borel space has analytic, non-(Borel-)measurable subsets [Kec95, Theorem 14.2]. Thus, with suitably chosen attribute domains, there are sets $\mathbf{T} \in \mathfrak{T}_Q$ such that, for instance, the projection of \mathbf{T} to \mathbb{T}_R is not in \mathfrak{T}_R . Therefore, the approach from (5.3) doesn't help us resolve the measurability of Q , even though it appeared promisingly similar to the arguments from the previous section. In particular, for the given query we cannot establish criterion (2) from Lemma 4.6.

For $t \in \mathbb{T}_Q$, we let $t_R \in \mathbb{T}_R$ and $t_S \in \mathbb{T}_S$ denote the projections of t to its R - and S -part. Then for all $r > 0$, $\mathbf{W}_r(t) := B_r(t_R) \times B_r(t_S) \subseteq \mathbb{T}_Q$ is a *measurable rectangle* containing t with $B_r(t_R) \subseteq \mathbb{T}_R$ and $B_r(t_S) \subseteq \mathbb{T}_S$. Note that $B_r(t_R)$ is a ball with respect to the (Polish) metric on \mathbb{T}_R , and $B_r(t_S)$ is a ball with respect to the (Polish) metric on \mathbb{T}_S .

Remark 5.13. Let us briefly comment on the intuition of the setup. The sets $\mathbf{W}_r(t)$ should be thought of as a small *windows* that we can use to make our considerations local around $t = (t_R, t_S)$. We use these windows to show query measurability by exploiting the finiteness of our database instances: since every D is finite, D is ε -coarse for $\varepsilon > 0$ small enough (see Section 4.4). Then for small enough radius r , the balls $B_r(t_R)$ and $B_r(t_S)$ both contain at most one R - or S -tuple from D , respectively. Thus, the image of D under the cross product query also contains at most one tuple in $\mathbf{W}_r(t) = B_r(t_R) \times B_r(t_S)$. Then, as $\mathbf{W}_r(t)$ has the appropriate shape, Lemma 5.11 can be used again.

Recall that for all $\varepsilon > 0$, $\mathbb{D}\mathbb{B}|_{\varepsilon}$ is the set of ε -coarse instances in $\mathbb{D}\mathbb{B}$. Let $\mathbf{F} \in \mathfrak{F}_Q$ and $n \in \mathbb{N}$. For all $\varepsilon > 0$, we let

$$\mathbf{D}_{\mathbf{F}, n, \varepsilon} := \{D \in \mathbb{D}\mathbb{B}|_{\varepsilon} : |Q(D)|_{\mathbf{F}} = n\}.$$

Then $\mathbf{D}_{\mathbf{F}, n, \varepsilon}$ is the ε -coarse preimage of the event $\{D_Q \in \mathbb{D}\mathbb{B}_Q : |D|_{\mathbf{F}} = n\}$ from $\mathbb{D}\mathbb{B}_Q$.

Lemma 5.14. *For all $\mathbf{F} \in \mathfrak{F}_Q$, $n \in \mathbb{N}$, $t \in \mathbb{T}_Q$ and all $\varepsilon, r > 0$, with $r < \frac{\varepsilon}{3}$ it holds that $\mathbf{D}_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t)), n, \varepsilon} \in \mathfrak{B}$.*

Proof. We prove the lemma as follows. Let

$$\mathfrak{F}_{\varepsilon} := \{\mathbf{F} \in \mathfrak{F}_Q : \mathbf{D}_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t)), n, \varepsilon} \in \mathfrak{B} \text{ for all } t \in \mathbb{T}_Q, n \in \mathbb{N}, \text{ and } r < \frac{\varepsilon}{3}\}. \quad (5.4)$$

Intuitively, this is the set of measurable sets \mathbf{F} with the property that all sufficiently fine “window approximations” of the ε -coarse preimages of $\{D_Q \in \mathbb{D}\mathbb{B}_Q : |D_Q|_{\mathbf{F}} = n\}$ are measurable in \mathfrak{B} .

It then suffices to show for all $\varepsilon > 0$ that

- (1) for every set $\mathbf{F} = R_Q(\mathbf{T})$ with $\mathbf{T} = \mathbf{T}_R \times \mathbf{T}_S \in \mathfrak{T}_R \times \mathfrak{T}_S \subseteq \mathfrak{T}_Q$ it holds that $\mathbf{F} \in \mathfrak{F}_\varepsilon$, and
 (2) the family \mathfrak{F}_ε is a σ -algebra on \mathbb{F}_Q .

From these two it follows that $\mathfrak{F}_\varepsilon = \mathfrak{F}_Q$. That is, indeed every measurable set $\mathbf{F} \in \mathfrak{F}_Q$ satisfies the property from (5.4). This line of argument is occasionally referred to as the *good sets principle* [Ash72, p. 5].

We fix $\varepsilon > 0$ arbitrary and show (1) and (2).

- (1) Suppose $\mathbf{F} = R_Q(\mathbf{T})$ for some $\mathbf{T} \in \mathfrak{T}_R \times \mathfrak{T}_S$, say $\mathbf{T} = \mathbf{T}_R \times \mathbf{T}_S$. Let $t \in \mathbb{T}_Q$, $n \in \mathbb{N}$, and $r < \frac{\varepsilon}{3}$ be arbitrary. With $t_R \in \mathbb{T}_R$ and $t_S \in \mathbb{T}_S$, we denote the R - and the S -part of t , respectively. Then

$$\begin{aligned} \mathbf{F} \cap R_Q(\mathbf{W}_r(t)) &= R_Q(\mathbf{T}_R \times \mathbf{T}_S) \cap R_Q(B_r(t_R) \times B_r(t_S)) \\ &= R_Q((\mathbf{T}_R \cap B_r(t_R)) \times (\mathbf{T}_S \cap B_r(t_S))). \end{aligned}$$

By Lemma 5.11, it follows that $\{D \in \mathbb{DB} : |Q(D)|_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t))} = n\} \in \mathfrak{DB}$. With Lemma 4.11 this entails that

$$\mathbf{D}_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t)), n, \varepsilon} = \mathbb{DB}|_\varepsilon \cap \{D \in \mathbb{DB} : |Q(D)|_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t))} = n\} \in \mathfrak{DB}.$$

As t , n and r were arbitrary, it follows that $\mathbf{F} \in \mathfrak{F}_\varepsilon$.

- (2) We show that \mathfrak{F}_ε is a σ -algebra on \mathbb{F}_Q by showing $\mathbb{F}_Q \in \mathfrak{F}_\varepsilon$, and that it is closed under complements and countable intersections.

First note that $\mathbb{F}_Q \in \mathfrak{F}_\varepsilon$ follows from Lemma 5.11 because $\mathbb{F}_Q = R_Q(\mathbb{T}_R \times \mathbb{T}_S)$. Now let $t \in \mathbb{T}_Q$, $n \in \mathbb{N}$, and $r < \frac{\varepsilon}{3}$ be arbitrary but fixed, and let $D \in \mathbb{DB}|_\varepsilon$. Figure 3 provides visualizations for the remaining cases.




FIGURE 3. Illustrations for the proof of Lemma 5.14. The indications of $B_r(t_R)$ and $B_r(t_S)$ are only to convey the intuition behind the construction of $\mathbf{W}_r(t)$.

- (a) Let $\mathbf{F} \in \mathfrak{F}_\varepsilon$. Then it holds that

$$\begin{aligned} D \in \mathbf{D}_{\mathbf{F}^c \cap R_Q(\mathbf{W}_r(t)), n, \varepsilon} &\iff |Q(D)|_{\mathbf{F}^c \cap R_Q(\mathbf{W}_r(t))} = n \\ &\iff |Q(D)|_{R_Q(\mathbf{W}_r(t))} - |Q(D)|_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t))} = n \end{aligned}$$

Recall that D is ε -coarse. In particular, as $r < \frac{\varepsilon}{3}$, the ball $B_r(t_R)$ contains at most one R -tuple from D and the ball $B_r(t_S)$ contains at most one S -tuple from D .

Hence, $\mathbf{W}_r(t) = B_r(t_R) \times B_r(t_S)$ contains at most one R_Q -tuple from $Q(D)$. Thus, $D \in \mathbf{D}_{\mathbf{F}^c \cap R_Q(\mathbf{W}_r(t)), n, \varepsilon}$ if and only if

$$|Q(D)|_{R_Q(\mathbf{W}_r(t))} = n \text{ and } |Q(D)|_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t))} = 0.$$

From (1) and the assumption $\mathbf{F} \in \mathfrak{F}_\varepsilon$, it thus follows that $\mathbf{F}^c \in \mathfrak{F}_\varepsilon$.

(b) Now let $(\mathbf{F}_i)_{i=1,2,\dots}$ be a sequence of sets $\mathbf{F}_i \in \mathfrak{F}_\varepsilon$ and let $\mathbf{F} = \bigcap_{i=1}^\infty \mathbf{F}_i$. Then it holds that

$$\begin{aligned} D \in \mathbf{D}_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t)), n, \varepsilon} & \\ \iff |Q(D)|_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t))} = n & \\ \iff |Q(D)|_{\bigcap_{i=1}^\infty (\mathbf{F}_i \cap R_Q(\mathbf{W}_r(t)))} = n. & \end{aligned}$$

As in the previous case, because D is ε -coarse, the set $\mathbf{W}_r(t)$ contains at most one R_Q -tuple from $Q(D)$. Thus, $D \in \mathbf{D}_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t)), n, \varepsilon}$ if and only if

$$|Q(D)|_{\mathbf{F}_i \cap R_Q(\mathbf{W}_r(t))} = n \text{ for all } i = 1, 2, \dots$$

Thus, $\bigcap_{i=1}^\infty \mathbf{F}_i \in \mathfrak{F}_\varepsilon$ using the assumption.

Together, \mathfrak{F}_ε is indeed a σ -algebra on \mathbb{F}_Q .

From (1) and (2) it follows that $\mathfrak{F}_\varepsilon = \mathfrak{F}_Q$ for all $\varepsilon > 0$. □

Lemma 5.15. *For all $\mathbf{F} \in \mathfrak{F}_Q$, all $n \in \mathbb{N}$, and all $\varepsilon > 0$, it holds that $\mathbf{D}_{\mathbf{F}, n, \varepsilon} \in \mathfrak{DB}$.*

Proof. Let $\mathbf{F} \in \mathfrak{F}_Q$, $n \in \mathbb{N}$, and $\varepsilon > 0$. It suffices to show that

$$\mathbf{D}_{\mathbf{F}, \geq n, \varepsilon} := \bigcup_{m \geq n} \mathbf{D}_{\mathbf{F}, m, \varepsilon} \in \mathfrak{DB}$$

where we may assume $n > 0$. We show that $D \in \mathbf{D}_{\mathbf{F}, \geq n, \varepsilon}$ is equivalent to D satisfying the following condition.

Condition 5.16. *The instance D is ε -coarse and for some $\ell \in \mathbb{N}_+$ there are $k_1, \dots, k_\ell \in \mathbb{N}_+$ with $k_1 + \dots + k_\ell \geq n$ such that for all $r \in (0, \frac{\varepsilon}{3})$ there are $t_{1,r}^*, \dots, t_{\ell,r}^* \in \mathbb{T}_R^* \times \mathbb{T}_S^*$ such that*

- (1) *for all $i \neq j$ it holds that $d_R(t_{i,r}^*, t_{j,r}^*) > \frac{\varepsilon}{3}$ or $d_S(t_{i,r}^*, t_{j,r}^*) > \frac{\varepsilon}{3}$ and*
- (2) *for all $i = 1, \dots, \ell$ it holds that $D \in \mathbf{D}_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t_{i,r}^*)), k_i, \varepsilon}$.*

We start with the easy direction (\Leftarrow).

\Leftarrow : Suppose D satisfies Condition 5.16. Then $D \in \mathbb{DB}|_\varepsilon$ and it remains to show $|Q(D)|_{\mathbf{F}} \geq n$. Note that it suffices to show that if r is small enough in Condition 5.16, then the sets $\mathbf{W}_r(t_{i,r}^*)$ are pairwise disjoint. In this case, the claim follows from (2).

By (1), for all $i, j = 1, \dots, n$ with $i \neq j$ it holds that at least one of $d_R(t_{i,r}^*, t_{j,r}^*)$ or $d_R(t_{i,S}^*, t_{j,S}^*)$ is larger than $\frac{\varepsilon}{3}$. Thus, for $r < \frac{\varepsilon}{6}$ it follows that at least one of $d_R(t_{i,R}^*, t_{j,R}^*)$ or $d_S(t_{i,S}^*, t_{j,S}^*)$ is larger than $2r$. Therefore, $\mathbf{W}_r(t_{i,r}^*) \cap \mathbf{W}_r(t_{j,r}^*) = \emptyset$.

\Rightarrow : Suppose that $D \in \mathbf{D}_{\mathbf{F}, \geq n, \varepsilon}$. Then D is ε -coarse and $|Q(D)|_{\mathbf{F}} \geq n$. Thus, there exist pairwise distinct $t_1, \dots, t_\ell \in Q(D)$ with $R_Q(t_i) \in \mathbf{F}$ for all $i = 1, \dots, n$ such that

$$|Q(D)|_{R_Q(t_1)} + \dots + |Q(D)|_{R_Q(t_\ell)} \geq n.$$

Let $r \in (0, \frac{\varepsilon}{3})$. Since \mathbb{T}_R^* is dense in \mathbb{T}_R and \mathbb{T}_S^* is dense in \mathbb{T}_S , for all $i = 1, \dots, n$ we can choose $t_{i,r}^* = (t_{i,r,R}^*, t_{i,r,S}^*) \in \mathbb{T}_R^* \times \mathbb{T}_S^*$ such that

$$d_R(t_{i,r,R}^*, t_{i,R}) < r \quad \text{and} \quad d_S(t_{i,r,S}^*, t_{i,S}) < r.$$

Thus, $t_i \in \mathbf{W}_r(t_{i,r}^*)$. Now let $t \in Q(D)$ with $t \neq t_i$. Because D is ε -coarse, it holds that

$$d_R(t_{i,R}, t_R) > \varepsilon \quad \text{or} \quad d_S(t_{i,S}, t_S) > \varepsilon.$$

For the rest of the proof we assume $d_R(t_{i,R}, t_R) > \varepsilon$ (the other case is completely symmetric). Recall that $r < \frac{\varepsilon}{3}$. Then it follows that

$$d_R(t_{i,r,R}^*, t_R) \geq d_R(t_{i,R}, t_R) - d_R(t_{i,R}, t_{i,r,R}^*) > \varepsilon - r > r.$$

Thus, $t \notin \mathbf{W}_r(t_{i,r}^*)$, so every $\mathbf{W}_r(t_{i,r}^*)$ contains no tuple from $Q(D)$ other than t_i . Thus, $|Q(D)|_{R_Q(t_i)} = |Q(D)|_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t_{i,r}^*))}$, and (2) follows.

Also for all $j \neq i$, we have that

$$d_R(t_{i,r,R}^*, t_{j,r,R}^*) \geq d_R(t_{i,R}, t_{j,R}) - d_R(t_{i,R}, t_{i,r,R}^*) - d_R(t_{j,R}, t_{j,r,R}^*) > \varepsilon - 2r > \frac{\varepsilon}{3},$$

establishing (1).

The equivalence still holds, if r is additionally required to be rational. That is,

$$D_{\mathbf{F}, \geq n, \varepsilon} = \bigcup_{\ell} \bigcup_{k_1, \dots, k_\ell} \bigcap_r \bigcup_{t_{1,r}^*, \dots, t_{\ell,r}^*} D_{\mathbf{F} \cap R_Q(\mathbf{W}_r(t_{i,r}^*)), k_i, \varepsilon} \in \mathfrak{DB}$$

using Lemma 5.14 (with the indices ranging as in our equivalence, and numbers being restricted to rationals). \square

Finally, the measurability of $Q = R \times S$ is a direct consequence of Lemma 5.15 and Corollary 4.13.

Lemma 5.17. *The query $Q = R \times S$ is measurable.* \square

Altogether, Lemmas 5.2, 5.7, 5.9 and 5.17 and Theorems 5.3 and 5.4 now prove the main result—the measurability of BALG^1 -queries—of this section (Theorem 5.1).

As a consequence, we also obtain the measurability of all kinds of derived operators, including the typical join operators. Note that it also follows from Theorem 5.1, that we can use finite Boolean combinations of predicates for selection queries.

6. AGGREGATION

There are practically relevant queries that are not already covered by our treatment in the previous section. For example, in our running example of temperature recordings, we might be interested in returning the *average* (or *minimum* or *maximum*) temperature per room, taken over all temperature records for this particular room.

In this section, we formalize aggregate operators and aggregate queries, possibly with grouping in the standard PDB framework in a possible worlds semantics style (cf. Section 3.2). In particular, we show that these queries are measurable in the standard PDB framework.

Remark 6.1. Often, when a separate treatment of aggregate queries over purely algebraic ones is motivated, it is mentioned that the correspondence of relational algebra and relational calculus limits expressive power to that of first-order logic. However, bag query languages based on relational algebra *do* allow expressing various kinds of aggregation based on exploiting the presence of multiplicities. For example, counting in a unary fashion is possible in BALG^1 [GM96]. Here, we follow a more general approach in allowing basically any measurable function over finite bags to be used for aggregation. This goes beyond the integer aggregation of BALG [GLMW96].

Let $(\mathbb{A}, \mathfrak{A})$ and $(\mathbb{B}, \mathfrak{B})$ be standard Borel spaces. An *aggregate operator* (or *aggregator*) from \mathbb{A} to \mathbb{B} is a function $\Phi: \mathcal{B}_{\text{fin}}(\mathbb{A}) \rightarrow \mathbb{B}$, that is, a function mapping (finite) bags of elements of \mathbb{A} to elements of \mathbb{B} . If $\mathbb{T}_R = \mathbb{A}$, then every such aggregator Φ induces an *aggregation query* $Q = \varpi_\Phi(R)$ with $\mathbb{T}_Q = \mathbb{B}$ as follows:

$$|Q(D)|_{R_Q(t)} := \begin{cases} 1 & \text{if } t = \Phi(R(D)), \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

for all $D \in \mathbb{DB}$.⁴ In essence, its output on a single instance D is the instance containing only the single “tuple” corresponding to the value of the aggregation over the relation R in D . Examples of common attribute operators Φ are shown in Table 8.

Proposition 6.2. *If Φ is a $(\mathbf{Count}(\mathbb{T}_R), \mathfrak{T}_Q)$ -measurable aggregator, then $\varpi_\Phi(R)$ is a measurable query.*

Proof. Let $\mathbf{F} = R_Q(\mathbf{T})$ where $\mathbf{T} \in \mathfrak{T}_Q$. Without loss of generality, we assume that R is the only relation in τ . Then for all $D \in \mathbb{DB}$, it holds that

$$|Q(D)|_{\mathbf{F}} = 1 \iff R(D) \in \Phi^{-1}(\mathbf{T}) \iff D \in R(\Phi^{-1}(\mathbf{T}))$$

(and $|Q(D)|_{\mathbf{F}} = 0$ otherwise). Since Φ is $(\mathbf{Count}(\mathbb{T}_R), \mathfrak{T}_Q)$ -measurable, the claim follows. \square

TABLE 8. Common aggregate operators. (We assume that **MIN**, **MAX** and **AVG** have a suitable definition on empty bags, so that their semantics are well-defined.)

Name	Definition
Count	$\mathbf{CNT}(\{a_1, \dots, a_m\}) := m$
Distinct Count	$\mathbf{CNTd}(\{a_1, \dots, a_m\}) := \{a_1, \dots, a_m\} $
Sum	$\mathbf{SUM}(\{a_1, \dots, a_m\}) := a_1 + \dots + a_m$
Minimum	$\mathbf{MIN}(\{a_1, \dots, a_m\}) := \min\{a_1, \dots, a_m\}$
Maximum	$\mathbf{MAX}(\{a_1, \dots, a_m\}) := \max\{a_1, \dots, a_m\}$
Average	$\mathbf{AVG}(\{a_1, \dots, a_m\}) := \frac{a_1 + \dots + a_m}{m}$

For $m \in \mathbb{N}$, we call a function $\varphi: \mathbb{T}_R^m \rightarrow \mathbb{T}_Q$ *symmetric* if $\varphi(t) = \varphi(t')$ for all $t \in \mathbb{T}_R^m$ and all permutations t' of t .

Lemma 6.3. *For all $m \in \mathbb{N}$, let $\varphi_m: \mathbb{T}_R^m \rightarrow \mathbb{T}_Q$ be a symmetric, measurable function. Then the aggregator $\Phi: \mathcal{B}_{\text{fin}}(\mathbb{T}_R) \rightarrow \mathbb{T}_Q$ with*

$$\Phi(\{a_1, \dots, a_m\}) := \varphi_m(a_1, \dots, a_m)$$

is $(\mathbf{Count}(\mathbb{T}_R), \mathfrak{T}_Q)$ -measurable.

Proof. It suffices to show that the restriction Φ_m of Φ to $\mathcal{B}_m(\mathbb{T}_R)$, the bags of cardinality m over \mathbb{T}_R , is measurable for all $m \in \mathbb{N}$. Note that $\mathcal{B}_0(\mathbb{T}_R)$ only contains the empty bag, and \mathbb{T}_R^0 only contains the empty tuple. That is, the statement is trivial for $m = 0$. Thus, let $m \in \mathbb{N}_+$. Since φ_m is $(\mathfrak{T}_R^{\otimes m}, \mathfrak{T}_Q)$ -measurable and symmetric, for all $\mathbf{T} \in \mathfrak{T}_Q$ it holds that

⁴Formally, $R(D)$ has been defined as the bag of R -facts in D whereas Φ should take bags of R -tuples. This small type mismatch is of no significance whatsoever.

$\varphi_m^{-1}(\mathbf{T})$ is a symmetric set in $\mathfrak{T}_R^{\otimes m} \subseteq \bigoplus_{m=0}^{\infty} \mathfrak{T}_R^{\otimes m}$. From Remark 2.8 and Fact 2.9 it follows that $\Phi_m^{-1}(\mathbf{T}) \in \mathbf{Count}(\mathbb{T}_R)$. \square

Example 6.4. All the aggregators in Table 8 yield measurable aggregation queries. The associated functions φ_m are all continuous under suitable choices of attribute domains. That is, for example, if $+$ in the definition of **SUM** is the addition of real numbers.

What we have introduced so far is only sufficient to express the aggregation over all tuples of a relation at once. Usually, we want to perform aggregation separately for parts of the data, as in our motivating example of returning the average temperature per room. For this, we need to *group* tuples before aggregating values. Suppose we want to group a relation R by attributes A_1, \dots, A_k and perform the aggregation over attribute A , separately for every occurring value of the attributes A_1, \dots, A_k . Without loss of generality, we assume that $\text{sort}(R) = (A_1, \dots, A_k, A)$. Then what we described is an *group-by aggregate query* $\varpi_{A_1, \dots, A_k, \Phi(A)}(R)$ and is defined by

$$|Q(D)|_{R_Q(a_1, \dots, a_k, b)} := \begin{cases} 1 & \text{if } b = \Phi(\mathbf{A}_{a_1, \dots, a_k}(D)) \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (6.1)$$

for all $D \in \mathbb{DB}$ where $\mathbf{A}(D)$ is the bag of values a such that $R(a_1, \dots, a_k, a) \in D$ with

$$|\mathbf{A}_{a_1, \dots, a_k}(D)|_a = |D|_{R(a_1, \dots, a_k, a)}.$$

Essentially, $\mathbf{A}_{a_1, \dots, a_k}(D)$ is obtained by selecting those tuples where the first k attributes have values a_1, \dots, a_k , and then projecting to the last attribute. Both kinds of aggregate queries (without, and with grouping) are shown in Table 9.

TABLE 9. Aggregate Queries.

Query	Semantics (for all $D \in \mathbb{DB}$, $t, (a_1, \dots, a_k, b) \in \mathbb{T}_Q$)
$Q = \varpi_{\Phi}(R)$	$ Q(D) _{R_Q(t)} := 1$ if $t = \Phi(R(D))$, and 0 otherwise
$Q = \varpi_{(A_1, \dots, A_k, \Phi(A))}(R)$	$ Q(D) _{R_Q(a_1, \dots, a_k, b)} := 1$ if $b = \Phi(\mathbf{A}_{a_1, \dots, a_k}(D))$ and 0 otherwise

Theorem 6.5. *If $\Phi: \mathcal{B}_{\text{fin}}(\mathbb{A}) \rightarrow \mathbb{B}$ is a $(\mathbf{Count}(\mathbb{A}), \mathfrak{B})$ -measurable aggregator, then the query $\varpi_{A_1, \dots, A_k, \Phi(A)}(R)$ is a measurable query.*

Proof. Let $\mathbb{T}_{\text{grp}} = \prod_{i=1}^k \mathbb{A}_i$. In the following, we fix a compatible Polish metric d_{grp} on \mathbb{T}_{grp} , and a countable dense set $\mathbb{T}_{\text{grp}}^*$ in \mathbb{T}_{grp} . For all $t \in \mathbb{T}_{\text{grp}}$ and all $r > 0$ define

$$Q_{t,r} = \pi_{A_1, \dots, A_k}(\sigma_{(A_1, \dots, A_k) \in B_r(t)}(R)) \times \varpi_{\Phi}(\pi_A(\sigma_{(A_1, \dots, A_k) \in B_r(t)}(R))).$$

Note that $Q_{t,r}$ is measurable by Theorem 5.1 for all particular choices of t and r . Let $\tilde{Q} = \pi_{A_1, \dots, A_k}(R)$ and

$$\widetilde{\mathbb{DB}}|_{\varepsilon} := \{D \in \mathbb{DB} : \tilde{Q}(D) \text{ is } \varepsilon\text{-coarse}\}$$

Then $\widetilde{\mathbb{DB}}|_{\varepsilon} \in \mathfrak{B}$, as $\widetilde{\mathbb{DB}}|_{\varepsilon} = \tilde{Q}^{-1}(\widetilde{\mathbb{DB}}'|_{\varepsilon})$ using Lemmas 4.11 and 5.9 where \mathbb{DB}' is the output instance space of \tilde{Q} . Intuitively, $\widetilde{\mathbb{DB}}|_{\varepsilon}$ are the instances that are ε -coarse in the (A_1, \dots, A_k) attributes of R . Similar to Corollary 4.13 it suffices to show that $\{D \in \widetilde{\mathbb{DB}}|_{\varepsilon} : |Q(D)|_{\mathbf{F}} \geq n\}$ is measurable for all positive ε , $\mathbf{F} \in \mathfrak{F}_Q$ and $n \in \mathbb{N}_+$.

We show that for all $D \in \widetilde{\mathbb{DB}}|_{\varepsilon}$, all $\mathbf{F} \in \mathfrak{F}_Q$, and all $n \in \mathbb{N}_+$ it holds that $|Q(D)|_{\mathbf{F}} \geq n$ is equivalent to D satisfying the following condition.

Condition 6.6. For all positive $r < \frac{\varepsilon}{3}$ there exist $t_{1,r}^*, \dots, t_{n,r}^* \in \mathbb{T}_{\text{grp}}^*$ with $d_{\text{grp}}(t_{i,r}^*, t_{j,r}^*) > \frac{\varepsilon}{3}$ such that $|Q_{t_{i,r}^*, r}^*(D)|_{\mathbf{F}} \geq 1$ for all $i = 1, \dots, n$.

\Rightarrow : Let $D \in \widetilde{\mathbb{DB}}|_{\varepsilon}$ such that $|Q(D)|_{\mathbf{F}} \geq n$. Then there exist $R_Q(t_1, b_1), \dots, R_Q(t_n, b_n) \in \mathbf{F}$ with $|Q(D)|_{R_Q(t_i, b_i)} = 1$ for all $i = 1, \dots, n$ such that $d_{\text{grp}}(t_i, t_j) > \varepsilon$ for all $i \neq j$. Since $\mathbb{T}_{\text{grp}}^*$ is dense in \mathbb{T}_{grp} , for all positive r (in particular $r < \frac{\varepsilon}{3}$) there exist $t_{i,r}^*$ with $d_{\text{grp}}(t_i, t_{i,r}^*) < r$ for all $i = 1, \dots, n$. Since $d_{\text{grp}}(t_i, t_j) > \varepsilon$, it follows that $d_{\text{grp}}(t_{i,r}^*, t_{j,r}^*) > \frac{\varepsilon}{3}$. In particular, every $B_r(t_{i,r}^*)$ contains no tuple among t_1, \dots, t_n other than t_i . Thus,

$$|Q_{t_{i,r}^*, r}^*(D)|_{R_Q(t_i, b_i)} = 1,$$

i. e. $|Q_{t_{i,r}^*, r}^*(D)|_{\mathbf{F}} \geq 1$ for all $i = 1, \dots, n$.

\Leftarrow : Suppose Condition 6.6 holds. Since D is finite, the tuples $t_{i,r}^*$ in Condition 6.6 converge to tuples t_i with

$$|Q_{t_{i,r}^*, r}^*(D)|_{R_Q(t_i, b_i)} \geq 1 \tag{6.2}$$

for some $b_i \in \mathbb{A}$ where $R_Q(t_i, b_i) \in \mathbf{F}$. Because $D \in \widetilde{\mathbb{DB}}|_{\varepsilon}$, and $d_{\text{grp}}(t_{i,r}^*, t_{j,r}^*) > \frac{\varepsilon}{3}$, the tuples t_1, \dots, t_n are pairwise distinct. Thus, (6.2) implies $|Q(D)|_{\mathbf{F}} \geq \sum_{i=1}^n |Q_{t_{i,r}^*, r}^*(D)|_{\mathbf{F}} \geq n$.

The equivalence still holds, when r is additionally required to be rational. Thus,

$$\{D \in \widetilde{\mathbb{DB}}|_{\varepsilon} : |Q(D)|_{\mathbf{F}} \geq n\} = \bigcap_r \bigcup_{t_{1,r}^*, \dots, t_{n,r}^*} \bigcap_i \{D \in \widetilde{\mathbb{DB}}|_{\varepsilon} : |Q_{t_{i,r}^*, r}^*(D)|_{\mathbf{F}} \geq 1\}.$$

with the indices ranging as in Condition 6.6 (and r rational). \square

By the observation of Example 6.4, Theorem 6.5 applies to the operators from Table 8.

Corollary 6.7. The query $\varpi_{A_1, \dots, A_k, \Phi(A)}(R)$ is measurable for all the aggregate operators $\Phi \in \{\text{CNT}, \text{CNTd}, \text{SUM}, \text{MIN}, \text{MAX}, \text{AVG}\}$. \square

7. DATALOG

The measurability results of the previous sections also allow us to say something about fixpoint queries. The key observation is the following lemma, which follows from Fact 2.5(2).

Lemma 7.1. Let $(Q_i)_{i \in \mathbb{N}}$ be a family of measurable queries such that $Q(D) := \bigcup_{i=0}^{\infty} Q_i(D)$ is finite for all $D \in \mathbb{D}$. Then Q is a measurable query.

Proof. For all $n \in \mathbb{N}$, let $Q^{(n)} := \bigcup_{i=0}^n Q_i$. As a finite (maximum-)union of measurable queries, $Q^{(n)}$ is measurable. As $Q(D)$ is finite for all D , it holds that $Q(D) = \lim_{n \rightarrow \infty} Q^{(n)}(D)$, and Q is the pointwise limit of the functions $Q^{(n)}$. Thus, $Q = \lim_{n \rightarrow \infty} Q^{(n)}$ is measurable as well. \square

We omit the definition of Datalog and related query languages. For simplicity, we consider set PDBs, and Datalog with sets semantics. Recall that if Q is a Datalog query, then Q can be written as a countable union of conjunctive queries [AHV95].⁵ Thus, combining our measurability results with the above lemma, we obtain the following.

⁵We note that a similar statement can be made for the bounded fixpoints semantics over bags featured in [CL97].

Corollary 7.2. *Every Datalog query is measurable.* □

In fact, our argumentation can be applied to all types of queries with operators that are based on countable iterative (or inductive, inflationary, or fixed-point) processes. All we need is that the iterative mechanism forms a converging sequence of measurable queries. For partial Datalog / fixed-point logic, we cannot directly use Lemma 7.2, but a slightly more complicated argument still based on countable limits works there as well.

8. BEYOND POSSIBLE WORLDS SEMANTICS

In Section 3.2, we introduced the notion of queries or views on probabilistic databases solely based on the existing notion of queries and views for traditional databases, which we referred to as the *possible worlds semantics* of queries or views. As the title of this section suggests, we explicitly leave this setup.⁶ Before, we have introduced views as functions mapping database instances to database instances and adopted a semantics based on possible worlds. Now, we want to discuss *PDB views* as functions that map probabilistic databases to probabilistic databases, for which no such semantics (to be precise, a definition as in the shape of Section 3.2) exists. Such “views” naturally arise in a variety of computational problems in probabilistic databases. For example, consider the following problems or “queries”:

- *probabilistic threshold queries* that intuitively return a deterministic table containing only those facts which have a marginal probability over some specified threshold [QJSP10];
- *probabilistic top-k queries* that intuitively return a deterministic table containing the k most probable facts [RDS07];
- *probabilistic skyline queries* [PJLY07] that consider how different instances compare to each other with respect to some notion of *dominance*; and
- *conditioning* [KO08] the probabilistic database to some event.

Note that the way we informally explained the first two queries above is only sensible if the space of facts is discrete. In a continuous setting, we interpret these queries with respect to a suitable countable partition of the fact space into measurable sets.

More such “queries” as the above can be found in [Agg09, WLLW13]. These queries (or views) still take as input a PDB and produce some output, but differ from the ones we have seen so far in that they have no reasonable semantics on *single* instances (i. e. per possible world). This can be, for example, because they explicitly take probabilities into account. The goal of this section is to interpret other kinds of problems on PDBs abstractly as functions on probability spaces in order fit them into a unified framework. Developing such an understanding has already been motivated in [WvK15] as yielding potential insight into common properties of the corresponding problems.

We now present our formal classification of views that are directly defined on the probability space level of a PDB (as opposed to the instance level as in Section 3.2). Let \mathbf{PDB}_τ denote the class of probabilistic databases of schema τ . Note that all PDBs in \mathbf{PDB}_τ have the same instance measurable space $(\mathbb{DB}, \mathfrak{DB})$. Queries and, more generally, views of input schema τ and output schema τ' are now mappings $V : \mathbf{PDB}_\tau \rightarrow \mathbf{PDB}_{\tau'}$.

⁶Note that we *do not* abandon our definition of PDBs as probability spaces over possible worlds. It is only that we broaden the notion of views to also incorporate mappings of PDBs that are not defined on a “per possible world” basis.

Definition 8.1 (View Types).

- (1) Every view is *type 1*.
- (2) A view V is *type 2*, or *pointwise local*, if (and only if) for every fixed input PDB $\mathcal{D} = (\mathbb{DB}, \mathfrak{DB}, P)$ there exists a measurable function $q_{\mathcal{D}}: \mathbb{DB} \rightarrow \mathbb{DB}_V$ such that $P_V = P \circ q_{\mathcal{D}}^{-1}$.
- (3) A view V is *type 3*, or *uniformly local*, if there exists a measurable function $q: \mathbb{DB} \rightarrow \mathbb{DB}_V$ such that for every input PDB $\mathcal{D} = (\mathbb{DB}, \mathfrak{DB}, P)$, it holds that $P_V = P \circ q^{-1}$.
- (4) A view V is *type 4*, or *pointwise*, if (and only if) there exists measurable function $q: \mathbf{F}_q \rightarrow \mathbb{F}_Q$ such that V is composed of functions of the shape of Q as in Theorem 4.4 (where Q depends on q).

We let \mathbf{V}_I , \mathbf{V}_{II} , \mathbf{V}_{III} and \mathbf{V}_{IV} denote the classes of views of type 1 throughout 4.

Remark 8.2. Let us shed some more light on these classes and their names:

- (1) Class \mathbf{V}_I does not require further explanation, as it contains *every* view.
- (2) The class \mathbf{V}_{II} is described via functions that may depend on the measurable structure of the input PDB. Specifically, these views are functions that directly transform input to output PDBs. We dub this “pointwise local”, because this function is applied per instance (hence, local) but the function itself depends on the concrete PDB (hence, is only pointwise local with respect to probability spaces). In general, type 2 views may take the probability space level into account. For example, probabilistic threshold or probabilistic top- k queries can be viewed as views of this class: they transform any input PDB to a single database instance (that is, a PDB with only one possible world of probability 1) containing the respective output tuples along with their probability in a separate attribute. Another example is conditioning a PDB, as the probability measure of a conditioned PDB involves a normalization term that depends on the probability mass of an event.
- (3) The class \mathbf{V}_{III} captures the lifting of typical database queries to PDBs under the possible worlds semantics. Hence, there is a *single* (measurable) function that is applied “locally” on every database instance. The term “uniformly” expresses that it does not depend on the concrete PDB. In general, type 3 views only take the instance level into account. All the views that we investigated in Sections 5 to 7 fall into this category.
- (4) The class \mathbf{V}_{IV} is the class of views corresponding to the measurability criterion of the Mapping Theorem. That is, there is a (measurable) function that is applied “pointwise”⁷ on every fact. This transformation naturally lifts to instances and, thus, to PDBs. In general, type 4 views only take the fact level into account. We have seen an example of such a view in Example 4.5.

Example 8.3. Recall our running example of temperature measurements. We use this to introduce an example of a view that performs “out-of-world aggregation” [WvK15] (and that, in particular, is of type 2, but not of type 3). The relation `TEMPREC` stores triples of room numbers (`RoomNo`), recording dates (`Date`) and recorded temperatures (`Temp`). Assume that the pair (`RoomNo`, `Date`) acts as a key, so that with probability 1 there is at most one temperature recording per pair. Moreover, assume that the PDB is modelled with independent tuples (r, d, θ) where θ is Normally distributed per pair (r, d) , but such that the existence of an record belonging to (r, d) is subject to uncertainty. That is, we have a mix of attribute- and tuple-level uncertainty. A possible representation of such a PDB is

⁷Note that the “pointwise” in the term “pointwise local” from \mathbf{V}_{II} refers to probability spaces as “points”, whereas the term “pointwise” alone, as here in the definition of \mathbf{V}_{IV} refers to facts as “points”

shown in Figure 4. Therein, we have two possible tuples (one recording for room 4108 and one recording for room 4108a), but the value of the temperature recording is specified as a Normally distributed random variable, parameterized with its mean and variance. We assume the existence of tuples with room 4108 or 4108a to be independent from one another, and that these events carry probability 0.4 and 0.8, respectively.

TEMPREC			
RoomNo	Date	Temp [°C]	Prob
4108	2021-07-12	$\mathcal{N}(20.5, 0.1)$	0.4
4108a	2021-07-12	$\mathcal{N}(21.0, 0.1)$	0.8

VIEW			
RoomNo	Date	Temp [°C]	Prob
4108	2021-07-12	20.5	0.4
4108a	2021-07-12	21.0	0.8

FIGURE 4. Representation of a PDB along with the representation of a view that goes beyond the instance level.

A possible view could ask to return, per room and date, the *expected* temperature, under preserving the tuple-level uncertainty. The output of this view is shown on the right-hand side of Figure 1. (It is here trivially to obtain, because the expected value was already part of the parametrization. The idea however also applies to more complicated setups.) The possible worlds of the input PDB $\mathcal{D} = (\mathbb{DB}, \mathfrak{DB}, P)$ are partitioned into four cases, depending on the presence of a tuple with room 4108 and 4108a, respectively. Our function $q_{\mathcal{D}}$ (which depends on \mathcal{D}) maps instances $D \in \mathbb{DB}$ as follows: the tuple (4108, 2021-07-12, 20.5) is present in the view result if and only if D contains a tuple with room 4108, and the tuple (4108, 2021-07-12, 21.0) is present in the view result if and only if D contains a tuple with room 4108a. Then the output PDB has probability measure $P \circ q_{\mathcal{D}}^{-1}$ and, in particular, four possible worlds. This is a so-called *tuple-independent* PDB with given marginal probabilities (0.4 and 0.8), hence it can be represented as shown on the right-hand side of Figure 1.⁸

Let us come back to the relationships between the classes of views we defined in Definition 8.1. Clearly, $\mathbf{V}_I \supseteq \mathbf{V}_{II} \supseteq \mathbf{V}_{III} \supseteq \mathbf{V}_{IV}$. We already informally argued in Section 4.2 that queries that depend on multiple tuples per instance are not captured by \mathbf{V}_{III} , i. e. $\mathbf{V}_{III} \not\subseteq \mathbf{V}_{IV}$. We provide two examples to expose that also the remaining inclusions are strict.

Proposition 8.4. *There exists a view that is not type 2, i. e. $\mathbf{V}_{II} \subsetneq \mathbf{V}_I$.*

We demonstrate this using a view that *conditions* its input PDB on an event.

Proof. Let \mathcal{D} be a PDB with three possible worlds D_1 , D_2 , and D_3 , such that

$$P(\{D_1\}) = \frac{1}{6}, \quad P(\{D_2\}) = \frac{1}{2}, \quad \text{and} \quad P(\{D_3\}) = \frac{1}{3}.$$

Consider the view V that conditions a PDB on the event $\{D_1, D_2\}$. Note that

$$P_{V(\mathcal{D})}(\{D_1\}) = \frac{1}{6} / \left(\frac{1}{6} + \frac{1}{2}\right) = \frac{1}{4} \quad \text{and} \quad P_{V(\mathcal{D})}(\{D_2\}) = \frac{1}{2} / \left(\frac{1}{6} + \frac{1}{2}\right) = \frac{3}{4}.$$

Yet, there is no event D in \mathcal{D} with the property that $P(D) = \frac{1}{4}$. Thus, a function $q_{\mathcal{D}}$ as required in Definition 8.1(2) does not exist and, hence, V is not type 2. \square

⁸The input PDB in this case is not tuple-independent, as any two facts with the same room number are mutually exclusive. Instead, this is a so-called (uncountable) *block-independent disjoint* PDB with two independent blocks, both representing a single random tuple, such that each block specifies the (probability distribution over) possible manifestations for this tuple.

Proposition 8.5. *There exists a type 2 view that is not type 3, so $\mathbf{V}_{\text{III}} \subsetneq \mathbf{V}_{\text{II}}$.*

We have already discussed such a view in Example 8.3, but have not actually shown that it is not of type 3. We demonstrate the proposition in considering another example, namely that of a probabilistic threshold query.

Proof. Let $\alpha \in (0, 1]$. We consider a function V_α that maps an input PDB $\mathcal{D} = (\mathbb{DB}, \mathfrak{DB}, P)$ to an output PDB $V_\alpha(\mathcal{D})$ such that

$$D_{\mathcal{D},\alpha} := \left\{ f \in \mathbb{F} : P(\{D \in \mathbb{DB} : |D|_f > 0\}) \geq \alpha \right\}$$

has probability 1 in $V_\alpha(\mathcal{D})$. Note that $D_{\mathcal{D},\alpha}$ is finite for all $\alpha \in (0, 1]$, so V_α is well-defined. The view V_α is a probabilistic threshold query with threshold α . Consider the function $q_{\mathcal{D}}$ with $D \mapsto D_{\mathcal{D},\alpha}$ for all $D \in \mathbb{DB}$. Then $q_{\mathcal{D}}$ is measurable and witnesses that V_α is type 2 for all $\alpha \in (0, 1]$.

Now let $\alpha > \frac{1}{2}$ and consider the following two PDBs \mathcal{D}_1 and \mathcal{D}_2 , with probability measures P_1 and P_2 , respectively, over distinct facts f and g :

D	$\{f\}$	$\{g\}$	D	$\{f\}$	$\{g\}$
$P_1(D)$	α	$1 - \alpha$	$P_2(D)$	$1 - \alpha$	α

FIGURE 5. Definition of \mathcal{D}_1 and \mathcal{D}_2 (given as an explicit list of the instances together with their probabilities).

Then $D_{\mathcal{D}_1,\alpha} = \{f\}$ and $D_{\mathcal{D}_2,\alpha} = \{g\}$. Suppose q exists such that $P_V = P \circ q$ for every input PDB $\mathcal{D} = (\mathbb{DB}, \mathfrak{DB}, P)$. Let $\mathbf{D}_f := \{D \in \mathbb{DB}_V : f \in D\}$. Then $P_{\mathcal{D}_1}(q^{-1}(\mathbf{D}_f)) = 1$ implies $\{f\} \in q^{-1}(\mathbf{D}_f)$. On the contrary, $P_{\mathcal{D}_2}(q^{-1}(\mathbf{D}_f)) = 0$, implies $\{f\} \notin q^{-1}(\mathbf{D}_f)$, a contradiction. Thus, V_α is not of type 3. \square

Together, we have that $\mathbf{V}_{\text{I}} \supsetneq \mathbf{V}_{\text{II}} \supsetneq \mathbf{V}_{\text{III}} \supsetneq \mathbf{V}_{\text{IV}}$. Before closing this section, let us highlight two key insights of the arguments used in the examples for Propositions 8.4 and 8.5. In essence, we separated type 2 from type 3 by arguing about the structure of possible worlds without really taking probabilities into account. For the separation of type 1 from type 2, we argued about the structure of the probability measure instead. This highlights again the different levels within the hierarchical structure of a PDB (fact level—instance level—probability space level) that the views of the different classes operate on.

9. CONCLUSIONS

We introduce the notion of standard PDBs, for which we rigorously describe how to construct suitable measurable spaces for infinite probabilistic databases, completing the picture of [GL19]. The viability of this model as a general and unifying framework for finite and infinite databases is supported by the well-definedness and compositionality of (typical) query semantics. Other kinds of PDB queries embed into the framework as well.

It is currently open, whether, and if so, how more in-depth results on point processes can be used for probabilistic databases, for example, to perform open-world query answering. Also, while we focused on relational algebra and aggregation, the queries of Section 8 deserve a systematic treatment in their own right in infinite PDBs.

ACKNOWLEDGMENTS

We thank Sam Staton, who initially brought the correspondence of probabilistic databases and finite point processes to our attention in personal correspondence.

This work is funded by *Deutsche Forschungsgemeinschaft* (DFG, German Research Foundation) under grants GR 1492/16-1 and GRK 2236 (UnRAVeL).

REFERENCES

- [ACK⁺11] Serge Abiteboul, T.-H. Hubert Chan, Evgeny Kharlamov, Werner Nutt, and Pierre Senellart. Capturing Continuous Data and Answering Aggregate Queries in Probabilistic XML. *ACM Transactions on Database Systems (TODS)*, 36(4):25:1–25:45, 2011. doi:10.1145/1804669.1804679.
- [Agg09] Charu C. Aggarwal, editor. *Managing and Mining Uncertain Data*, volume 35 of *Advances in Database Systems*. Springer, Boston, MA, USA, 2009. doi:10.1007/978-0-387-09690-2.
- [AHV95] Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, Boston, MA, USA, 1st edition, 1995.
- [Alb91] Joseph Albert. Algebraic Properties of Bag Data Types. In *Proceedings of the 17th International Conference on Very Large Databases (VLDB 1991)*, pages 211–219, San Francisco, CA, USA, 1991. Morgan Kaufmann Publishers Inc.
- [Ash72] Robert B. Ash. *Real Analysis and Probability*. Probability and Mathematical Statistics: A Series of Monographs and Textbooks. Academic Press, Inc., New York, NY, USA, 1972. doi:10.1016/C2013-0-06164-6.
- [AW09] Parag Agrawal and Jennifer Widom. Continuous Uncertainty in Trio. In *Proceedings of the 3rd VLDB Workshop on Management of Uncertain Data (MUD '09)*, pages 17–32, Enschede, The Netherlands, 2009. Centre for Telematics and Information Technology (CTIT).
- [Bad07] Adrian Baddeley. Spatial Point Processes and Their Applications. In Wolfgang Weil, editor, *Stochastic Geometry*, Lecture Notes in Mathematics, chapter 1, pages 1–75. Springer, Berlin, Heidelberg, Germany, 1st edition, 2007.
- [BCL17] Stefan Borgwardt, İsmail İlkan Ceylan, and Thomas Lukasiewicz. Ontology-Mediated Queries for Probabilistic Databases. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI '17)*, pages 1063–1069, Palo Alto, CA, USA, 2017. AAAI Press.
- [BCL18] Stefan Borgwardt, İsmail İlkan Ceylan, and Thomas Lukasiewicz. Recent Advances in Querying Probabilistic Knowledge Bases. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI '18)*, pages 5420–5426. International Joint Conferences on Artificial Intelligence, 2018. doi:10.24963/ijcai.2018/765.
- [BCL19] Stefan Borgwardt, İsmail İlkan Ceylan, and Thomas Lukasiewicz. Ontology-Mediated Query Answering over Log-Linear Probabilistic Data. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, volume 33, Palo Alto, CA, USA, 2019. AAAI Press. doi:10.1609/aaai.v33i01.33012711.
- [BDM⁺05] Jihad Boulos, Nilesh Dalvi, Bhushan Mandhani, Shobhit Mathur, Chris Ré, and Dan Suciu. MYSTIQ: A System for Finding more Answers by Using Probabilities. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data (SIGMOD '05)*, pages 891–893, New York, NY, USA, 2005. ACM. doi:10.1145/1066157.1066277.
- [Bel20] Vaishak Belle. Symbolic Logic Meets Machine Learning: A Brief Survey in Infinite Domains. In Jesse Davis and Karim Tabia, editors, *Scalable Uncertainty Management*, Lecture Notes in Computer Science, page 3–16. Springer International Publishing, 2020. doi:10.1007/978-3-030-58449-8_1.
- [BGMP92] Daniel Barbará, Héctor García-Molina, and Daryl Porter. The Management of Probabilistic Data. *IEEE Transactions on Knowledge and Data Engineering*, 4(5):487–502, 1992. doi:10.1109/69.166990.
- [BPVDB15] Vaishak Belle, Andrea Passerini, and Guy Van Den Broeck. Probabilistic Inference in Hybrid Domains by Weighted Model Integration. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI 2015)*, IJCAI'15, pages 2770–2776. AAAI Press, Jul 2015.

- [CDVdB16] İsmail İlkan Ceylan, Adnan Darwiche, and Guy Van den Broeck. Open-World Probabilistic Databases. In *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning (KR '16)*, pages 339–348, Palo Alto, CA, USA, 2016. AAAI Press.
- [CDVdB21] İsmail İlkan Ceylan, Adnan Darwiche, and Guy Van den Broeck. Open-World Probabilistic Databases: Semantics, Algorithms, Complexity. *Artificial Intelligence*, 295, 2021. doi:10.1016/j.artint.2021.103474.
- [CKP03] Reynold Cheng, Dmitri V. Kalashnikov, and Sunil Prabhakar. Evaluating Probabilistic Queries over Imprecise Data. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data (SIGMOD '03)*, pages 551–562, New York, NY, USA, 2003. ACM. doi:10.1145/872757.872823.
- [CL97] Latha S. Colby and Leonid Libkin. Tractable Iteration Mechanisms for Bag Languages. In *Proceedings of the 6th International Conference on Database Theory (ICDT 1997)*, volume 1186 of *Lecture Notes in Computer Science*, pages 461–475, Berlin, Germany, 1997. Springer-Verlag. doi:10.1007/3-540-62222-5_64.
- [CP87] Roger Cavallo and Michael Pittarelli. The Theory of Probabilistic Databases. In *Proceedings of the 13th International Conference on Very Large Data Bases (VLDB '87)*, pages 71–81, San Francisco, CA, USA, 1987. Morgan Kaufmann.
- [CVP⁺13] Zhuhua Cai, Zografoula Vagena, Luis Perez, Subramanian Arumugam, Peter J. Haas, and Christopher Jermaine. Simulation of database-valued markov chains using simsql. In *Proceedings of the 2013 international conference on Management of data - SIGMOD '13*, page 637, New York, New York, USA, 2013. ACM Press. doi:10.1145/2463676.2465283.
- [Deg17] Christoph Degen. *Finite Point Processes and Their Application to Target Tracking*. PhD thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, 2017.
- [DGK82] Umeshwar Dayal, Nathan Goodman, and Randy Howard Katz. An Extended Relational Algebra with Control over Duplicate Elimination. In *Proceedings of the 1st ACM SIGACT-SIGMOD Composium on Principles of Database Systems (PODS '82)*, pages 117–123, New York, NY, USA, 1982. ACM.
- [DGM⁺04] Amol Deshpande, Carlos Guestrin, Samuel R. Madden, Joseph M. Hellerstein, and Wei Hong. Model-Driven Data Acquisition in Sensor Networks. In *Proceedings of the 30th International Conference on Very Large Data Bases (VLDB '04)*, pages 588–599, St. Louis, 2004. Morgan Kaufmann. doi:10.1016/B978-012088469-8.50053-X.
- [DKM10] Daniel Deutch, Christoph Koch, and Tova Milo. On Probabilistic Fixpoint and Markov Chain Query Languages. In *Proceedings of the 29th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS '10)*, pages 215–226, New York, NY, USA, 2010. ACM.
- [DRKNP16] Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole. *Statistical Relational Artificial Intelligence: Logic, Probability, and Computation*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool, San Rafael, CA, USA, 2016.
- [DRS09] Nilesh Dalvi, Christopher Ré, and Dan Suciu. Probabilistic Databases: Diamonds in the Dirt. *Communications of the ACM*, 52(7):86–94, 2009. doi:10.1145/1538788.1538810.
- [DS96] Debabrata Dey and Sumit Sarkar. A Probabilistic Relational Model and Algebra. *ACM Transactions on Database Systems (TODS)*, 21(3), 1996. doi:10.1145/232753.232796.
- [DS07] Nilesh Dalvi and Dan Suciu. Efficient Query Evaluation on Probabilistic Databases. *The VLDB Journal*, 16(4):523–544, 2007. doi:10.1007/s00778-006-0004-3.
- [DS20] Swaraj Dash and Sam Staton. A Monad for Probabilistic Point Processes. *Proceedings 3rd Annual International Applied Category Theory Conference 2020 (ACT 2020)*, 2020. To appear.
- [DVJ03] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Models*. Probability and Its Applications. Springer-Verlag New York, New York, NY, USA, 2nd edition, 2003. doi:10.1007/b97277.
- [DVJ08] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes. Volume II: General Theory and Structure*. Probability and Its Applications. Springer Science+Business Media, LLC, New York, NY, USA, 2nd edition, 2008.
- [FGB02] Anton Faradjian, Johannes Gehrke, and Philippe Bonnet. GADT: A Probability Space ADT for Representing and Querying the Physical World. In *Proceedings of the 18th International*

- Conference on Data Engineering (ICDE '02)*, pages 201–211. IEEE Computing Society, 2002. doi:10.1109/ICDE.2002.994710.
- [FHO12] Robert Fink, Larisa Han, and Dan Olteanu. Aggregation in Probabilistic Databases via Knowledge Compilation. In *Proceedings of the 38th International Conference on Very Large Data Bases (VLDB '12)*, volume 5, pages 490–501. VLDB Endowment, 2012. doi:10.14778/2140436.2140445.
- [FR97] Norbert Fuhr and Thomas Rölleke. A Probabilistic Relational Algebra for the Integration of Information Retrieval and Database Systems. *ACM Transactions on Information Systems (TOIS)*, 15(1):32–66, 1997. doi:10.1145/239041.239045.
- [Fra82] Peter Franken. *Point Process Method in Queueing Theory*, page 163–188. Birkhäuser Boston, Boston, MA, USA, 1982. doi:10.1007/978-1-4612-5791-2_6.
- [Fuh95] Norbert Fuhr. Probabilistic Datalog—A Logic for Powerful Retrieval Methods. In *Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '95)*, pages 282–290, New York, NY, USA, 1995. ACM.
- [GH86] Erol Gelenbe and George Hebrail. A Probability Model of Uncertainty in Data Bases. In *1986 IEEE Second International Conference on Data Engineering*, pages 328–333. IEEE, 1986. doi:10.1109/ICDE.1986.7266237.
- [GHNR14] Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sriram K. Rajamani. Probabilistic Programming. In *Future of Software Engineering Proceedings (FOSE 2014)*, pages 167–181, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2593882.2593900.
- [GKT07] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. Provenance Semirings. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2007)*, pages 31–40. Association for Computing Machinery, 2007. doi:10.1145/1265530.1265535.
- [GL19] Martin Grohe and Peter Lindner. Probabilistic Databases with an Infinite Open-World Assumption. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '19)*, pages 17–31, New York, NY, USA, 2019. ACM. doi:10.1145/3294052.3319681.
- [GL20] Martin Grohe and Peter Lindner. Infinite Probabilistic Databases. In *23rd International Conference on Database Theory (ICDT 2020)*, pages 16:1–16:20, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICDT.2020.16.
- [GLMW96] Stéphane Grumbach, Leonid Libkin, Tova Milo, and Limsoon Wong. Query Languages for Bags: Expressive Power and Complexity. *ACM SIGACT News*, 1996(2):30–44, 1996. doi:10.1145/235767.235770.
- [GM96] Stéphane Grumbach and Tova Milo. Towards Tractable Algebras for Bags. *Journal of Computer and System Sciences*, 52(3):570–588, 1996. doi:10.1006/jcss.1996.0042.
- [Gre09] Todd J. Green. Models for Incomplete and Probabilistic Information. In Charu C. Aggarwal, editor, *Managing and Mining Uncertain Data*, volume 35 of *Advances in Database Systems*, chapter 2, pages 9–43. Springer, Boston, MA, USA, 2009. doi:10.1007/978-0-387-09690-2_2.
- [JXW⁺11] Ravi Jampani, Fei Xu, Mingxi Wu, Luis Perez, Chris Jermaine, and Peter J. Haas. The Monte Carlo Database System: Stochastic Analysis Close to the Data. *ACM Transactions on Database Systems (TODS)*, 36(3):18:1–18:41, 2011. doi:10.1145/2000824.2000828.
- [Kal02] Olav Kallenberg. *Foundations of Modern Probability*. Probability and its Applications. Springer Science+Business Media New York, New York, NY, USA, 2nd edition, 2002. doi:10.1007/978-1-4757-4015-8.
- [Kec95] Alexander S. Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag New York, Inc., softcover reprint of the hardcover 1st edition, 1995. URL: <http://link.springer.com/10.1007/978-1-4612-4190-4>, doi:10.1007/978-1-4612-4190-4.
- [KK10] Oliver Kennedy and Christoph Koch. PIP: A Database System for Great and Small Expectations. In *Proceedings of the 26th International Conference on Data Engineering (ICDE '10)*, pages 157–168, Washington, DC, USA, 2010. IEEE.

- [Kle14] Achim Klenke. *Probability Theory: A Comprehensive Course*. Universitext. Springer-Verlag London, London, UK, Translation from the German language edition; 2nd edition, 2014. doi:10.1007/978-1-4471-5361-0.
- [KLP00] Gabriel Kuper, Leonid Libkin, and Jan Paradaens, editors. *Constraint Databases*. Springer-Verlag, 2000. doi:10.1007/978-3-662-04031-7.
- [KO08] Christoph Koch and Dan Olteanu. Conditioning Probabilistic Databases. In *Proceedings of the 34th International Conference on Very Large Data Bases (VLDB '08)*, volume 1, pages 313–325. VLDB Endowment, 2008. doi:10.14778/1453856.1453894.
- [Koc08] Christoph Koch. On Query Algebras for Probabilistic Databases. *ACM SIGMOD Record*, 37(4):78–85, 2008.
- [Koc09] Christoph Koch. MayBMS: A System for Managing Large Probabilistic Databases. In Charu C. Aggarwal, editor, *Managing and Mining Uncertain Data*, volume 35 of *Advances in Database Systems*, chapter 6, pages 149–184. Springer, Boston, MA, USA, 2009. doi:10.1007/978-0-387-09690-2_6.
- [KT12] Alex Kulesza and Ben Taskar. Determinantal Point Processes for Machine Learning. *Foundations and Trends[®] in Machine Learning*, 5(2–3):123–286, 2012. doi:10.1561/22000000044.
- [LP17] Günter Last and Matthew Penrose. *Lectures on the Poisson Process*. Institute of Mathematical Statistics Textbook. Cambridge University Press, Cambridge, UK, 2017. doi:10.1017/9781316104477.
- [LW97] Leonid Libkin and Limsoon Wong. Query Languages for Bags and Aggregate Functions. *Journal of Computer and System Sciences*, 55(2):241–272, 1997. doi:10.1006/jcss.1997.1523.
- [Mac75] Odile Macchi. The Coincidence Approach to Stochastic Point Processes. *Advances in Applied Probability*, 7(1):83–122, 1975. doi:10.2307/1425855.
- [Mil06] Brian Christopher Milch. *Probabilistic Models with Unknown Objects*. PhD thesis, University of California, Berkeley, 2006.
- [MIW11] Raghotham Murthy, Robert Ikeda, and Jennifer Widom. Making Aggregation Work in Uncertain and Probabilistic Databases. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 23(8):1261–1273, 2011. doi:10.1109/TKDE.2010.166.
- [MMR⁺05] Brian Milch, Bhaskara Marthi, Stuart Russell, David Sontag, David L. Ong, and Andrey Kolobov. BLOG: Probabilistic Models with Unknown Objects. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI '05)*, St. Louis, MO, USA, 2005. Morgan Kaufmann.
- [Moy62] José Enrique Moyal. The General Theory of Stochastic Population Processes. *Acta Mathematica*, 108:1–31, 1962. doi:10.1007/BF02545761.
- [MPS19] Paolo Morettn, Andrea Passerini, and Roberto Sebastiani. Advanced SMT Techniques for Weighted Model Integration. *Artificial Intelligence*, 275:1–27, 2019. doi:10.1016/j.artint.2019.04.003.
- [Pit94] Michael Pittarelli. An Algebra for Probabilistic Databases. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 6(2):293–303, 1994.
- [PJLY07] Jian Pei, Bin Jiang, Xuemin Lin, and Yidong Yuan. Probabilistic Skylines on Uncertain Data. In *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB '07)*, pages 15–26. VLDB Endowment, 2007.
- [QJSP10] Yinian Qi, Rohit Jain, Sarvjeet Singh, and Sunil Prabhakar. Threshold Query Optimization for Uncertain Data. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data (SIGMOD 2010)*, pages 315–326. Association for Computing Machinery, 2010. doi:10.1145/1807167.1807203.
- [RD06] Matthew Richardson and Pedro Domingos. Markov Logic Networks. *Machine Learning*, 62(1–2):107–136, 2006. doi:10.1007/s10994-006-5833-1.
- [RDS07] Christopher Ré, Nilesh Dalvi, and Dan Suciu. Efficient Top-*k* Query Evaluation on Probabilistic Data. In *2007 IEEE 23rd International Conference on Data Engineering (ICDE 2007)*, pages 886–895. IEEE, 2007. doi:10.1109/ICDE.2007.367934.
- [Rei81] Raymond Reiter. *On Closed World Databases*, page 119–140. Morgan Kaufmann, 1981. doi:10.1016/B978-0-934613-03-3.50014-3.
- [RSG05] Robert Ross, V. S. Subrahmanian, and John Grant. Aggregate Operators in Probabilistic Databases. *Journal of the ACM (JACM)*, 52(1):54–101, 2005. doi:10.1145/1044731.1044734.

- [SD07] Parag Singla and Pedro Domingos. Markov Logic in Infinite Domains. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI '07)*, pages 368–375, Arlington, VA, USA, 2007. AUAI Press.
- [SMS⁺08] Sarvjeet Singh, Chris Mayfield, Rahul Shah, Sunil Prabhakar, Susanne Hambrusch, Jennifer Neville, and Reynold Cheng. Database Support for Probabilistic Attributes and Tuples. In *2008 IEEE 24th International Conference on Data Engineering (ICDE '08)*, pages 1053–1061, Washington, DC, USA, 2008. IEEE Computer Society. doi:10.1109/ICDE.2008.4497514.
- [SORK11] Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. *Probabilistic Databases*. Synthesis Lectures on Data Management. Morgan & Claypool, San Rafael, CA, USA, 1st edition, 2011. doi:10.2200/S00362ED1V01Y201105DTM016.
- [Sri98] Shashi Mohan Srivastava. *A Course on Borel Sets*. Graduate Texts in Mathematics. Springer, New York, NY, USA, 1st edition, 1998. doi:10.1007/b98956.
- [TPD⁺12] Thanh T. L. Tran, Liping Peng, Yanlei Diao, Andrew McGregor, and Anna Liu. CLARO: Modeling and Processing Uncertain Data Streams. *The VLDB Journal*, 21(5):651–676, 2012. doi:10.1007/s00778-011-0261-7.
- [VdBS17] Guy Van den Broeck and Dan Suciu. Query Processing on Probabilistic Data: A Survey. *Foundations and Trends[®] in Databases*, 7(3–4):197–341, 2017. doi:10.1561/19000000052.
- [Wid09] Jennifer Widom. Trio: A System for Data, Uncertainty, and Lineage. In Charu C. Aggarwal, editor, *Managing and Mining Uncertain Data*, volume 35 of *Advances in Database Systems*, chapter 5, pages 113–148. Springer, Boston, MA, USA, 2009. doi:10.1007/978-0-387-09690-2_5.
- [Wil04] Stephen Willard. *General Topology*. Dover Publications, Inc., Mineola, NY, USA, Dover reprint edition, 2004.
- [WLLW13] Yijie Wang, Xiaoyong Li, Xiaoling Li, and Yuan Wang. A Survey of Queries over Uncertain Data. *Knowledge and Information Systems*, 37(3):485–530, 2013. doi:10.1007/s10115-013-0638-6.
- [Won82] Eugene Wong. A Statistical Approach to Incomplete Information in Database Systems. *ACM Transactions on Database Systems (TODS)*, 7(3):470–488, 1982. doi:10.1145/319732.319747.
- [WSH⁺18] Yi Wu, Siddharth Srivastava, Nicholas Hay, Simon Du, and Stuart Russell. Discrete-Continuous Mixtures in Probabilistic Programming: Generalized Semantics and Inference Algorithms. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, volume 80 of *Proceedings of Machine Learning Research*, pages 5343–5352. PMLR, 2018.
- [WvK15] Brend Wanders and Maurice van Keulen. Revisiting the Formal Foundation of Probabilistic Databases. In *Proceedings of the 2015 Conference of the International Fuzzy Systems Association and the European Society for Fuzzy Logic and Technology (IFSA-EUSFLAT '15)*, Advances in Intelligent Systems Research, pages 289–296, Paris, France, 2015. Atlantis Press.
- [Zim97] Esteban Zimányi. Query Evaluation in Probabilistic Relational Databases. *Theoretical Computer Science*, 171(1):179–219, 1997. doi:10.1016/S0304-3975(96)00129-6.
- [ZP97] Esteban Zimányi and Alain Pirotte. Imperfect Information in Relational Databases. In Amihai Motro and Philippe Smets, editors, *Uncertainty Management in Information Systems*, pages 35–87. Springer Science+Business Media, LLC, reprint of 1st edition, 1997. doi:10.1007/978-1-4615-6245-0_3.

APPENDIX A. TOPOLOGIES AND METRIC SPACES

The brief account of topological notions is based upon [Wil04].

Let \mathbb{X} be some set. A *topology* on \mathbb{X} is a family \mathfrak{D} of subsets of \mathbb{X} that contains both \emptyset and \mathbb{X} , and is closed under finite intersections and (arbitrary, i. e. possibly uncountable) unions. The pair $(\mathbb{X}, \mathfrak{D})$ is called a *topological space*. The sets in \mathfrak{D} are called *open* sets, and their complements (with respect to \mathbb{X}) are called *closed* sets.

For $\mathbf{X} \subseteq \mathbb{X}$, the *closure* of \mathbf{X} is the intersection of all closed sets that contain \mathbf{X} . A set \mathbf{X} is called *dense* in $(\mathbb{X}, \mathfrak{D})$ if its closure is \mathbb{X} . The topological space $(\mathbb{X}, \mathfrak{D})$ is called *separable* if there exists a countable subset of \mathbb{X} that is dense in $(\mathbb{X}, \mathfrak{D})$.

Let \mathbb{X} be a non-empty set and let $d: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ be a function satisfying

- $d(X, Y) \geq 0$ with $d(X, Y) = 0$ if and only if $X = Y$;
- $d(X, Y) = d(Y, X)$; and
- $d(X, Z) \leq d(X, Y) + d(Y, Z)$ (the *triangle inequality*)

for all $X, Y, Z \in \mathbb{X}$. Then d is called a *metric* on \mathbb{X} and the pair (\mathbb{X}, d) is called a *metric space*. For $\varepsilon > 0$, the set

$$B_\varepsilon(X) := \{Y \in \mathbb{X} : d(X, Y) < \varepsilon\}$$

is the *open ball of radius ε around X* . The *metric topology* on \mathbb{X} with respect to d is the topology \mathfrak{D} on \mathbb{X} that is defined by

$$\mathfrak{D} := \{\mathbf{X} \subseteq \mathbb{X} : \text{for all } X \in \mathbf{X} \text{ there exists some } \varepsilon > 0 \text{ such that } B_\varepsilon(X) \subseteq \mathbf{X}\}.$$

A sequence (X_1, X_2, \dots) of elements in a metric space (\mathbb{X}, d) is called a *Cauchy sequence* if $\lim_{i \rightarrow \infty} d(X_i, X_{i+1}) = 0$. A sequence (X_1, X_2, \dots) in (\mathbb{X}, d) *converges* to an element $X \in \mathbb{X}$ in (\mathbb{X}, d) if for all $\varepsilon \in (0, \infty)$ there exists some i , large enough, such that $d(X, X_i) < \varepsilon$. A metric d on \mathbb{X} is called *complete*, if all Cauchy sequences over (\mathbb{X}, d) converge to elements of \mathbb{X} . A metric space is called *complete*, if its metric is complete. In a metric space, a set \mathbf{X} is *dense* (with respect to the metric topology) if and only if for all $X \in \mathbb{X}$ and all $\varepsilon > 0$ there exists some $Y \in \mathbf{X}$ such that $d(X, Y) < \varepsilon$.

A topological space $(\mathbb{X}, \mathfrak{D})$ is called *metrizable* if there exists a metric d on \mathbb{X} such that \mathfrak{D} is the corresponding metric topology. If such d exists that is additionally complete, the space is called *completely metrizable*. A *Polish space* is a completely metrizable, separable topological space.