# On Probabilistic Models for Uncertain Sequential Pattern Mining 

Muhammad Muzammal and Rajeev Raman<br>Department of Computer Science, University of Leicester, UK.<br>\{mm386,r.raman\}@mcs.le.ac.uk


#### Abstract

We study uncertainty models in sequential pattern mining. We consider situations where there is uncertainty either about a source or an event. We show that both these types of uncertainties could be modelled using probabilistic databases, and give possible-worlds semantics for both. We then describe "interestingness" criteria based on two notions of frequentness (previously studied for frequent itemset mining) namely expected support and probabilistic frequentness. We study the interestingness criteria from a complexity-theoretic perspective, and show that in case of source-level uncertainty, evaluating probabilistic frequentness is \#P-complete, and thus no polynomial time algorithms are likely to exist, but evaluate the interestingness predicate in polynomial time in the remaining cases.


Keywords: Mining Uncertain Data, Sequential Pattern Mining, Probabilistic Databases, Novel Algorithms for Mining, Theoretical Foundations of Data Mining.

## 1 Introduction

The problem of sequential pattern mining (SPM), or finding frequent sequences of events in data with a temporal component, was originally proposed by Agrawal and Srikant for analyzing retail data [16], [3]. Since then, it has gained a great deal of attention [23], [15], [4] and has found applications in other domains, including discovering web access patterns and motif extraction [9], [20].The classical SPM problem assumes that the data is deterministic, or entirely determined. However, it is recognized that data obtained from a wide range of data sources is inherently uncertain [1], [17]. An approach to modelling uncertainty that has seen increased interest recently is that of probabilistic databases [17], [1]; even more recently data mining and ranking problems have been studied in the context of probabilistic databases, including top- $k[24],[7]$ and frequent itemset mining (FIM) [2], [5], [6]. This paper considers the SPM problem in probabilistic databases.

Although SPM is superficially similar to FIM, we observe that generalizing SPM to probabilistic databases gives rise to some subtle issues. As we show, different ways of formalizing the problem can lead to very different outcomes in terms of computational complexity. We now outline the main issues, for which we
give a high-level overview of SPM (a formal definition is in Section 2). In SPM, events (retail transactions, observations of objects/persons) are associated with sources (customers, sensors/cameras). Events have a time-stamp, so the event database is easily re-organized into a collection of source sequences, each of which is a sequence of events (ordered by time stamp) associated with a given source. The SPM problem is to find sequential patterns - patterns of events that have a temporal order - that occur in a significant number of source sequences. Unlike FIM, uncertainty in SPM can occur in three different aspects: the source, the event and the time may all be uncertain. Uncertainty in time seems not wellsuited to the probabilistic database approach, and we focus on uncertainty in (any one of) the source and the event.

Source-Level Uncertainty (SLU). In the following situations, an event is recorded deterministically, but the source is not readily identifiable:
(a) a customer (source) purchases some items (event) from a store, and provides identity information (loyalty programme, filling a form etc.). However, multiple matches may emerge in the customer database: the customer's details may be incomplete or incorrect, or the customer database itself may be probabilistic as a result of "deduplication" or cleaning [11].
(b) a vehicle/person is identified by a sensor/camera using methods such as biometrics or automatic number plate recognition (ANPR), which are inherently noisy. To mine patterns such as " $10 \%$ of cars pass camera $X$, then camera $Y$ and later camera $Z "$, we consider each car as a source and each sighting as an event.

In such scenarios, it is certain that an event occurred (e.g. a customer bought some items, a vehicle/person entered an area) but the source associated with that event is uncertain. The software would typically assign confidence values to the various alternative matches. Whilst large databases of type (a) are readily available in commercial organizations, large databases of type (b) are now becoming more common, a notable example being the UK's police ANPR database [21]. We model the above scenarios by assuming that each event is associated with a probability distribution over possible sources that could have resulted in the event. This formulation thus shows attribute-level uncertainty [17], as the "source" attribute of each event is uncertain.

Event-Level Uncertainty (ELU). In some scenarios, the source of the data is known, but the events are uncertain. Such scenarios could be modelled as a sequence of events output by a source, but with each event only having a certain probability of having truly occurred (i.e. tuple-level uncertainty [17]), or as above, using attribute-level uncertainty. For example:
(a) Employees movements are tracked in a building using RFID sensors [13]. The stream of tags read by various sensors is stored in a relation SIGHTING( $t$, tID, aID), which denotes that the RFID tag tID was detected by antenna aID at time $t$. Since an RFID antenna has only a certain probability of reading a tag within its range, the PEEX system [13] processes the

SIGHTING relation to output an uncertain higher-level event relation such as MEETING(time, person1, person2, room, prob). An example tuple in MEETING could be (103, 'Alice', 'Bob', 435, 0.4), which means that at time 103, PEEX believes that Alice and Bob are having a meeting (event) with probability 0.4 in room 435 (source) [13]; since antennae are at fixed locations, the source is certain but the event is uncertain. ${ }^{1}$
(b) A logged-in user (source) enters terms into a search engine (events), which are disambiguated in many different ways e.g. a search term "Tiger" could be disambiguated as $\langle($ Animal, 0.5$)$, (Sports Personality, 0.3 ), (Airplane, 0.1 ), ... $\rangle$.

Our Contributions. The problem of SPM in probabilistic databases was introduced in [14]. Our main contributions are as follows.
(1) We formalize the notion of event-level uncertainty, provide a possibleworlds semantics, and contrast it with source-level uncertainty. From the possibleworlds semantics, we obtain definitions for frequent sequences under two measures: expected support and probabilistic frequentness. These measures were used earlier for FIM in probabilistic data [5], [6]. For SPM [14] addressed source-level uncertainty with the expected support; the other three combinations are new.
(2) We discuss the computational complexity of the fundamental question:

Given a sequence $s$ and a probabilistic database, is $s$ frequent?
In the framework described in [10], which includes not only SPM but also FIM and a host of other database optimization and machine learning problems, the above question is the interestingness or quality predicate. As noted in [10], given such a predicate as a "black box", one can embed it into a variety of candidate generate-and-test frameworks for finding not only frequent sequential patterns, but also maximal frequent sequential patterns. Many popular classical SPM algorithms such as GSP [16], SPADE [23] or SPAM [4] all fit into this framework, and algorithms such as PrefixSpan [15], which do not explicitly generate candidates, also implicitly evaluate the interestingness predicate.

Whilst the interestingness predicate can often be readily evaluated in polynomial time in the deterministic setting, this is not so clear in the probabilistic setting, as a naive approach would involve enumerating exponentially many possible worlds. We show a significant difference from a complexity-theoretic viewpoint between the four variants of the problem (two models and two measures); whilst for three variants, it takes polynomial time to check if a sequence is frequent, for the fourth, this task is \#P-complete, and thus no polynomial-time algorithms are likely to exist.
Related Work. Classical SPM has been studied extensively [16], [23], [15], [4]. Modelling uncertain data as probabilistic databases [17], [1] has led to several ranking/mining problems being studied in this context. The top- $k$ problem (a ranking problem) was intensively studied recently (see [12], [24], [7] and references therein). In particular [7] highlights subtle issues in the semantics of

[^0]the term "top- $k$ " when applied to probabilistic databases. FIM in probabilistic databases was studied under the expected support measure in [2], [6]; and under the probabilistic frequentness measure in [5]. A dichotomy result, showing that certain queries on probabilistic databases are either \#P-complete or in PTIME was given in [8], but this applies to "tuple-level independent" probabilistic structures, which do not model SPM (upper bounds) or source-level uncertainty for SPM (for hardness). To the best of our knowledge, apart from [14], the SPM problem in probabilistic databases has not been studied. However, uncertainty in the time-stamp attribute was considered in [18] - we do not consider time to be uncertain. Also [22] studies SPM in "noisy" sequences, but the model proposed there is very different to ours and does not fit in the probabilistic database framework.

## 2 Problem Statement

### 2.1 Models of Uncertainty

In this section, we give formal definitions of the source-level and event-level uncertain models. We begin by formally defining the classical SPM problem, then the event-level uncertain model. When defining the event-level uncertain model, for simplicity we consider only the tuple-level uncertain case (example in Table 1). We then recap the source-level uncertain case from [14].

Deterministic SPM. We begin by recapitulating the definition of standard SPM. Let $\mathcal{I}=\left\{i_{1}, i_{2}, \ldots, i_{q}\right\}$ be a set of items and $\mathcal{S}=\left\{\sigma_{1}, \ldots, \sigma_{m}\right\}$ be a set of sources. An event $e \subseteq \mathcal{I}$ is a collection of items. A database $D=\left\langle r_{1}, r_{2}, \ldots, r_{n}\right\rangle$ is an ordered list of records such that each $r_{i} \in D$ is of the form $\left(e i d_{i}, e_{i}, \sigma_{i}\right)$, where eid $d_{i}$ is an event-id, $e_{i}$ is an event and $\sigma_{i}$ is a source. A sequence $s=\left\langle s_{1}, s_{2}, \ldots, s_{a}\right\rangle$ is an ordered list of events. Let $s=\left\langle s_{1}, s_{2}, \ldots, s_{q}\right\rangle$ and $t=\left\langle t_{1}, t_{2}, \ldots, t_{r}\right\rangle$ be two sequences. We say that $s$ is a subsequence of $t$, denoted $s \preceq t$, if there exist integers $1 \leq i_{1}<i_{2}<\cdots<i_{q} \leq r$ such that $s_{k} \subseteq t_{i_{j}}$, for $k=1, \ldots, q$. The source sequence corresponding to a source $i$, denoted by $D_{i}$, is just the multiset $\{e \mid(e i d, e, i) \in D\}$, ordered by eid. For a sequence $s$ and source $i$, let $X_{i}(s, D)$ be an indicator variable, whose value is 1 if $s \preceq D_{i}$, and 0 otherwise. The objective is to find all sequences $s$ whose support (Supp) is at least some user-defined threshold $\theta, 1 \leq \theta \leq m$, where:

$$
\begin{align*}
& \text { here: }  \tag{1}\\
& \qquad \operatorname{Supp}(s, D)=\sum_{i=1}^{m} X_{i}(s, D) .
\end{align*}
$$

Event-Level Uncertainty. A probabilistic database $D^{p}$ is a collection of $p$ sequences $D_{1}^{p}, \ldots, D_{m}^{p}$, where $D_{i}^{p}$ is the p-sequence of source $i \in \mathcal{S}$. A p-sequence is the natural analogue of the source sequence in classical SPM, and each psequence is of the form $\left\langle\left(e_{1}, c_{1}\right) \ldots\left(e_{k}, c_{k}\right)\right\rangle$, where the $e_{j}$ 's are events (ordered by eid) and $c_{j}$ is the confidence that $e_{j}$ actually occurred. In examples, we write a p-sequence $\langle(\{a, d\}, 0.4),(\{a, b\}, 0.2)\rangle$ as $(a, d: 0.4)(a, b: 0.2)$. An example of an event-level uncertain database is in Table 1(L). The possible worlds semantics
of $D^{p}$ is as follows. For each event $e_{j}$ in a p-sequence $D_{i}^{p}$ there are two kinds of worlds; one in which $e_{j}$ occurs and the other where it does not. Let occur $=\left\{x_{1}, \ldots, x_{l}\right\}$, where $1 \leq x_{1}<\ldots<x_{l} \leq k$, be the indices of events that occur in $D_{i}^{*}$. Then $D_{i}^{*}=\left\langle e_{x_{1}}, \ldots, e_{x_{l}}\right\rangle$, and $\operatorname{Pr}\left(D_{i}^{*}\right)=\prod_{j \in o c c u r} c_{j} * \prod_{j \notin o c c u r}\left(1-c_{j}\right)$. In other words, we assume that the events in a p-sequence are stochastically independent. For example, all possible worlds of $D_{Y}^{p}$ are shown in Table 1(R) along with detailed probabilities. The set of possible worlds of $D_{i}^{p}$, denoted by $P W\left(D_{i}^{p}\right)$, is obtained by taking all possible $2^{l}$ alternatives for occur, and we say $P W\left(D^{p}\right)=P W\left(D_{1}^{p}\right) \times \ldots \times P W\left(D_{m}^{p}\right)$. The full set of possible worlds for each source is shown in Table $2(\mathrm{~L})$. For any $D^{*} \in P W\left(D^{p}\right)$ such that $D^{*}=\left(D_{1}^{*}, \ldots, D_{m}^{*}\right)$, the probability of $D^{*}$ is given by $\operatorname{Pr}\left[D^{*}\right]=\prod_{i=1}^{m} \operatorname{Pr}\left(D_{i}^{*}\right)$, i.e. we assume that the p-sequences of all sources are mutually independent. An example possible world $D^{*}$ is shown in Table 2(R), which is obtained by taking one possible world each from the worlds of every p-sequences in Table 2(L). The probability of $D^{*}$ is the product of probabilities of all the worlds in it, $\operatorname{Pr}\left[D^{*}\right]$ $=0.29 \times 0.32 \times 0.35=0.03$.
Table 1. An event-level uncertain database (L), and all possible worlds of $D_{Y}^{p}$ along with their probabilities ( R ).

$$
\begin{array}{|l|c|}
\hline & \text { p-sequence } \\
\hline D_{X}^{p} & (a, d: 0.6)(a, b: 0.3)(b, c: 0.7) \\
\hline D_{Y}^{p} & (a, d: 0.4)(a, b: 0.2) \\
\hline D_{Z}^{p} & (a: 1.0)(a, b: 0.5)(b, c: 0.3) \\
\hline
\end{array} \quad \begin{array}{|c|l|}
\hline\{(a, d)\} & (1-0.4) \times(1-0.2)=0.48 \\
\hline\{(a, d)(a, b)\} & (0.4) \times(0.2)=0.08 \\
\hline
\end{array}
$$

Table 2. The set of possible worlds for every p-sequence in Table $1(\mathrm{~L})$. In all, there are $8 \times 4 \times 4=128$ possible worlds of $D^{p}$, one such possible world is shown (R).

| $P W\left(D_{X}^{p}\right)$ | $\{\rangle=0.084\} ;\{(a, d)=0.126\} ;\{(a, b)=0.036\} ;\{(b, c)=0.196\} ;$ <br>  <br>  <br>  <br> $\{(a, d)(a, b)=0.054\} ;\{(a, d)(b, c)=0.294\} ;$ |
| :--- | :--- |
| $P W\left(D_{Y}^{p}\right)$ | $\{\rangle=0.48\} ;\{(a, d)=0.32\} ;\{(a, b)=0.12\} ;\{(a, d)(a, b)=0.08\}$ |
| $P W\left(D_{Z}^{p}\right)$ | $\{(a)=0.35\} ;\{(a)(a, b)=0.35\} ;\{(a)(b, c)=0.15\} ;$ <br>  <br>  <br> $\{(a)(a, b)(b, c)=0.15\}$ |


| $D_{X}^{*}$ | $\{(a, d)$ <br> $(\mathrm{b}, \mathrm{c})\}$ | 0.294 |
| :---: | :--- | :---: |
| $D_{Y}^{*}$ | $\{(a, d)\}$ | 0.32 |
| $D_{Z}^{*}$ | $\{(a)$ <br> $(a, b)\}$ | 0.35 |

Source-Level Uncertainty. A probabilistic database $D^{p}$ is an ordered list of records $\left\langle r_{1}, \ldots, r_{n}\right\rangle$ of the form (eid, $e, W$ ) where eid is an event-id, $e$ is an event and $W$ is a probability distribution over $\mathcal{S}$. The distribution $W$ contains pairs of the form $(\sigma, c)$, where $\sigma \in \mathcal{S}$ and $0<c \leq 1$ is the confidence that the event $e$ is associated with source $\sigma$; we assume $\sum_{(\sigma, c) \in W} c=1$. An example of a sourcelevel uncertain database is in Table $3(\mathrm{~L})$. A possible world $D^{*}$ of $D^{p}$ is generated by taking each event $e_{i}$ in turn, and assigning it to one of the possible sources $\sigma_{i} \in W_{i}$, where $\sigma_{i} \in \mathcal{S}$. Thus every record $r_{i}=\left(e i d_{i}, e_{i}, W_{i}\right) \in D^{p}$ takes the form $r_{i}^{\prime}=\left(e i d_{i}, e_{i}, \sigma_{i}\right)$, for some $\sigma_{i} \in \mathcal{S}$ in $D^{*}$. By enumerating all such possible combinations, we get the complete set of possible worlds. Assuming that the distributions associated with each record $r_{i}$ in $D^{p}$ are stochastically independent, the probability of a possible world $D^{*}$ is $\operatorname{Pr}\left[D^{*}\right]=\prod_{i=1}^{n} \operatorname{Pr}_{W_{i}}\left[\sigma_{i}\right]$. For example, a possible world $D^{*}$ for the database of Table 3 can be generated by assigning events $e_{1}, e_{3}$ and $e_{4}$ to $X$ with probabilities $0.6,0.3$ and 0.7 respectively, and $e_{2}$ to $Z$ with probability 1.0 , and $\operatorname{Pr}\left[D^{*}\right]=0.6 \times 1.0 \times 0.3 \times 0.7=0.126$. The complete set of possible worlds for database of Table 3 is in Table 4.

Of course, a source-level uncertain database can be transformed into a collection of p-sequences, and it will be useful to do so. Specifically, for $i=1, \ldots, m$, let $D_{i}^{p}$ be a list of those events in $D^{p}$ that have non-zero confidence of being associated with source $i$, ordered by eid, together with the associated confidence. However, the resulting p-sequences are not independent; as shown in Table 3. Thus, one may view a source-level uncertain database as a collection of p-sequences with dependencies in the form of x -tuples [7].

Table 3. A source-level uncertain database (L) transformed to p-sequences (R). Note that the $p$-sequence representation is the same as the event-level uncertain database of Table 1. However, events like $e_{1}$ (marked with $\dagger$ on (R)) can only be associated with one of the sources $X$ and $Y$ in any possible world.

| eid | event | $W$ |
| :---: | :---: | :---: |
| $e_{1}$ | $(a, d)$ | $(X: 0.6)(Y: 0.4)$ |
| $e_{2}$ | $(a)$ | $(Z: 1.0)$ |
| $e_{3}$ | $(a, b)$ | $(X: 0.3)(Y: 0.2)(Z: 0.5)$ |
| $e_{4}$ | $(b, c)$ | $(X: 0.7)(Z: 0.3)$ |


|  | p-sequence |
| :---: | :---: |
| $D_{X}^{p}$ | $(a, d: 0.6)^{\dagger}(a, b: 0.3)(b, c: 0.7)$ |
| $D_{Y}^{p}$ | $(a, d: 0.4)^{\dagger}(a, b: 0.2)$ |
| $D_{Z}^{p}$ | $(a: 1.0)(a, b: 0.5)(b, c: 0.3)$ |

Table 4. All possible worlds for the database of Table 3 along with their probabilities. The right-most column shows the support of the sequence (a)(b) in each possible world.

| $D^{*}$ | $X$ | $Y$ | $Z$ | $\operatorname{Pr}\left(D^{*}\right)$ |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| $D_{1}^{*}$ | $(a, d: 0.6)(a, b: 0.3)(b, c: 0.7)$ | $\rangle$ | $(a: 1.0)$ | 0.126 | 1 |
| $D_{2}^{*}$ | $(a, d: 0.6)(a, b: 0.3)$ | $\rangle$ | $(a: 1.0)(b, c: 0.3)$ | 0.054 | 2 |
| $D_{3}^{*}$ | $(a, d: 0.6)(b, c: 0.7)$ | $(a, b: 0.2)$ | $(a: 1.0)$ | 0.084 | 1 |
| $D_{4}^{*}$ | $(a, d: 0.6)$ | $(a, b: 0.2)$ | $(a: 1.0)(b, c: 0.3)$ | 0.036 | 1 |
| $D_{5}^{*}$ | $(a, d: 0.6)(b, c: 0.7)$ | $\rangle$ | $(a: 1.0)(a, b: 0.5)$ | 0.210 | 2 |
| $D_{6}^{*}(a, d: 0.6)$ | $\rangle$ | $(a: 1.0)(a, b: 0.5)(b, c: 0.3)$ | 0.090 | 1 |  |
| $D_{7}^{*}$ | $(a, b: 0.3)(b, c: 0.7)$ | $(a, d: 0.4)$ | $(a: 1.0)$ | 0.084 | 1 |
| $D_{8}^{*}$ | $(a, b: 0.3)$ | $(a, d: 0.4)$ | $(a: 1.0)(b, c: 0.3)$ | 0.036 | 1 |
| $D_{9}^{*}$ | $(b, c: 0.7)$ | $(a, d: 0.4)(a, b: 0.2)$ | $(a: 1.0)$ | 0.056 | 1 |
| $D_{10}^{*}$ | $\rangle$ | $(a, d: 0.4)(a, b: 0.2)$ | $(a: 1.0)(b, c: 0.3)$ | 0.024 | 2 |
| $D_{11}^{*}$ | $(b, c: 0.7)$ | $(a, d: 0.4)$ | $(a: 1.0)(a, b: 0.5)$ | 0.140 | 1 |
| $D_{12}^{*}\langle \rangle$ | $(a, d: 0.4)$ | $(a: 1.0)(a, b: 0.5)(b, c: 0.3)$ | 0.060 | 1 |  |

### 2.2 Notions of Frequentness

We now use the possible-worlds semantics to give two definitions of frequentness.
Expected Support. As every possible world $D^{*}$ is a (deterministic) database, $\operatorname{Supp}\left(s, D^{*}\right)$ is defined as in Eq. 1. We then define the expected support of a sequence $s$ in $D^{p}$ as follows:

$$
\begin{equation*}
\operatorname{ESupp}\left(s, D^{p}\right)=\sum_{D^{*} \in P W\left(D^{p}\right)} \operatorname{Pr}\left[D^{*}\right] * \operatorname{Supp}\left(s, D^{*}\right) . \tag{2}
\end{equation*}
$$

We illustrate these concepts by computing the expected support of a sequence $s=(a)(b)$ and the event-level uncertain database of Table 1. Since there are too many possible worlds, we do not use Eq. 2 and compute $\operatorname{ESupp}(s)$ as follows. We first compute, for every source, the probability that it supports $s$. E.g. from Table 2, all worlds in $P W\left(D_{Z}^{p}\right)$ support $s$ except the first one, so the probability that $Z$ supports $s$ is $(0.35+0.15+0.15)=0.65$, and the probability that it does not is 0.35 . Similarly, the probabilities that $X$ and $Y$ support $s$ are 0.558 and 0.08 . Now for $i=0,1,2,3$, we use the independence of p -sequences to compute
the probability that exactly $i$ sources support $s$ as shown in Table 5(L), e.g. the probability that $s$ is supported by all three sources is $(0.558 \times 0.08 \times 0.65)$ $=0.029$. Then we get $\operatorname{ESupp}(s)=(0 \times 0.142+\ldots+3 \times 0.029)=1.288$.

To compute $\operatorname{ESupp}(s)$ for the source-level uncertain case (Table 3), we directly use Eq. 2 together with Table 4 (and get that $\operatorname{ESupp}(s)=(1 \times 0.126+2 \times$ $0.054+\ldots+1 \times 0.060)=1.288$. Note that the event- and source-level uncertain databases of Table 1 and 3 , which have the same p-sequence form, have the same expected support, even though the possible worlds are very different.

We now formalize the computational task of finding all frequent sequences in a probabilistic database $D^{p}[14]$ :

Definition 1. Given a probabilistic database $D^{p}$, determine all sequences s such that $\operatorname{ESupp}\left(s, D^{p}\right) \geq \theta$, for some user-specified threshold $\theta$.

Table 5. The SPD for the probabilistic DBs of Table 1 (left) and Table 3 (right).

| No. of sources | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| support probability | 0.142 | 0.456 | 0.372 | 0.029 | | No. of sources | 0 | 1 | 2 | 3 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | support probability | 0.0 | 0.712 | 0.288 | 0.0 |

Probabilistic Frequent Sequences. A criticism of expected support [5] is that the expectation of a random variable is only one of its measures, and it does not provide confidence bounds that the support of a sequence is high. Following [5], we define the notion of probabilistic frequent sequences.

Given a probabilistic database $D^{p}$ and its set of possible worlds $P W\left(D^{p}\right)$, the support probability for a sequence $s$ and support value $k$ is denoted by:

$$
\begin{equation*}
\operatorname{Pr}_{k}(s)=\sum_{D^{*} \in P W\left(D^{p}\right),\left(\operatorname{Supp}\left(s, D^{*}\right)=k\right)} \operatorname{Pr}\left(D^{*}\right) \tag{3}
\end{equation*}
$$

where $\operatorname{Supp}\left(s, D^{*}\right)$ is the support of $s$ in $D^{*}$. In other words, $\operatorname{Pr}_{k}(s)$ is the probability that the support of $s$ is exactly $k$. Next define the support probability distribution (SPD) as the vector $\left\langle\operatorname{Pr}_{0}(s), \ldots, \operatorname{Pr}_{m}(s)\right\rangle$. The SPDs for the databases of Table 1 and Table 3 are shown in Table 5 and are very different, even though the p-sequences are the same, e.g. for the event-level DB in Table 1, $\operatorname{Pr}_{3}(s)=0.029$, but in the source-level DB of Table $3, \operatorname{Pr}_{3}(s)=0$, as no such world exists where all three sources support $s$ (see Table 4). Finally, denote by $\operatorname{Pr}_{\geq \theta}(s)=\sum_{k=\theta}^{m} \operatorname{Pr}_{k}(s)$ the probability that the support of $s$ is at least $\theta$. We now define:

Definition 2. Given a probabilistic database $D^{p}$ and two user-specified thresholds, a support threshold $\theta, 1 \leq \theta \leq m$ and a confidence threshold $\tau \in(0,1]$, find all probabilistic frequent sequences (PFSes), which are sequences s s.t. $\operatorname{Pr}_{\geq \theta}(s) \geq \tau$ (i.e. $s$ is a PFS if it has probability $\geq \tau$ of having support $\geq \theta$ ).

Observe that the SPD gives far more detailed information than expected support; from the SPD of a sequence one can easily compute not only the expected support, but also the variance and higher moments.

## 3 Support Computation

We now discuss the computational complexity of computing frequent sequences based on the definitions in the previous section. As indicated earlier, we focus on the interestingness predicate, which when specialized to our problem and definitions of frequentness, yield the following questions. Given a probabilistic database $D^{p}$ (either source-level or event-level uncertain):

- For a given sequence $s$ and a threshold $\theta$, is $\operatorname{ESupp}\left(s, D^{p}\right) \geq \theta$ ?
- For a given sequence $s$ and thresholds $\theta$ and $\tau$, is $\operatorname{Pr}\left(\operatorname{Supp}\left(s, D^{p}\right) \geq \theta\right) \geq \tau$ ?

We assume that the database, whether source-level or event-level uncertain, is given as a list of p-sequences. Denote by $m$ the number of sources (as before), by $N_{i}$ the number of events in the p-sequence of the $i$-th source, by $N=\sum_{i=1}^{m} N_{i}$ the total size of all p-sequences and by $k$ the number of events in $s .^{2}$

### 3.1 Source Support Probability

The first task is to determine the source support probability, namely the probability that a given source $\sigma_{i}$ supports a sequence $s$, or $\operatorname{Pr}\left(s \preceq D_{i}^{p}\right)$. An important issue is that $s$ may be a subsequence of $D_{i}^{p}$ in many different ways. For example, if $s=\langle\underbrace{\langle(a)(a) \ldots(a)}_{k \text { times }}\rangle$ and $D_{i}^{p}=\left\langle\left(a: c_{1}\right),\left(a, c_{2}\right), \ldots,\left(a, c_{N_{i}}\right)\right\rangle$, then any subset of
$k$ positions from $D_{i}^{p}$ could be the (sole) basis for the $i$-th source to support $s$. In [14] a dynamic programming recurrence is given that computes this probability in polynomial time:

Theorem 1 ([14]). Given s and $D_{i}^{p}$, we can calculate $\operatorname{Pr}\left(s \preceq D_{i}^{p}\right)$ in $O\left(k \cdot N_{i}\right)$ time.

### 3.2 Expected Support

For both event-level and source-level uncertainty, it was shown in [14] that:

$$
\begin{equation*}
\operatorname{ESupp}\left(s, D^{p}\right)=\sum_{i=1}^{m} \operatorname{Pr}\left(s \preceq D_{i}^{p}\right) . \tag{4}
\end{equation*}
$$

This holds even though p-sequences are not independent in the source-level uncertain case, due to the principle of linearity of expectation. To calculate $\operatorname{ESupp}\left(s, D^{p}\right)$ using Eq. 4, we can apply Theorem 1 to each source in turn, which takes $O\left(k \cdot \sum_{i=1}^{m} N_{i}\right)=O(k N)$ time, and obtain:

Theorem 2. Given s and $D^{p}$, we can calculate $\operatorname{ESupp}\left(s, D^{p}\right)$, and hence evaluate the interestingness predicate, in $O(k N)$ time, for both event-level and sourcelevel uncertainty.

[^1]
### 3.3 Probabilistic Frequentness

Event-Level Uncertainty. As in [5], we compute the entire support probability distribution (SPD), namely the vector $\left\langle\operatorname{Pr}_{0}(s), \ldots, \operatorname{Pr}_{m}(s)\right\rangle$, where $\operatorname{Pr}_{k}(s)$ is the probability that the support of $s$ is exactly $k$. Given the SPD, we then compute the cumulative probabilities $\operatorname{Pr}_{\geq \theta}(s)=\sum_{k=\theta}^{m} \operatorname{Pr}_{k}(s)$ in $O(m)$ time and thereby answer the interestingness predicate.

In the case of event-level uncertainty, the p-sequences are independent. This allows the SPD to be calculated as a dynamic programming recurrence as follows. We first compute $\operatorname{Pr}\left(s \preceq D_{i}^{p}\right)$ for all sources $i$ in $O(k N)$ time as in Thm. 2. Next, we define $\operatorname{Pr}_{i, j}(s)$, for $0 \leq i, j \leq m$, as the probability that exactly $i$ of the first $j$ sources support $s$. We then use the formula:

$$
\begin{equation*}
\operatorname{Pr}_{i, j}(s)=\operatorname{Pr}_{i-1, j-1}(s) \cdot \operatorname{Pr}\left(s \preceq D_{i}^{p}\right)+\operatorname{Pr}_{i, j-1}(s) \cdot\left(1-\operatorname{Pr}\left(s \preceq D_{i}^{p}\right)\right), \tag{5}
\end{equation*}
$$

where $\operatorname{Pr}_{0, j}(s)=1,0 \leq j \leq m$ and $\operatorname{Pr}_{i, j}(s)=0, \forall i>j$, to compute all the values $\operatorname{Pr}_{i, j}$ in $O\left(m^{2}\right)$ time. Since $\operatorname{Pr}_{i, m}(s)=\operatorname{Pr}_{i}(s)$, we get the full SPD and can use this to determine if $s$ is a PFS; the overall time is $O\left(k N+m^{2}\right)$. In summary:

Theorem 3. Given $s$ and $D^{p}$, if $D^{p}$ is an event-level uncertain database, we can calculate the support probability distribution, and hence the interestingness predicate in $O\left(k N+m^{2}\right)$ time.

Remark 1. When computing $\operatorname{Pr}_{i, j}(s)$ using Eq. 5, we consider two cases: either $\sigma_{i}$ supports $s$, and exactly $j-1$ of $\sigma_{1}, \ldots, \sigma_{i-1}$ support $s$ (the first term) or $\sigma_{i}$ does not support $s$ and exactly $j$ of $\sigma_{1}, \ldots, \sigma_{i-1}$ support $s$. The correctness of Eq. 5 depends crucially on the fact that we assume independence among p -sequences, so $\operatorname{Pr}\left(s \preceq D_{i}^{p}\right)\left(\right.$ resp. $\left.\operatorname{Pr}\left(s \npreceq D_{i}^{p}\right)\right)$ is unchanged even when conditioned on knowing that exactly $j-1$ (resp. $j$ ) of the first $i-1$ sources support $s$.

Remark 2. Since $N / m$ is the average length of a p-sequence, $N / m \ll m$ and (since $k$ is not very large as well) the $m^{2}$ term will often dominate the $k N$ term. This means the interestingness predicate for probabilistic frequentness will often be computationally more expensive than for expected support.

Source-Level Uncertainty. In source-level uncertainty, we cannot use Eq. 5 to efficiently compute the SPD, since the p-sequences are not independent (see Remark 1). Consider the very simple probabilistic database which consists of just the event $\{a\}$, associated with source $\sigma_{1}$ and $\sigma_{2}$ with probabilities 0.5 each. There are only two possible worlds, the first with the event $\{a\}$ associated with $\sigma_{1}$ and nothing with $\sigma_{2}$, and the second the other way around. Clearly, if $s$ is the sequence $\langle\{a\}\rangle$, then $\operatorname{Pr}\left(s \preceq D_{1}^{p}\right)=\operatorname{Pr}\left(s \preceq D_{2}^{p}\right)=0.5$. However, applying Eq. 5 gives that $\operatorname{Pr}_{1,1}=0.5$ (correct) and $\operatorname{Pr}_{2,2}=0.25$ (incorrect). The probability that both sources support $s$ is zero, not 0.25 - there is no possible world in which both sources support $s$. To see what goes wrong, consider the computation of $\operatorname{Pr}_{2,2}(s)$ as $\operatorname{Pr}_{1,1}(s) \cdot 0.5+\operatorname{Pr}_{2,1}(s) \cdot 0.5=0.5 \cdot 0.5+0 \cdot 0.5=0.25$. Looking at the first term, if $\sigma_{1}$ supports $s$, then conditioned on this knowledge, the probability that $\sigma_{2}$ supports $s$ is zero. Thus, the correct computation for $\operatorname{Pr}_{2,2}(s)$ would be


| eid | event | $W$ |
| :---: | :---: | :---: |
| $e_{1}$ | a | $\left(\sigma_{1}: 0.5\right)\left(\sigma_{2}: 0.5\right)$ |
| $e_{2}$ | a | $\left(\sigma_{1}: 0.33\right)\left(\sigma_{2}: 0.33\right)\left(\sigma_{3}: 0.33\right)$ |
| $e_{3}$ | a | $\left(\sigma_{1}: 0.5\right)\left(\sigma_{3}: 0.5\right)$ |

(a)

Fig. 1. A sample bipartite graph (a) transformed to a probabilistic database (b).


Fig. 2. Two possible worlds in $P W\left(D^{p}\right)$ and their bipartite graph representations. A perfect matching (b), when every vertex in $U \cup V$ is adjacent to a single edge.
as $0.5 \cdot 0+0 \cdot 0.5=0$. Unfortunately, the difficulty is not with one particular approach but is intrinsic to the problem: we now show that computing even a single entry of the SPD is provably hard. Define the following problem:
Definition 3 (Exact- $k$ SUPport problem). The input is a source-level uncertain probabilistic database $D^{p}$, a sequence $s$ and a number $k, 0 \leq k \leq m$. The output is $\operatorname{Pr}_{k}(s)$, i.e. the probability that exactly $k$ sources support $s$.

Theorem 4. Exact- $k$ Support is \#P-complete.
Proof. We reduce the problem of computing the number of perfect matchings in a bipartite graph, a known \#P-complete problem [19] to Exact- $k$ Support.

Let $G(U, V, E)$ be an undirected bipartite graph, where $U$ and $V$ are disjoint sets of vertices and $E$ is the set of edges between them, $E \subseteq U \times V$. We assume that $|U|=|V|=n$. A perfect matching $M$ is a subset of edges in $E$ such that each vertex in $U \cup V$ is adjacent to exactly a single edge in $M$. Given a bipartite graph $G$, the problem of counting how many perfect matchings there are in $G$ is \#P-complete [19].

Given a bipartite graph $G$, to compute the number of matchings in $G$, we create an instance of Exact- $k$ Support (a probabilistic database $D^{p}$, a sequence $s$ and a number $k$ ) in polynomial time such that solving the latter instance gives the number of perfect matchings in $G$. Given $G=(U, V, E)$ where $U=\left\{u_{1}, \ldots, u_{n}\right\}$ and $V=\left\{v_{1}, \ldots, v_{n}\right\}$, we create a set of sources $\mathcal{S}=\left\{\sigma_{1}, \ldots, \sigma_{n}\right\}$ such that $\sigma_{i} \in \mathcal{S}$ represents $u_{i} \in U$. The probabilistic database $D^{p}$ is a set of records $r_{i}=\left(e_{i}, e, W_{i}\right)$, where $e_{i}$ is a event id, $e$ is an event and $W_{i}$ is a distribution. The record $r_{i}$ represents $v_{i}$ in $V$ together with all the edges from $v_{i}$ to the neighborhood of $v_{i}$, i.e. the set of vertices in $U$ adjacent
to $v_{i}$. In what follows, we denote the neighborhood of $v_{i}$ as $N\left(v_{i}\right)$. The event contained in every record is a set containing just the singleton element $\{a\}$. In the $i$-th record $r_{i}$, the distribution $W_{i}$ contains only the sources $\sigma_{j}$ that represent vertices $u_{j} \in N\left(v_{i}\right)$. All sources in $W_{i}$ have the same probability, i.e. $\left(1 /\left|N\left(v_{i}\right)\right|\right)$. An example of such a transformation is shown in Fig. 1. Finally, we choose $k=0$ and the sequence $s=(a)(a)$, and ask to compute $\operatorname{Pr}_{0}(s)$, i.e. the probability that no sources support $s$. This completes creation of the instance of Exact- $k$ Support, and the transformation is clearly polynomial-time. Clearly, every possible world $D^{*} \in P W\left(D^{p}\right)$ is equally likely, and the probability of a possible world $D^{*}$ is $\phi=\left(1 /\left|P W\left(D^{p}\right)\right|\right)$, where $\left|P W\left(D^{p}\right)\right|=\prod_{i=1}^{n}\left|N\left(v_{i}\right)\right|$. For example, there are 12 possible worlds for the database in Fig. 1(b), and the probability of each world is $(1 / 12)$. In each possible world, each record $r_{i}$ is associated with some source $\sigma_{j}$, so a possible world can be viewed as a sub-graph of $G$ where every vertex in $V$ has degree exactly one. Two possible worlds and their corresponding graphs are shown in Fig. 2. Those possible worlds where each source is associated with exactly one record corresponds to a perfect matching (Fig. 2(b)); in such possible worlds, the support of the sequence $s=(a)(a)$ will clearly be zero. Thus, we see that $\operatorname{Pr}_{0}(s)=\phi \cdot(\#$ matchings in $G$ ), and once we are given $\operatorname{Pr}_{0}(s)$, we obtain the number of matchings in $G$ by multiplying by the total number of possible worlds. For example, in database in 1(b), there are only three possible worlds where each source is associated with exactly one event, which are $\left(\sigma_{1}: e_{1}, \sigma_{2}: e_{2}, \sigma_{3}: e_{3}\right),\left(\sigma_{1}: e_{2}, \sigma_{2}: e_{1}, \sigma_{3}: e_{3}\right)$ and $\left(\sigma_{1}: e_{3}, \sigma_{2}: e_{1}, \sigma_{3}: e_{2}\right)$. Hence, $\left.\operatorname{Supp}\left(s, D^{*}\right)=0\right)$ in three worlds and therefore, $\operatorname{Pr}_{0}(s)=3 \times(1 / 12)=0.25$. We multiply the answer by the number of possible worlds, 12 , to get 3 , the number of perfect matchings in $G$.

Hence, we have shown that if the value $\operatorname{Pr}_{k}(s)$ can be computed for $s=$ (a) (a) and $k=0$ in $D^{p}$, we can also find number of perfect matchings in $G$, thus reducing the problem of counting perfect matchings in a bipartite graph to Exact- $k$ Support, and showing that Exact- $k$ Support is \#P-complete.

## 4 Conclusions and Future Work

We studied uncertainty models for sequential pattern mining and defined the notions of frequentness under the expected suppot and probabilistic frequentness measures. We elaborated on these measures from complexity-theoretic viewpoint, and discussed the computational cost of evaluating these measure for our considered models. Thus we were able to show, that whilst dynamic programming could be used to find frequent sequences efficiently, computing probabilistic frequentness for source-level uncertainty is \#-P complete. An empirical evaluation and comparison of our considered measures in terms of computational cost and quality of the solution should be an interesting direction to explore.

## References

1. Aggarwal, C.C. (ed.): Managing and Mining Uncertain Data. Springer (2009)
2. Aggarwal, C.C., Li, Y., Wang, J., Wang, J.: Frequent pattern mining with uncertain data. In: KDD. pp. 29-38 (2009)
3. Agrawal, R., Srikant, R.: Mining sequential patterns. In: ICDE. pp. 3-14 (1995)
4. Ayres, J., Flannick, J., Gehrke, J., Yiu, T.: Sequential pattern mining using a bitmap representation. In: KDD. pp. 429-435 (2002)
5. Bernecker, T., Kriegel, H.P., Renz, M., Verhein, F., Züfle, A.: Probabilistic frequent itemset mining in uncertain databases. In: KDD. pp. 119-128 (2009)
6. Chui, C.K., Kao, B., Hung, E.: Mining frequent itemsets from uncertain data. In: PAKDD. pp. 47-58 (2007)
7. Cormode, G., Li, F., Yi, K.: Semantics of ranking queries for probabilistic data and expected ranks. In: ICDE. pp. 305-316 (2009)
8. Dalvi, N.N., Suciu, D.: The dichotomy of conjunctive queries on probabilistic structures. In: PODS. pp. 293-302 (2007)
9. El-Ramly, M., Stroulia, E., Sorenson, P.G.: From run-time behavior to usage scenarios: an interaction-pattern mining approach. In: KDD. pp. 315-324 (2002)
10. Gunopulos, D., Khardon, R., Mannila, H., Saluja, S., Toivonen, H., Sharm, R.S.: Discovering all most specific sentences. ACM ToDS 28(2), 140-174 (2003)
11. Hassanzadeh, O., Miller, R.J.: Creating probabilistic databases from duplicated data. The VLDB Journal 18(5), 1141-1166 (2009)
12. Hua, M., Pei, J., Zhang, W., Lin, X.: Ranking queries on uncertain data: a probabilistic threshold approach. In: SIGMOD Conference. pp. 673-686 (2008)
13. Khoussainova, N., Balazinska, M., Suciu, D.: Probabilistic event extraction from rfid data. In: ICDE. pp. 1480-1482 (2008)
14. Muzammal, M., Raman, R.: Mining sequential patterns from probabilistic databases. Tech. Rep. CS-10-002, Dept. of Computer Science, Univ. of Leicester (2010), available from http://www.cs.le.ac.uk/people/mm386/pSPM.pdf
15. Pei, J., Han, J., Mortazavi-Asl, B., Wang, J., Pinto, H., Chen, Q., Dayal, U., Hsu, M.: Mining sequential patterns by pattern-growth: The prefixspan approach. IEEE Trans. Knowl. Data Eng. 16(11), 1424-1440 (2004)
16. Srikant, R., Agrawal, R.: Mining sequential patterns: Generalizations and performance improvements. In: EDBT. pp. 3-17 (1996)
17. Suciu, D., Dalvi, N.N.: Foundations of probabilistic answers to queries. In: SIGMOD Conference. p. 963 (2005)
18. Sun, X., Orlowska, M.E., Li, X.: Introducing uncertainty into pattern discovery in temporal event sequences. In: ICDM. pp. 299-306 (2003)
19. Valiant, L.G.: The complexity of computing the permanent. Theor. Comput. Sci. 8, 189-201 (1979)
20. Wang, K., Xu, Y., Yu, J.X.: Scalable sequential pattern mining for biological sequences. In: CIKM. pp. 178-187 (2004)
21. Wikipedia: http://en.wikipedia.org/wiki/anpr - Wikipedia, the free encyclopedia (2010), \url \{http://en.wikipedia.org/wiki/ANPR\}, [Online; accessed 30-April2010]
22. Yang, J., 0010, W.W., Yu, P.S., Han, J.: Mining long sequential patterns in a noisy environment. In: SIGMOD Conference. pp. 406-417 (2002)
23. Zaki, M.J.: Spade: An efficient algorithm for mining frequent sequences. Machine Learning 42(1/2), 31-60 (2001)
24. Zhang, Q., Li, F., Yi, K.: Finding frequent items in probabilistic data. In: SIGMOD Conference. pp. 819-832 (2008)

[^0]:    ${ }^{1}$ Some formulations of this and similar problems may even exhibit SLU.

[^1]:    ${ }^{2}$ We assume that an event consists of at most a constant number of items.

