



*Consiglio Nazionale delle Ricerche
Istituto di Calcolo e Reti ad Alte Prestazioni*

Discovering Expressive Process Models by Clustering Log Traces

Gianluigi Greco¹, Antonella Guzzo², Luigi
Pontieri², Domenico Saccà³

RT-ICAR-CS-06-04

Aprile 2006



Consiglio Nazionale delle Ricerche, Istituto di Calcolo e Reti ad Alte Prestazioni (ICAR)
– Sede di Cosenza, Via P. Bucci 41C, 87036 Rende, Italy, URL: www.icar.cnr.it



Consiglio Nazionale delle Ricerche
Istituto di Calcolo e Reti ad Alte Prestazioni

Discovering Expressive Process Models by Clustering Log Traces

Gianluigi Greco¹, Antonella Guzzo², Luigi
Pontieri², Domenico Saccà³

Rapporto Tecnico N.: 1
RT-ICAR-CS-06-04

Data:
Aprile 2006

¹ Dipartimento di Matematica, Università degli Studi della Calabria, via P. Bucci 30B, Rende (CS)

² Istituto di Calcolo e Reti ad Alte Prestazioni, ICAR-CNR, Sede di Cosenza, via P. Bucci 41C, 87036 Rende(CS)

³ DEIS, Università degli Studi della Calabria, via P. Bucci 41C, Rende (CS)

I rapporti tecnici dell'ICAR-CNR sono pubblicati dall'Istituto di Calcolo e Reti ad Alte Prestazioni del Consiglio Nazionale delle Ricerche. Tali rapporti, approntati sotto l'esclusiva responsabilità scientifica degli autori, descrivono attività di ricerca del personale e dei collaboratori dell'ICAR, in alcuni casi in un formato preliminare prima della pubblicazione definitiva in altra sede.

Discovering Expressive Process Models by Clustering Log Traces

Gianluigi Greco¹, Antonella Guzzo², Luigi Pontieri², Domenico Saccà³

¹Gianluigi Greco is at Dipartimento di Matematica, University of Calabria, Via Bucci 30b, Rende (CS), Italy.

Phone: +39 0984 496429. Fax: +39 0984 496410. E-mail: ggreco@mat.unical.it.

Corresponding author.

²Antonella Guzzo is at ICAR-CNR, Via Bucci 41c, Rende (CS), Italy.

Phone: +39 0984 494618. Fax: +39 0984 839054. E-mail: guzzo@icar.cnr.it.

²Luigi Pontieri is at ICAR-CNR, Via Bucci 41c, Rende (CS), Italy.

Phone: +39 0984 831733. Fax: +39 0984 839054. E-mail: pontieri@icar.cnr.it.

³Domenico Saccà is at ICAR-CNR and DEIS, University of Calabria, Via Bucci 41c, Rende (CS), Italy.

Phone: +39 0984 494750. Fax: +39 0984 839054. E-mail: sacca@unical.it.

A preliminary version of portions of this paper appeared in *Proc. of the 8th Pacific-Asia Conference "Advances in Knowledge Discovery and Data Mining", PAKDD 2004*, Sydney, Australia, May 26-28, 2004, LNCS 3056, pp. 52-62.

Discovering Expressive Process Models

by Clustering Log Traces

Gianluigi Greco, Antonella Guzzo, Luigi Pontieri, Domenico Saccà

Abstract

Process mining techniques have recently received considerable attention in the literature for their ability to assist the (re)design of complex processes, by means of algorithms that automatically discover models capable to explain all the events registered in some log traces provided as input. Following this line of research, the paper investigates an extension of such basic approaches where the identification of different variants for the process is explicitly accounted for, by means of a technique for clustering log traces. Indeed, modelling each group of similar executions with a different schema is the proposed solution to single out “conformant” models for which, in particular, the number of the possible enactments that are however extraneous to the process semantics is minimized.

Therefore, a novel framework for process mining is introduced, and some relevant computational issues are deeply studied. Since it is proven that finding an exact solution to such an enhanced process mining problem requires high computational costs in most practical cases, a greedy approach is devised on the basis of an iterative, hierarchical refinement of the process model, in which, at each step, traces sharing similar behaviors are clustered together and eventually equipped with a specialized schema. The algorithm is designed in a sophisticated way which guarantees that each refinement leads to a sounder schema, so that a monotonic search is attained. Experimental results evidence the validity of the proposed approach with respect to both effectiveness and scalability.

Keywords: H.2.m: Process Mining; H.2.8.d: Data Mining; H.2.4.p Workflow management; H.2.8.b: Clustering, classification, and association rules.

I. INTRODUCTION

Even though workflow management systems are more and more utilized in enterprises, their actual impact in managing complex processes is still limited by the difficulties encountered in the design phase. Indeed, processes have often complex dynamics whose modelling requires expensive and long analysis, which may be eventually unfeasible under an economic viewpoint. Therefore, several approaches have been recently proposed in the literature to accommodate the design of complex workflows by means of *process mining* techniques (see Section VII, for an overview of different proposals).

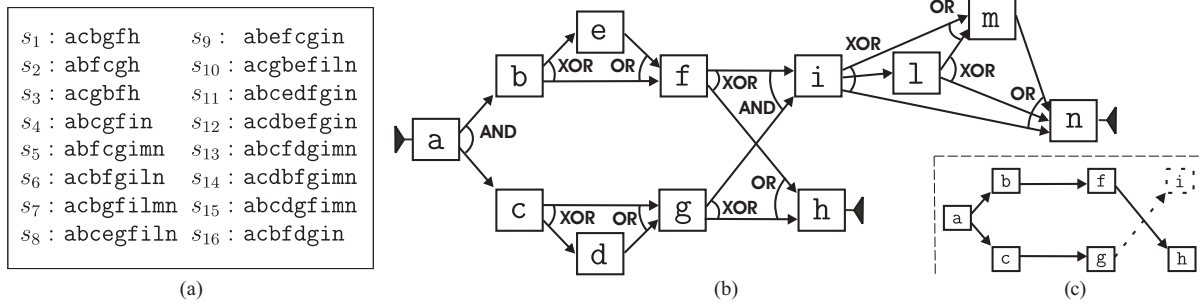


Fig. 1. An example process: (a) some traces; (b) workflow schema WS_0 ; (c) an example instance.

These approaches are devoted to automatically derive a model that can explain all the episodes recorded in the event *log* (a.k.a. transactional log or audit trail [1]), which is in fact collected and stored by most information systems while the activities of a given process are executed — this is the case of transactional information systems such as WFM, ERP, CRM, SCM, and B2B systems. Eventually, the “mined” model is used to (re)design a detailed workflow schema, capable of supporting further coming enactments.

Example 1.1: As a sample applicative scenario, we shall consider throughout the paper the automatization of the (*OrderManagement*) process of handling customers’ orders within a business company. The process consists of the following activities: (a) receiving the order, (b) authenticating the client, (c) checking in the stock the availability of the required product, (d) verifying the availability of external supplies, (e) registering the client in the company database, (f) evaluating the trustworthiness of the client, (g) evaluating the plan of the production, (h) rejecting the order, (i) accepting the order, (l) contacting the mail department in order to speed up the shipment of the goods, (m) applying some discount, and (n) preparing the bill.

In this scenario, some execution traces may be available, for they are collected by the information systems of the business company (e.g., the log in Figure 1.(a)). And, in fact, the aim of process mining techniques is to look at these traces and to automatically identify a model for the *OrderManagement* process. \triangleleft

Actually, since several alternative models may be possibly singled out for any fixed log, techniques have been proposed for identifying the most *conformant* one, i.e., the model which is the best “aligned” with the log. In this respect, the most widely considered requirement for conformance is the *completeness* of the model (*fitness* in [2]) that is, roughly speaking, the percentage of the *traces* stored in the log that may be the result of

some enactment supported by the mined model. Then, differences among various proposals for process mining often come in the language used for representing process models, and in the specific algorithms used for mining the best fitting one.

In this paper, we continue on the way of enhancing the process mining framework by particularly taking care of the following two key issues:

- (i) A good completeness of the model does not necessarily entail its conformance with the underlying process P . For instance, a fully complete model may support not only all the traces provided in input, but also an arbitrary number of execution patterns that are not registered in the log. In many cases, such a situation might strongly undermine the significance of the mined model w.r.t. the application needs. Therefore, it is recommended to focus on models that are not too generic and that appropriately encode the behavior of the process. This can be accomplished by taking care of another parameter, called *soundness* (*minimality* or *behavioral appropriateness* in [2]), which measures the percentage of enactments of the mined model that do find some correspondence in the log. Intuitively, the lower is the factor, the more extraneous (w.r.t. the process semantics) executions the model supports.
- (ii) The output of process mining techniques is usually too complex for being practically exploited by the designer, as it emerged in several real applications domains (see, e.g., [3], [4]). This aspect is particularly relevant in the case of processes consisting of a large number of activities, with several specific execution-scenarios and complex behavioral rules, for which the mined models may be made of hundreds of tasks and constraints. Nonetheless, very few efforts have been paid to support some kinds of abstraction (e.g., techniques in [5], [6], [7] are able to produce models that focus on the main behavior as reported in the log, by properly dealing with noise). In this respect, a viable solution for a better understanding of the process semantics is to take care of the existence of variants of the process, so that each variant may be independently mined and equipped with a specialized model. Clearly, the most relevant variants should be identified only, in order to prevent the model for becoming too much specific. Then, suitable strategies supporting some kinds of trade-off in the identification of the variants and, possibly, the ability of incrementally refining the model should be conceived.

Though some of the approaches in the literature already exploit some minimality conditions for increasing the soundness (e.g., [8], [9]), very few efforts have been paid to discover

variants and usage scenarios (cf. [10], [11]). However, the isolation of the variants would allow to make it clear the “implicit” model (cf. [9]) underlying a process, i.e., the knowledge (organizational/behavioral rules) that actors have in mind during the enactments. In particular, if variants for the process are not discovered and properly classified, the mined model will eventually mix different usage scenarios and, hence, will likely be unsound.

A. Overview of the Approach and Contributions

In this paper, we face the above issues by proposing a novel process mining algorithm that is able to cope with complex processes, by discovering conformant (both sound and complete) models. This is carried out by producing a clear and modular representation of the process, where its relevant variants are explicitly singled out by means of a technique for clustering log traces. Specifically, each variant is modelled by means of a distinct workflow schema \mathcal{WS}_i , so that the resulting model for the whole process, called *disjunctive workflow schema*, is precisely the set of all these individual schemas.

Actually, the higher expressiveness of disjunctive schemas comes with a cost. Indeed, the problem of computing conformant disjunctive schemas and the problem of checking whether a given disjunctive schema at hand is sound *enough* turn out to be intractable.

Therefore, we pragmatically face the process mining problem by means of a greedy strategy based on an iterative, hierarchical refinement of a disjunctive schema \mathcal{WS}^\vee , in which, at each step, a schema $\mathcal{WS}_i \in \mathcal{WS}^\vee$ is selected, and the traces supported by \mathcal{WS}_i are partitioned into clusters. Eventually, each cluster is equipped with a refined workflow schema, which is meant to model a specific usage scenario specializing \mathcal{WS}_i . Notably, the algorithm is designed in a sophisticated way guaranteeing that each refinement leads to a sounder schema for the given process, so that a monotonic search is attained.

Moreover, in order to efficiently partition traces into clusters by means of well-known methods, we investigate an approach for producing a “flat” relational representation of the traces, by projecting them onto a set of suitable *features*, which are meant to characterize the traces which are not properly modelled by the current workflow schema being refined. Therefore, the paper also accounts for the definition of an abstract representation for the features, and of an efficient level-wise algorithm for their computation.

We conclude by noticing that all the algorithms above have been implemented, and the results of a thorough experimental activity are also discussed in the paper.

B. Organization

The rest of the paper is organized as follows. In Section II we report a few preliminaries on workflows, processes, and logs. The problem of discovering a disjunctive process model from some log data is introduced in Section III, where some relevant computational issues have been also studied. In Section IV a solution for the basic case where the mined model is non-disjunctive is discussed. Its generalization is reported in Section V, where the clustering-based approach to process mining and the feature extraction algorithm are illustrated. The two next sections, i.e., Section VI and Section VII, are devoted to present results of experiments and major related works in the literature, respectively. A few concluding remarks are, finally, sketched in Section VIII.

II. FORMAL FRAMEWORK: WORKFLOW SCHEMA AND LOGS

In this section, we introduce the basic notions and notations for formally representing process models that will be used throughout the paper. We point out that a comprehensive formalization of all the facets of workflow modelling is beyond the scope of this paper; the reader interested in expanding on this subject is referred to, e.g., [12], [13], [14], [15], [16], [17]. Specifically, the modelling language is kept simple and less expressive than other existing languages, such as *Petri nets* [18] or *event driven process chains* EPCs [19], [20].

For instance, it does not account for loops, i.e., recursive workflow schemas and iterated executions, which can be instead coped with, e.g., by the WF-net model and the α -algorithm [21], and — as commonly done in the process mining framework — it forces all outgoing and ingoing edges of a given node to have the same type. On the other hand, the modelling language allows for the definition of non-free choice constructs, which are generally avoided in other formalizations related to process mining.

Moreover, an important feature of our approach is that there is no conceptual limitation in accommodating more sophisticated models and techniques for their manipulation, because it is to large extent independent of the underlying workflow model, which mainly affects the way each cluster of traces is equipped with a schema (cf. algorithm `MineWorkflow` in Section IV). In this respect, the proposed approach is modular and may allow to use any pre-existing process mining algorithm (provided, as we shall see in the following, it satisfies the properties in Theorem IV.3). And, in fact, the assumptions

exploited in the paper have been required by the necessity of starting from a simple model enjoying these properties, yet covering important features in workflow specification.

Definition II.1: Let P be a process. A *workflow schema* \mathcal{WS} for P is a tuple $\langle A, E, a_0, A_F, \mathbf{Fork}, \mathbf{Join} \rangle$, where A is a finite set of *activities* (also called nodes, or tasks), $E \subseteq (A - A_F) \times (A - \{a_0\})$ is an acyclic relation of precedences among activities (whose elements are simply called edges), $a_0 \in A$ is the starting activity, and $A_F \subseteq A$ is the set of final activities. The tuple $\langle A, E, a_0, A_F \rangle$ is often referred to as the *control flow graph* of \mathcal{WS} : it states the orderings among the activities involved in P that must be respected during their execution.

The functions $\mathbf{Fork} : (A - A_F) \mapsto \{\mathbf{AND}, \mathbf{OR}, \mathbf{XOR}\}$ and $\mathbf{Join} : (A - \{a_0\}) \mapsto \{\mathbf{AND}, \mathbf{OR}\}$, also called *local constraints* in the literature, relate the execution of activities that are neighbor of each other in the control flow graph. A node a with $\mathbf{Join}(a) = \mathbf{OR}$ (resp. \mathbf{AND}) is an *or-join* (resp. *and-join*); moreover, a node a with $\mathbf{Fork}(a) = \mathbf{OR}$ (resp. $\mathbf{AND}, \mathbf{XOR}$) is an *or-fork* (resp. *and-fork, xor-fork*). \square

An example workflow schema is reported in Figure 1.(b) by means of an intuitive graphical notation. Actually, this is a schema (possibly, the result of some process mining algorithm) for the *OrderManagement* process presented in the Introduction.

Each time a workflow is enacted in a workflow management system, it produces an instance, that is a suitable subgraph $I = \langle A_I, E_I \rangle$ of the schema satisfying all the constraints. Intuitively, A_I contains those activities that are executed in I and that, therefore, activate their outgoing edges. Notice that the target of these edges are not required to be executed as well, i.e., some branches of the control flow graph may be only partially executed because a final activity, belonging to another branch, was completed in the meanwhile. Therefore, in the following definition, these kinds of task that are not executed in the instance I are explicitly distinguished and included in the set A'_I .

Definition II.2: Let $\mathcal{WS} = \langle A, E, a_0, A_F, \mathbf{Join}, \mathbf{Fork} \rangle$ be a workflow schema. Let $I' = \langle A_I \cup A'_I, E'_I \rangle$ be a connected subgraph of $\langle A, E \rangle$, such that the following conditions hold:

- (i) $a_0 \in A_I$; (ii) $A_I \cap A_F \neq \emptyset$; (iii) $\{(a, b) \mid (a, b) \in E'_I \wedge a \in A'_I\} = \emptyset$;
- (iv) for each $(a, b) \in E'_I$, $a \in A_I$ and $b \in (A'_I \cup A_I)$;
- (v) for each $a \in A_I \cup A'_I - \{a_0\}$, $\{b \mid (b, a) \in E'_I \wedge b \in A_I\} \neq \emptyset$;
- (vi) for each $a \in (A_I - \{a_0\})$ s.t. $\mathbf{Join}(a) = \mathbf{AND}$, $\{(b, a) \mid (b, a) \in E\} \subseteq E'_I$;

- (vii) for each $a \in (A_I - A_F)$ s.t. $Fork(a) = \text{AND}$, $\{(a, b) \mid (a, b) \in E\} \subseteq E'_I$;
- (viii) for each $a \in (A_I - A_F)$ s.t. $Fork(a) = \text{XOR}$, $|\{(a, b) \mid (a, b) \in E'_I\}| = 1$; and
- (ix) for each $a \in (A_I - A_F)$ s.t. $Fork(a) = \text{OR}$, $|\{(a, b) \mid (a, b) \in E'_I\}| \geq 1$;

Then, the graph $I = \langle A_I, E'_I \cap (A_I \times A_I) \rangle$ is an *instance* of \mathcal{WS} (denoted as $\mathcal{WS} \models I$). \square

An example instance for the *OrderManagement* process is reported in Figure 1.(c), according to a graphical notation where dashed nodes and edges represent activities and precedences that are activated but not included in the instance. In particular, the instance models a case where the order is rejected (activity **h** is executed) because the client is not reliable (checked by **f**). However, given that the required product is available in the stock, the activity **i** is the target of an activated edge, but it is never executed.

We point out that most process-oriented commercial systems only store partial information about process instances, by tracing some events related to the execution of its activities. In particular, the *logs* stored by such systems may range from a simple task sequence (recall, e.g., Figure 1.(a)) to richer formats (evidencing, e.g, the start and/or the completion of a task). We next describe an abstract representation of a process log which is commonly adopted in the literature.

Let A be the set of task identifiers for the process P ; then, a *workflow trace* s over A is a string in A^* , representing a sequence of task executions. Given a trace s , we denote by $s[i]$ the i -th task in the corresponding sequence, and by $length(s)$ the length of s . The set of all the tasks in s is denoted by $tasks(s) = \bigcup_{1 \leq i \leq length(s)} s[i]$. Finally, a *workflow log* for P , denoted by \mathcal{L}_P , is a bag of traces over A , i.e., $\mathcal{L}_P = [s \mid s \in A^*]$.

We conclude this section by relating the notion of trace with that of instance.

Definition II.3: Let $I = \langle A_I, E_I \rangle$ be an instance of a workflow schema $\mathcal{WS} = \langle A, E, a_0, A_F, Join, Fork \rangle$ for a process P , and let s be a trace in \mathcal{L}_P . Then, s is *compliant with \mathcal{WS} through I* , denoted by $s \models^I \mathcal{WS}$, if s is a topological sort of I . \square

For instance, the trace **abfcgh** is compliant with the instance reported in Figure 1.(c), while the traces **afbcgh** and **abflm** are not.

Notice that the definition above entails that if $s \models^I \mathcal{WS}$, then: $s[1] = a_0$, i.e., the first task in s is the initial activity in \mathcal{WS} ; $s[length(s)] \in A_F$, i.e., the last task in s is a final activity in \mathcal{WS} ; and, $tasks(s) = A_I$, i.e., all the tasks in s are executed in I .

We conclude by saying that s is simply *compliant with \mathcal{WS}* , denoted by $s \models \mathcal{WS}$, if

there exists I such that $s \models^I \mathcal{WS}$. The reader may check, for instance, that the traces in Figure 1.(a) are compliant with the workflow schema in Figure 1.(b), as each of them is a topological sort of some instance of that schema.

Before leaving the section, we note that, while logs in real life are often characterized by multiple executions of the same task, our workflow model admits only one execution for each activity in the schema. However, multiple executions can be coped with by means of a syntactic expedient. Let P be a process and let \mathcal{L}_P be a log for it where multiple executions of the same activity are allowed in the same trace. Let \mathbf{a} be an activity in P , and let $n(\mathbf{a}) \geq 0$ be the maximum number of times it is executed over all the traces in \mathcal{L}_P . Then, we can mine a workflow schema $\overline{\mathcal{WS}}$, where each activity \mathbf{a} is (virtually) substituted by $n(\mathbf{a})$ distinct activities, say $\mathbf{a}_1, \dots, \mathbf{a}_{n(\mathbf{a})}$, by exploiting the log $\overline{\mathcal{L}}_P$ derived from \mathcal{L}_P by replacing the i -th occurrence of \mathbf{a} in any trace by \mathbf{a}_i . By construction, $\overline{\mathcal{WS}}$ contains no duplicated activities and, hence, it conforms Definition II.1.

III. THE PROCESS MODEL DISCOVERY PROBLEM

In this section, we preliminary introduce and discuss the problem of automatically (re)constructing a workflow schema modelling a given process, on the basis of some log data collected from previous executions. Afterwards, a number of variants of the process mining problem are defined and studied from a computational point of view.

A. Disjunctive Schemas

Let A be the set of task identifiers for the process P . We assume that the actual workflow schema for P is unknown, and we consider the problem of properly identifying it, in the set of all the possible workflow schemas having A as set of activities. To this aim, we first derive from the log traces an initial workflow schema, and we subsequently iteratively refine it into a number of specific schemas, each of them modelling a class of traces having the same *behavioral* characteristics. Therefore, the resulting model is a collection of workflow schemas as defined below.

Definition III.1: Let P be a process. A *disjunctive workflow schema* for P , denoted by \mathcal{WS}^\vee , is a set $\{\mathcal{WS}^1, \dots, \mathcal{WS}^m\}$ of workflow schemas for P . The number m is referred to as the *size* of \mathcal{WS}^\vee , denoted by $|\mathcal{WS}^\vee|$. □

Any instance I of some $\mathcal{WS}^j \in \mathcal{WS}^\vee$ is also said to be an instance of \mathcal{WS}^\vee , denoted by $I \models \mathcal{WS}^\vee$. And, in fact, the schema is said *disjunctive* precisely because the set of its instances is the union of the sets of the instances of each workflow schema in \mathcal{WS}^\vee .

Similarly, a trace s is compliant with \mathcal{WS}^\vee , denoted by $s \models \mathcal{WS}^\vee$, if s is compliant with some $\mathcal{WS}^j \in \mathcal{WS}^\vee$. Note that, since a trace s does not provide any information about the edges exploited to execute the registered activities, deciding whether s is compliant with the schema \mathcal{WS}^\vee may be, in principle, a complex task. Indeed, while constructing an instance for s , it is crucial that activating edges are properly chosen to satisfy all the local constraints (in particular, difficulty is with the mixing of exclusive choices with synchronization) and to execute all the activities in the trace. The following proposition, whose proof is in [22], makes it clear that a smart strategy can be conceived to this aim, so that deciding whether a trace s is compliant with a schema \mathcal{WS}^\vee can be done in polynomial time (in the size of the schema).

Proposition III.2: Let $\mathcal{WS}^\vee = \{\mathcal{WS}^1, \dots, \mathcal{WS}^m\}$ be a disjunctive workflow schema, and s be a trace. Let e , n_f , and e_f denote the maximum number of edges, xor-fork nodes and edges originating in xor-fork nodes, respectively, over all schemas in \mathcal{WS}^\vee . Then, deciding whether $s \models \mathcal{WS}^\vee$ is feasible in $O(m \times \max(e, n_f \times e_f))$.

Let us now assume that a log \mathcal{L}_P for the process P is given. Then, we aim at discovering the disjunctive schema \mathcal{WS}^\vee for P which is as “close” as possible to the actual unknown schema that generated \mathcal{L}_P . In the following, the quality of the mined model is evaluated according to two different criteria, namely the *completeness* and the *soundness*, constraining the discovered model to admit exactly the traces in the log. In particular, a (fully) complete workflow is such that all traces in the log at hand are compliant with some instance of it, whereas a (fully) sound workflow is such that all of its possible enactments have been actually registered in the log. These two criteria are formalized below.

Definition III.3: Let \mathcal{WS}^\vee be a disjunctive workflow model, and \mathcal{L}_P be a log for the process P . We define:

- $soundness(\mathcal{WS}^\vee, \mathcal{L}_P) = \frac{|\{s | s \in \mathcal{L}_P \wedge s \models \mathcal{WS}^\vee\}|}{|\{s | s \models \mathcal{WS}^\vee\}|}$, i.e., the percentage of traces compliant with \mathcal{WS}^\vee that have been registered in the log – the larger the sounder;
- $completeness(\mathcal{WS}^\vee, \mathcal{L}_P) = \frac{|\{s | s \in \mathcal{L}_P \wedge s \models \mathcal{WS}^\vee\}|}{|\{s | s \in \mathcal{L}_P\}|}$, i.e., the percentage of traces in the log that are compliant with \mathcal{WS}^\vee – the larger the more complete.




Fig. 2. The two schemas constituting the disjunctive workflow schema \mathcal{WS}^V .

Given two real numbers α and β between 0 and 1 (desirable values should be close to 1), we say that \mathcal{WS}^V is α -sound w.r.t. \mathcal{L}_P , if $\text{soundness}(\mathcal{WS}^V, \mathcal{L}_P) \geq \alpha$; moreover, \mathcal{WS}^V is β -complete w.r.t. \mathcal{L}_P , if $\text{completeness}(\mathcal{WS}^V, \mathcal{L}_P) \geq \beta$. \square

We next apply the notions introduced above to a simple explicative example.

Example III.4: Consider again the workflow schema \mathcal{WS}_0 in Figure 1.(b). We leave to the careful reader the task of checking that \mathcal{WS}_0 admits 20 instances and 276 traces. Let L be the log shown in Figure 1.(a). Thus, we have: $\text{soundness}(\{\mathcal{WS}_0\}, L) = \frac{16}{276} = 5.797\%$, and $\text{completeness}(\{\mathcal{WS}_0\}, L) = \frac{16}{16} = 100\%$.

A sounder representation for the traces in L is given by the disjunctive workflow schema which consists of the two workflow schemas shown in Figure 2. In fact, we have: $\text{soundness}(\{\mathcal{WS}_1^V, \mathcal{WS}_2^V\}, L) = \frac{11}{97} = 11.34\%$, and $\text{completeness}(\{\mathcal{WS}_1^V, \mathcal{WS}_2^V\}, L) = \frac{11}{16} = 68.75\%$. Indeed, $\{\mathcal{WS}_1^V, \mathcal{WS}_2^V\}$ models 97 distinct instances, 64 through \mathcal{WS}_1^V and 33 through \mathcal{WS}_2^V . However, the completeness value is lower than 1 as some traces in L (namely s_8, \dots, s_{12}) are not compliant with $\{\mathcal{WS}_1^V, \mathcal{WS}_2^V\}$. Interestingly, this example shows how the soundness of a given schema can be increased by replacing it with a set of more specific schemas. Such a strategy is exploited in the mining algorithm presented in Section V, which also prevents completeness losses. \triangleleft

Since soundness and completeness are the parameters to be taken into account while mining process models, it is relevant to study their “intrinsic” complexity. As proven next, checking soundness is a hard task, surprisingly even in the case the model is given at hand and does not need to be discovered. The result evidences that our mining problem is structurally harder than traditional learning problems, where the functions to be optimized can be checked in polynomial time for any candidate solution.

Proposition III.5: Let \mathcal{WS}^V be a disjunctive workflow model, and \mathcal{L}_P be a log for a process P . Then, deciding whether \mathcal{WS}^V is 1-complete w.r.t. \mathcal{L}_P is feasible in time

$O(|\{s \mid s \in \mathcal{L}_P\}| \times m \times \max(e, n_f \times e_f))$,¹ with one scan of \mathcal{L}_P only. Moreover, deciding whether \mathcal{WS}^\vee is 1-sound w.r.t. \mathcal{L}_P is co-NP-complete. Hardness holds even for $|\mathcal{WS}^\vee| = 1$.

Proof: To check 1-completeness, we can repeat the procedure in Proposition III.2, for each trace s in \mathcal{L}_P ; therefore, the result easily follows.

Let us now prove that checking for 1-soundness is co-NP-complete. Consider the complementary problem, say $\overline{\text{soundness}}$, of deciding whether there is a trace s such that $s \models \mathcal{WS}^\vee$ and $s \notin \mathcal{L}_P$. Problem $\overline{\text{soundness}}$ is obviously in NP; we next show that it is NP-complete. To this purpose, we recall that in [23] it has been proven that deciding whether a given workflow schema \mathcal{WS} admits an instance is NP-complete. We construct the following instance of the problem $\overline{\text{soundness}}$: $\mathcal{WS}^\vee = \{\mathcal{WS}\}$ and $\mathcal{L}_P = \emptyset$. Then, the answer of $\overline{\text{soundness}}$ is “yes” if and only if \mathcal{WS} admits an instance; hence $\overline{\text{soundness}}$ is NP-complete as well. It turns out that deciding 1-soundness is co-NP-complete. ■

B. Exact and Maximum Process Discovery

Armed with the framework depicted so far, we are in the position of formalizing the problem we want to deal with. Basically, we aim at discovering a disjunctive schema \mathcal{WS}^\vee for a given process P which is α -sound and β -complete, for some given α and β . However, it is easy to see that a trivial schema satisfying the above conditions always exists (even for $\alpha = 1$ and $\beta = 1$), consisting in the union of exactly one workflow modelling each distinct trace in \mathcal{L}_P . Unfortunately, such a model would be an overly detailed and complex representation of the process P , for its size being $|\mathcal{WS}^\vee| = |\mathcal{L}_P| = |\{s \mid s \in \mathcal{L}_P\}|$. We therefore introduce a bound on the number of schemas in \mathcal{WS}^\vee for defining the basic problem we shall study in the paper.

Let \mathcal{L}_P be a workflow log for the process P . Given two real numbers α and β , and a natural number $m > 0$, the *Exact Process Discovery problem*, denoted by $\text{EPD}(\mathcal{L}_P, \alpha, \beta, m)$, consists in finding (if any) an α -sound and β -complete disjunctive workflow schema \mathcal{WS}^\vee , such that $|\mathcal{WS}^\vee| \leq m$. Actually, we can show (see [22]) that the Exact Process Discovery problem can be solved in polynomial time only for trivial cases (unless $\text{P} = \text{NP}$).

Theorem III.6: Given a log \mathcal{L}_P for a process P and a natural number $m > 0$, $\text{EPD}(\mathcal{L}_P, 1, 1, m)$ always admits a trivial solution if $|\mathcal{L}_P| \leq m$. Otherwise, i.e., if $|\mathcal{L}_P| > m$, deciding whether $\text{EPD}(\mathcal{L}_P, 1, 1, m)$ admits a solution is in Σ_2^{P} and NP-hard.

¹See notation in Proposition III.2.

Since, it could happen that EPD does not have solutions (cf. Theorem III.6), and since it is difficult to check whether we are in such a situation (cf. Proposition III.5), we restate the process discovery problem in a way it always admits a solution. Specifically, we allow to sacrifice enough portions of soundness to get a result; however, we impose a strict bound on the completeness (which is usually 1), since modelling all the traces in the log often is an important requirement for the mined schema.

Definition III.7: Let \mathcal{L}_P be a workflow log for the process P . Given a natural number m , the *Maximum Process Discovery problem*, denoted by $\text{MPD}(\mathcal{L}_P, m)$, consists in finding a 1-complete disjunctive workflow schema \mathcal{WS}^\vee , such that $|\mathcal{WS}^\vee| \leq m$ and $\text{soundness}(\mathcal{WS}^\vee, \mathcal{L}_P)$ is the maximum over all the 1-complete schemas. \square

However, also the problem MPD is untractable by a straightforward reduction from EPD. Moreover, we suspect an hardness result for the second level of the polynomial hierarchy. Therefore, it comes no surprise that we pragmatically face the problem by using a heuristic approach, in which we mine a preliminary, possibly not sound enough, model, and iteratively refine it. In particular, the approach is designed in a way that each time a refinement is done, the model is guaranteed to become sounder than before. The careful reader at this point may understand the importance of this property, since it allows to monotonically search in the space of the solutions by never checking for the soundness of the current model, which is unfeasible.

IV. MINING A WORKFLOW SCHEMA

In this section, we start the description of the algorithm for $\text{MPD}(\mathcal{L}_P, m)$ by proposing a solution for the case $m = 1$. Our approach is simple, but it enjoys some nice properties that are crucial for the extension to the case $m > 1$, and which have been not fully considered while designing previous approaches in the literature.

Throughout the section we assume that a log \mathcal{L}_P for the process P over the tasks A is given. For the sake of exposition, we assume that each trace contains the initial activity a_0 as the first task. We start by introducing notions for expressing precedence relationships between activities, which are derived from occurrences in the logs:

- The *dependency graph* for \mathcal{L}_P is the graph $D_{\mathcal{L}_P} = \langle A, E \rangle$, where $E = \{(a, b) \mid \exists s \in \mathcal{L}_P, i \in \{1, \dots, \text{length}(s) - 1\} \text{ s.t. } a = s[i] \wedge b = s[i + 1]\}$;
- Two activities a and b in A are *parallel* in \mathcal{L}_P , if they occur in some cycle of $D_{\mathcal{L}_P}$.




Fig. 3. Example IV.1: (a) Dependency graph; (b) Control flow graph.

- Given two activities a and b in A , we say that a precedes b in \mathcal{L}_P , denoted by $a \rightarrow b$, if a and b are not parallel, and there is a path from a to b in $D\mathcal{L}_P$.

Example IV.1: Consider the log $L = \{abcde, adbce, ae\}$, concerning a process P . The dependency graph for L is shown in Figure 3.(a). Notice, for instance, that a , b and c are parallel activities in L , and that $a \rightarrow c$ and $b \rightarrow e$ hold. \triangleleft

Armed with these notions, we can now describe the algorithm `MineWorkflow` which is shown in Figure 4. The algorithm starts by constructing the dependency graph which is subsequently modified in Steps 2–11 in order to remove the cycles, which in fact correspond to bunches of parallel activities. Actually, removing an edge is a critical step, since it must be carried out by preserving connectivity. Therefore, each time an edge, say (a, b) , is removed from E , we preserve the paths by connecting a and b with some preceding activity (Step 5) and with some following one (Step 8), for each of the traces in \mathcal{L}_P .

Notice that there are several alternative graphs that may support all the log at hand; the algorithm has been designed for introducing as less spurious traces as possible, so that soundness is heuristically maximized. For instance, a schema with a_0 connected to all the other activities will be a complete schema as well; but it is likely to be terribly unsound. Our idea is instead to preserve the paths in the flow (ideally, in absence of cycles we can get even 1-sound schemas). Accordingly, in the Steps 12–18, the algorithm builds the local constrains by applying the most stringent possible choices, for each activity. In particular, notice that or-forks and or-joins are likely to deteriorate the soundness, and should be therefore avoided if possible.

Example IV.2: Let us turn back to Example IV.1, and apply the `MineWorkflow` algorithm. In the Steps 2–10, the edges (b, c) , (c, d) , and (d, b) are removed from the graph in Figure 3.(a), since they are parallel in L . The connectivity of the graph is reestablished by connecting a to c , and b to e ; actually, other edges are processed that were already in the graph. In Step 11, a becomes the starting activity and e the only final one. The

Input: A log \mathcal{L}_P .

Output: A workflow schema $\mathcal{WS} = \langle A, E, a_0, A_F, \text{Join}, \text{Fork} \rangle$.

Method: Perform the following steps:

```

1   $\langle A, E \rangle := D_{\mathcal{L}_P}$ ;           //nodes and edges are initially those of the dependency graph
2  for each  $(a, b) \in E$  s.t.  $a$  and  $b$  are parallel in  $\mathcal{L}_P$  do           //remove cycles
3     $E := E - \{(a, b)\}$ ;
4  for each  $s \in \mathcal{L}_P$  s.t.  $\{a, b\} \subseteq \text{tasks}(s)$  do           //update edges
5     $pre := s[i]$ , where  $s[i] \rightarrow a \wedge s[i] \rightarrow b$  and not exists  $s[k]$  with  $k > i$  s.t.  $s[k] \rightarrow a \wedge s[k] \rightarrow b$ ;
6     $E := E \cup \{(pre, a)\} \cup \{(pre, b)\}$ ;
7     $post := s[j]$ , where  $a \rightarrow s[j] \wedge b \rightarrow s[j]$  and not exists  $s[h]$  with  $h < j$  s.t.  $a \rightarrow s[h] \wedge b \rightarrow s[h]$ ;
8     $E := E \cup \{(a, post)\} \cup \{(b, post)\}$ ;
9  end for
10 end for
11  $a_0 := s[1]$ , no matter of which trace  $s \in \mathcal{L}_P$  is selected;    $A_F := \{a \in A \mid \nexists b \in A \text{ s.t. } a \rightarrow b\}$ ;
12 for each  $a \in A$  do           //construction of local constraints
13   if  $\forall s \in \mathcal{L}_P$  s.t.  $a \in \text{tasks}(s)$ , it holds that  $\forall c$  s.t.  $(a, c) \in E, c \in \text{tasks}(s)$  then  $\text{Fork}(a) = \text{AND}$ ;
14   else if  $\forall s \in \mathcal{L}_P$  s.t.  $a \in \text{tasks}(s), |\{c \mid (a, c) \in E \wedge c \in \text{tasks}(s)\}| = 1$  then  $\text{Fork}(a) = \text{XOR}$ ;
15   else  $\text{Fork}(a) = \text{OR}$ ;
16   if  $\forall s \in \mathcal{L}_P$  s.t.  $a \in \text{tasks}(s), (c, a) \in E \Rightarrow c \in \text{tasks}(s)$  then  $\text{Join}(a) = \text{AND}$ ;
17   else  $\text{Join}(a) = \text{OR}$ ;
18 end for
19 return  $\langle A, E, a_0, A_F, \text{Join}, \text{Fork} \rangle$ ;

```

Fig. 4. **Algorithm MineWorkflow:** A solution for $\text{MPD}(\mathcal{L}_P, 1)$.

resulting control flow is shown in Figure 3.(b). Finally, the associated constraints are such that: *Fork* and *Join* both assign the value AND to each of the activities **b**, **c**, and **d**, for they having just one predecessor and one successor. Moreover, the initial activity **a** (resp. the final activity **e**) is associated with an OR value by the *Fork* (resp. *Join*) function, since the trace **ae** in the log contains only one of the successors of **a** (resp. predecessors of **e**). \triangleleft

We can now conclude the description of the algorithm, by stating a number of relevant properties. First, it is not difficult to see that the mined schema is 1-complete; more interestingly, we are able to show (see [22]) that it satisfies a kind of “monotonicity” property that is crucial for guaranteeing the correctness of the greedy strategy we shall exploit in the following section for the general problem $\text{MPD}(\mathcal{L}_P, m)$, with $m > 1$.

Theorem IV.3: *The MineWorkflow algorithm on input \mathcal{L}_P computes a workflow schema \mathcal{WS} in linear time in the size of \mathcal{L}_P and satisfies the following conditions:*

Maximum Completeness Condition: \mathcal{WS} is a 1-complete workflow schema w.r.t. \mathcal{L}_P

Monotonicity Condition: Let \mathcal{L}'_P be a log such that $\mathcal{L}'_P \subseteq \mathcal{L}_P$, and let \mathcal{WS}' be the output of `MineWorkflow` on input \mathcal{L}'_P . Then, the set of traces compliant with \mathcal{WS}' is a subset of the set of traces compliant with \mathcal{WS} , i.e., $\{s \mid s \models \mathcal{WS}'\} \subseteq \{s \mid s \models \mathcal{WS}\}$.

As a further remark, note that the proposed algorithm can be extended to cope with noise by slightly modifying the construction of the dependency graph. Specifically, we can use a threshold ϱ , such that an edge (a, b) is in a suitable dependency graph, say $D_{\mathcal{L}_P}^\varrho$, if and only if a occurs before b in at least $\varrho \times |\mathcal{L}_P|$ traces. Clearly enough, this construction avoids the introduction of spurious dependencies that are unfrequent in \mathcal{L}_P .

V. CLUSTERING WORKFLOW TRACES

In order to solve the $\text{MPD}(\mathcal{L}_P, m)$ problem, we exploit the idea of iteratively and incrementally refining a schema, starting with a preliminary non-disjunctive model which can be mined by the algorithm `MineWorkflow`. In a nutshell, we propose a greedy solution, implemented in the algorithm `ProcessDiscover`, that computes the mined schema \mathcal{WS}^\vee through a hierarchical clustering in which the current disjunctive schema (equipped with a subset of the input log) is stepwise refined into k clusters of traces modelled by possibly different schemas, with the aim of increasing its soundness yet preserving maximum completeness. In this section, we discuss the details of such an algorithm.

A. Algorithm `ProcessDiscovery`

The algorithm `ProcessDiscover` is shown in Figure 5. It first mines, in Step 1, a workflow schema \mathcal{WS}_0^1 , by means of the procedure `MineWorkflow` described in Section IV. Such a workflow is associated with all the traces in \mathcal{L}_P and becomes the starting point of the computation; indeed, the disjunctive workflow schema \mathcal{WS}^\vee initially contains \mathcal{WS}_0^1 only (see Step 2). Then, the algorithm starts refining the schema. In particular, each workflow schema, say \mathcal{WS}_i^j , eventually inserted in \mathcal{WS}^\vee , is identified by the number i of refinements occurred since the beginning of the computation and by an index j for distinguishing the schemas at the same refinement level. Each schema \mathcal{WS}_i^j is also equipped with a set of traces that \mathcal{WS}_i^j is able to model, denoted by $\mathcal{L}(\mathcal{WS}_i^j)$. Such sets of traces can be viewed as clusters of the original log.

At each step, `ProcessDiscover` selects a schema $\mathcal{WS}_i^j \in \mathcal{WS}^\vee$ for being refined (Step 4) by exploiting the function `refineWorkflow`. To this aim, the most natural strategy is

Input: Problem $\text{MPD}(\mathcal{L}_P, m)$, natural numbers maxFeatures , k and ℓ , thresholds σ and γ .

Output: A process model.

Method: Perform the following steps:

```

1   $\mathcal{WS}_0^1 := \text{MineWorkflow}(\mathcal{L}_P);$  //See Section IV
2   $\mathcal{WS}^\vee := \{\mathcal{WS}_0^1\};$   $\mathcal{L}(\mathcal{WS}_0^1) := \mathcal{L}_P;$   $\text{noRefinement} = \text{false};$ 
3  while  $|\mathcal{WS}^\vee| < m$  or  $\text{noRefinement} = \text{true}$  do
4     $\mathcal{WS}_i^j := \text{selectSchema}(\mathcal{WS}^\vee);$ 
5     $\text{oldSize} := |\mathcal{WS}^\vee|;$ 
6     $\text{refineWorkflow}(i, j, k, \sigma, \gamma, \ell, \text{maxFeatures})$  //See Theorem V.1
7    if  $\text{oldSize} = |\mathcal{WS}^\vee|$  then  $\text{noRefinement} = \text{true};$ 
8  end while
9  return  $\mathcal{WS}^\vee;$  //See Theorem V.1

```

Procedure $\text{refineWorkflow}(i: \text{refinement}, j: \text{schema}, k: \text{branching factor}, \sigma, \gamma: \text{thresholds}, \ell, \text{maxFeatures}: \text{natural numbers})$

```

P1   $\mathcal{F} := \text{IdentifyRelevantFeatures}(\mathcal{L}(\mathcal{WS}_i^j), \mathcal{WS}_i^j, \sigma, \gamma, \ell, \text{maxFeatures});$  //See Section V-B
P2  if  $|\mathcal{F}| > 1$  then
P3     $\mathcal{R}(\mathcal{WS}_i^j) := \text{Project}(\mathcal{L}(\mathcal{WS}_i^j), \mathcal{F});$  //See Section V-B
P4     $h := \max\{h \mid \mathcal{WS}_{i+1}^h \in \mathcal{WS}^\vee\};$ 
P5     $\langle \mathcal{L}(\mathcal{WS}_{i+1}^{h+1}), \dots, \mathcal{L}(\mathcal{WS}_{i+1}^{h+k}) \rangle := k\text{-means}(\mathcal{R}(\mathcal{WS}_i^j));$ 
P6     $\mathcal{WS}^\vee := \mathcal{WS}^\vee - \{\mathcal{WS}_i^j\};$ 
P7    for each  $\mathcal{WS}_{i+1}^h$  do
P8       $\mathcal{WS}_{i+1}^h := \text{MineWorkflow}(\mathcal{L}(\mathcal{WS}_{i+1}^h));$  //See Section IV
P9       $\mathcal{WS}^\vee = \mathcal{WS}^\vee \cup \{\mathcal{WS}_{i+1}^h\};$ 
P10   end for
P11  end if;

```

Fig. 5. **Algorithm** ProcessDiscover

to select the schema having the minimum value of soundness over all the schemas in \mathcal{WS}^\vee . In order to get an efficient implementation, we pragmatically suggest to exploit an approximation of such an approach, where the schema having the maximum number of or-forks is chosen — these nodes introduces, in fact, nondeterminism and possibly spurious traces. Notably, in our current implementation (see Section VI), the user is allowed to interactively select at each step the cluster to refine, so that any arbitrary, application-dependent strategy may be adopted.

The refinement is carried out by ‘partitioning’ the traces associated with \mathcal{WS}_i^j in a way that guarantees the resulting schema to become sounder than before. Actually, in order to reuse classical clustering methods, and specifically in our implementation the k -means algorithm, the procedure refineWorkflow translates the log $\mathcal{L}(\mathcal{WS}_i^j)$ into flat relational data, denoted by $\mathcal{R}(\mathcal{WS}_i^j)$, by means of the procedures $\text{IdentifyRelevantFeatures}$ and Project , which will be discussed in Section V-B. The basic idea is to identify a set of relevant features that are assumed to characterize the traces in the cluster, thereby leading

to view each trace as a boolean tuple over the space of such features. In particular, if more than one feature is identified, it computes the clusters $\mathcal{WS}_{i+1}^{h+1}, \dots, \mathcal{WS}_{i+1}^{h+k}$, where h is the maximum index of the schemas already inserted in \mathcal{WS}^\vee at the level $i+1$, by applying the *k-means* algorithm on the traces in $\mathcal{L}(\mathcal{WS}_i^j)$, and add them to the disjunctive schema \mathcal{WS}^\vee .

Finally, for each schema added to \mathcal{WS}^\vee , the procedure `MineWorkflow` described in Section IV is applied, so that the control flow and the local constraints are mined as well.

Note that a main point of the algorithm is fixing the number k of new schemas to be added at each refinement step. The range of k goes from a minimum of 2, which will require several steps for the computation, to an unbounded value, which will return the result in only one step. One could then expect that the latter case is the most desirable. This is not necessarily true; rather there are three basic reasons for preferring a hierarchical computation, with several refinements:

- 1) Exploiting a large value for k would be beneficial only if it were possible to encode traces in a metric space by guaranteeing that standard clustering algorithms produce a solution maximizing the soundness; recall, indeed, that our basic aim is to identify bunches of execution traces which can be soundly modelled by a workflow schema. However, given that the notion of soundness is computationally intractable (cf. Proposition III.5), and that the MPD problem is likely to be hard for the second level of the polynomial hierarchy (cf. Theorem III.6), we believe that this cannot be done efficiently. To see this from another perspective, we can say that, differently from traditional clustering problems where the compactness of the clusters is measured according to the same metric used for evaluating the similarity between pair of entities (e.g., euclidean distance in a metric space), in the MPD problem the objective function (soundness) cannot be related in a straightforward manner with some kind of likeness among traces. An iterative approach based on stepwise refinements guaranteeing that each refinement leads to a sounder schema is quite an effective solution to this problem. Moreover, it allows to monotonically search in the space of the solutions by never checking for the soundness of the current model.
- 2) The result of the hierarchical clustering is a taxonomy of workflow schemas whose leaves encode, in fact, the mined disjunctive model (recall that each workflow is in fact equipped with a level i and is the result of a refinement of some workflow with level $i-1$). Such a tree-based representation is relevant because it gives more insights on the properties of

the modelled workflow instances and provides an intuitive and expressive description of the process behavior at different levels of detail. The exploitation of this structure for semantic knowledge consolidation tasks has been recently discussed in [24].

3) Disjunctive models can be used not only to get an effective comprehension of complex processes, but also as executable models supporting further coming instances of the process. Indeed, as soon as there is a new enactment, the most appropriate mined variant of the process should be selected based on the environment (e.g., users and data values). Clearly enough, shifting the choice of the variant to the very first step of the enactment may be undesirable in several situations (see, e.g., studies on branching bisimulation [25], [26]). And, in fact, an interesting strategy to support the enactment is to exploit the hierarchical structure of the mined model. At the very beginning, the most general model (i.e., the root of the hierarchy) may be selected for being enacted. Then, as soon as a choice is done which allows to univocally determine the specific variant (of the current schema) being actually executed, the hierarchy can be traversed, and one child of the current node can be selected to be the actual workflow schema. The technique can be iterated till a leaf of the hierarchy is selected. Therefore, the more the schema has a tree-like structure, the more we are free to decide the actual moment in which a specific schema has to be associated with the current enactment.

We can now conclude the description of the algorithm `ProcessDiscover` by stating its main properties. The careful reader will notice that the properties of `MineWorkflow` are in fact crucial now.

Theorem V.1: Let \mathcal{WS}^\vee be the output of `ProcessDiscover` applied on input $\text{MPD}(\mathcal{L}_P, m)$. Then, (1) \mathcal{WS}^\vee is 1-complete w.r.t. \mathcal{L}_P . (2) Let \mathcal{WS}_b^\vee be an α -sound disjunctive schema, and let \mathcal{WS}_a^\vee be the α' -sound disjunctive schema obtained updating \mathcal{WS}_b^\vee by means of the invocation of `refineWorkflow` in Step 6. Then, $\alpha' \geq \alpha$. (3) The main loop (Step 3–8) is repeated m times at most.

Proof: (1) Assume that the output \mathcal{WS}^\vee of `ProcessDiscover` on input $\text{MPD}(\mathcal{L}_P, m)$ is of the form $\{\mathcal{WS}_1, \dots, \mathcal{WS}_n\}$. Recall, preliminary, that each schema \mathcal{WS}_i is also equipped with a subset of \mathcal{L}_P , denoted by $\mathcal{L}(\mathcal{WS}_i)$. It is easy to see that the logs $\mathcal{L}(\mathcal{WS}_1), \dots, \mathcal{L}(\mathcal{WS}_n)$ form, in fact, a partition of \mathcal{L}_P . Indeed, the traces associated with \mathcal{WS}_0^1 coincide with those in \mathcal{L}_P (see, Step 2); moreover, each time a schema is refined, its associated traces are

simply clustered in step P5, so that the property of being a partition of \mathcal{L}_P is preserved after any invocation of the *refineWorkflow* procedure.

Finally, the result follows because of the 1-completeness (by the property of *MineWorkflow* in Step 1 and Step P8) of each schema $\mathcal{WS}_i \in \{\mathcal{WS}_1, \dots, \mathcal{WS}_n\}$.

(2) Let \mathcal{WS}_b^\vee be an α -sound disjunctive schema, and let \mathcal{WS}_a^\vee be the α' -sound disjunctive schema obtained updating \mathcal{WS}_b^\vee by means of the invocation of *refineWorkflow*. Let us preliminary recall that: $\alpha = \frac{|\{s | s \in \mathcal{L}_P \wedge s \models \mathcal{WS}_b^\vee\}|}{|\{s | s \models \mathcal{WS}_b^\vee\}|}$; $\alpha' = \frac{|\{s | s \in \mathcal{L}_P \wedge s \models \mathcal{WS}_a^\vee\}|}{|\{s | s \models \mathcal{WS}_a^\vee\}|}$.

Then, by property 1 above, both \mathcal{WS}_b^\vee and \mathcal{WS}_a^\vee are 1-complete. Therefore, $\frac{|\{s | s \in \mathcal{L}_P \wedge s \models \mathcal{WS}_b^\vee\}|}{|\{s | s \in \mathcal{L}_P\}|} = \frac{|\{s | s \in \mathcal{L}_P \wedge s \models \mathcal{WS}_a^\vee\}|}{|\{s | s \in \mathcal{L}_P\}|} = 1$. Now, it is the case that $\frac{\alpha'}{\alpha} = \frac{|\{s | s \models \mathcal{WS}_b^\vee\}|}{|\{s | s \models \mathcal{WS}_a^\vee\}|}$. Assume now, w.l.o.g., that *refineWorkflow* finds at least two features, otherwise $\mathcal{WS}_b^\vee = \mathcal{WS}_a^\vee$ holds. Then, let $\mathcal{WS}_b^\vee = \{\mathcal{WS}_1, \dots, \mathcal{WS}_n\}$, and let \mathcal{WS}_n be the schema refined by *refineWorkflow*, so that $\mathcal{WS}_a^\vee = \{\mathcal{WS}_1, \dots, \mathcal{WS}_{n-1}, \mathcal{WS}_{n+1}, \dots, \mathcal{WS}_{n+k}\}$ is the disjunctive schema after that the refinement of \mathcal{WS}_n is performed — notice that we are assuming, w.l.o.g., that \mathcal{WS}_n is the schema chosen for being refined. Now, observe that $\{s \mid s \models \mathcal{WS}_b^\vee\} = \bigcup_{i=1..n} \{s \mid s \models \mathcal{WS}_i\}$, and $\{s \mid s \models \mathcal{WS}_a^\vee\} = \bigcup_{i=1..n-1, n+1..n+k} \{s \mid s \models \mathcal{WS}_i\}$. Recall that each schema of the form \mathcal{WS}_i , with $n+1 \leq i \leq n+k$, is obtained by mining $\mathcal{L}(\mathcal{WS}_i)$ by means of *MineWorkflow*, and that $\{\mathcal{L}(\mathcal{WS}_{n+1}), \dots, \mathcal{L}(\mathcal{WS}_{n+k})\}$ is a partition of $\mathcal{L}(\mathcal{WS}_n)$, by result of clustering in Step P5. It follows that $\mathcal{L}(\mathcal{WS}_i) \subseteq \mathcal{L}(\mathcal{WS}_n)$, for each index i with $n+1 \leq i \leq n+k$. Then, since *MineWorkflow* satisfies the monotonicity condition (cf. Theorem IV.3), the set $\{s \mid s \models \mathcal{WS}_i\}$ is contained in the set $\{s \mid s \models \mathcal{WS}_n\}$. Clearly, this entails that $\{s \mid s \models \mathcal{WS}_a^\vee\} \subseteq \{s \mid s \models \mathcal{WS}_b^\vee\}$, and, consequently, $|\{s \mid s \models \mathcal{WS}_a^\vee\}| \leq |\{s \mid s \models \mathcal{WS}_b^\vee\}|$. Thus, $\frac{\alpha'}{\alpha} \geq 1$ holds as well.

(3) Termination in m iterations at most is guaranteed by the check in Step 3. ■

B. Dealing with Relevant Features

The last aspect to be analyzed in the algorithm *ProcessDiscovery* is the way the procedures *IdentifyRelevantFeatures* and *Project* are carried out. Roughly speaking, the former identifies a set \mathcal{F} of relevant features, whereas the latter projects the traces into a vectorial space whose components correspond to such features. This is done with the aim of reducing the problem of clustering workflow traces to a standard clustering problem, for which efficient algorithms have been already proposed in literature.

Actually, the idea of representing the data set at hand by using a proper set of features

has been exploited for efficiently handling data mining problems (see, e.g., [27], [28], [29], [30], [31]); as an instance, in [30], the problem of classifying sequences is dealt with by considering frequent subsequences of them as relevant features.

Clearly enough, changing the domain of interest dramatically affects the notion of relevant feature, which is strongly application dependent. In particular, in the case of workflow executions, the identification of relevant features is aimed at having clusters which can be modelled by means of sound schemas. Generally, a schema with a low value of soundness is such that it mixes different execution scenarios that cannot be kept separate by means of dependencies in the control flow and local constraints, only. Thus, a simple way for increasing the soundness is to precisely single out those kinds of (frequent) ‘behavior’ that are not properly modelled by the workflow schema. This is explained below.

Example V.2: Consider again the sample workflow schema and the associated log reported in Figure 1. Let us try to identify some relevant features to be used for the clustering. Consider, for instance, the sequences **abefi** and **im**, and notice that they frequently occur in the log (five times each). However, their combination, i.e., **abefim**, never occurs in the log, which is unexpected by looking at the control flow only. Similar considerations apply also when noticing the absence in the log of the sequence **acdgi**, which is strange due to frequent occurrence of both the sequential patterns **acdgi** and **il**.

Intuitively, the situations above witness some kinds of behavioral constraint. For instance, it may be the case that a fidelity discount is not applied for new clients and that the mail department can be contacted only when it was not necessary to check the availability of external suppliers. As a result of these behaviors, since both **abefim** and **acdgi** are expected but not in the log, the workflow schema is likely to be unsound. The notion of relevant feature reported below is aimed at capturing such scenarios. \triangleleft

Let \mathcal{L}_P be a log, $\mathcal{WS} = \langle A, E, a_0, A_F, Fork, Join \rangle$ be a workflow schema and σ be a threshold, i.e., a real number such that $0 \leq \sigma \leq 1$. Then, we say that a sequence $\mathbf{a}_1\mathbf{a}_2\dots\mathbf{a}_h$ of activities (i.e., a string in A^*) is σ -frequent in \mathcal{L}_P if for each pair of consecutive activities $(\mathbf{a}_i, \mathbf{a}_{i+1})$, with $1 \leq i \leq h - 1$, there exists a path from \mathbf{a}_i to \mathbf{a}_{i+1} in $\langle A, E \rangle$, and if $|\{s \in \mathcal{L}_P \mid \mathbf{a}_1 = s[i_1], \dots, \mathbf{a}_h = s[i_h] \wedge i_1 < \dots < i_h\}| / |\mathcal{L}_P| > \sigma$.

Definition V.3: A *discriminant rule* (or *feature*) with threshold $\langle \sigma, \gamma \rangle$ in a log \mathcal{L}_P is an expression ϕ of the form $[\mathbf{a}_1\dots\mathbf{a}_h] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{a}$ such that: (i) $\mathbf{a}_1\dots\mathbf{a}_h$ is σ -frequent in \mathcal{L}_P , (ii)

$\mathbf{a}_h \mathbf{a}$ is σ -frequent in \mathcal{L}_P , and (iii) $\mathbf{a}_1 \dots \mathbf{a}_h \mathbf{a}$ is not γ -frequent in \mathcal{L}_P . \square

Notice that in the above definition, we considered a second threshold (γ) for evaluating whether the resulting string is frequent: the lower γ , the more unexpected the rule. In the extreme case, i.e., $\gamma = 0$, the workflow mixes two scenarios that are completely independent of each other. This is the case, for instance, of the schema and the log in Figure 1, where both $[\mathbf{abefi}] \not\rightarrow_{\langle \sigma, 0 \rangle} \mathbf{m}$ and $[\mathbf{acdgi}] \not\rightarrow_{\langle \sigma, 0 \rangle} \mathbf{l}$ are discriminant rules for $\sigma < 5/16$.

Actually, while discovering discriminant rules we are interested in those satisfying some additional minimality requirements.

Definition V.4: A feature $\phi : [\mathbf{a}_1 \dots \mathbf{a}_h] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{a}$ is *minimal* if the following conditions (aimed at avoiding redundancies) are satisfied:

(iv) does not exist \mathbf{b} with $[\mathbf{a}_1 \dots \mathbf{a}_h] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{b}$, such that \mathbf{ab} is σ -frequent in \mathcal{L}_P , and

(v) does not exist $[\mathbf{c}_1 \dots \mathbf{c}_k] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{a}$, such that $tasks(\mathbf{c}_1 \dots \mathbf{c}_k) \subseteq tasks(\mathbf{a}_1 \dots \mathbf{a}_h)$. \square

Once that minimal features have been discovered, the procedure **Project** maps each trace s in the log \mathcal{L}_P into a point of a suitable vectorial space where the k -means algorithm can operate. The vectorial space has as many dimensions as the number of features considered, and each trace s is mapped to the point \vec{s} as follows. Let $\phi : [\mathbf{a}_1 \dots \mathbf{a}_h] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{a}$ be a feature computed by **IdentifyRelevantFeatures**, then the value of the component of \vec{s} associated with ϕ is: 0, if $\mathbf{a} \in tasks(s)$, or $\frac{\sum_i |s[i] \cap \{\mathbf{a}_1, \dots, \mathbf{a}_h\}| \times length(s)^{length(s)-i}}{\sum_i length(s)^{length(s)-i}}$, otherwise.

Intuitively, **Project** tries to map the traces by splitting them according to the occurrence of the feature ϕ . Indeed, the lowest value (cf. 0) is assigned in the case where \mathbf{a} occurs in the trace s , while the highest value (cf. 1) is assigned if all the nodes in $\{\mathbf{a}_1, \dots, \mathbf{a}_h\}$ are in s , but \mathbf{a} is not. Otherwise, i.e., if $\mathbf{a} \notin tasks(s)$ but some node in $\{\mathbf{a}_1, \dots, \mathbf{a}_h\}$ is not in s , the value is obtained by lexicographically weighting the occurrences of such nodes, by giving preference to those occurring first in the control flow.

Example V.5: Recall that $[\mathbf{abefi}] \not\rightarrow_{\langle \sigma, 0 \rangle} \mathbf{m}$ and $\phi_2 : [\mathbf{acdgi}] \not\rightarrow_{\langle \sigma, 0 \rangle} \mathbf{l}$, for $\sigma < 5/16$, are discriminant rules in the example in Figure 1. Then, it is not difficult to check that $\phi_1 : [\mathbf{efi}] \not\rightarrow_{\langle \sigma, 0 \rangle} \mathbf{m}$ and $\phi_2 : [\mathbf{dgi}] \not\rightarrow_{\langle \sigma, 0 \rangle} \mathbf{l}$ are minimal discriminant rules, which allow to project the log into a 2-dimensional vectorial space, as shown in Figure 6.(a).

By looking at this feature space in Figure 6.(b), four clusters can be identified: $C_1 = \{s_1, s_2, s_3, s_4, s_5, s_6, s_7\}$, $C_2 = \{s_8, s_9, s_{10}\}$, $C_3 = \{s_{11}, s_{12}\}$, and $C_4 = \{s_{13}, s_{14}, s_{15}, s_{16}\}$. \triangleleft

Let us now focus on the algorithm **IdentifyRelevantFeatures** for computing the set

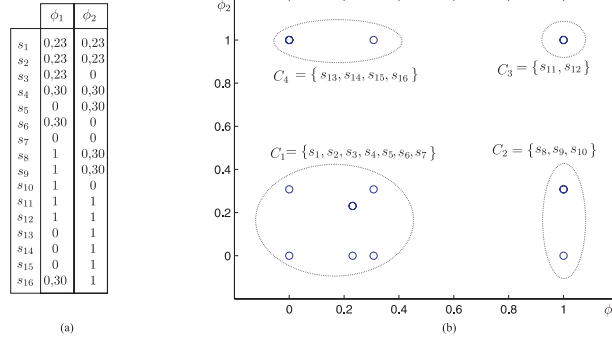


Fig. 6. Sample traces: (a) projection in the feature space; (b) clusters.

of all the minimal features. The algorithm takes in input a log \mathcal{L}_P , a workflow schema \mathcal{WS} , the thresholds σ and γ in Definition V.3, a number ℓ bounding the length of the features to be discovered, and a number $maxFeatures$ bounding the number of features that should be returned as output. The algorithm exploits a level-wise search in the space of all the possible features, in the *a-priori* style (see, e.g., [32], [33], [34]).

To see how this exploration is possible, consider two strings $s : \mathbf{a}_1 \dots \mathbf{a}_h$ and $s' : \mathbf{a}'_1 \dots \mathbf{a}'_{h'}$. We say that s *directly precedes* s' denoted by $s \prec s'$ if s is a prefix of s' and $|tasks(s)| = |tasks(s')| - 1$. Moreover, we say that s *precedes* s' , denoted by $s \prec^* s'$, if either $s \prec s'$ or there exists a string s'' such that $s \prec^* s''$ and $s'' \prec^* s'$. It is not difficult to see that all frequent sequences can be constructed by means of a chain over the \prec relation and, therefore, the space of frequent sequences forms a lower semi-lattice that can be explored in a bottom-up fashion: At each iteration of the main loop (Steps 3–15) the algorithm generates all the possible σ -frequent sequences whose length is len , by exploiting previously computed σ -frequent sequences with length $len - 1$, and stores them in $Cand_{len}$. In particular, in Steps 8 and 9, it scans the log for singling out the sequences in $Cand_{len}$ that are σ -frequent and γ -frequent in \mathcal{L}_P , by storing them in the sets L_{len}^σ and L_{len}^γ , respectively. Actually, the computation starts in Step 1, where L_2^σ is initialized to contain all the σ -frequent sequences of length 2. Then, candidates of length len are obtained by combining any sequence of the form $\mathbf{a}_1 \dots \mathbf{a}_j$ in L_{len-1}^σ with any sequence of the form $\mathbf{a}_j \mathbf{a}$ in L_2^σ .

Finally, after that frequent sequences are discovered, in Steps 10–13, the features consisting of len nodes are identified and stored in F_{len} . To this aim, the sequence of the form $\mathbf{a}_1 \dots \mathbf{a}_j \mathbf{a}$ must be not γ -frequent and the minimality conditions must be satisfied (Step 11). At the end of each iteration, the discovered minimal features are eventually inserted into the set \mathcal{F} (Step 14). The process is repeated until no other frequent sequences are found

Input: A log \mathcal{L}_P , a schema $\mathcal{WS} = \langle A, E, a_0, A_F, Fork, Join \rangle$, thresholds σ and γ , natural number ℓ and $maxFeatures$.

Output: A set of minimal discriminant rules.

Method: Perform the following steps:

```

1   $L_2^\sigma := \{\mathbf{ab} \mid \mathbf{ab} \text{ is } \sigma\text{-frequent in } \mathcal{L}_P\};$ 
2   $len := 3; \mathcal{F} := \emptyset;$ 
3  while  $len \leq \ell$  and  $L_{len}^\sigma \neq \emptyset$  do           //iterations on the length of the features
4     $Cand_{len} := \emptyset; F_{len} := \emptyset;$ 
5    for each sequence  $\mathbf{a}_1 \dots \mathbf{a}_j \in L_{len-1}^\sigma$  do           //construction of the candidates
6      for each  $\mathbf{a}_j \mathbf{a} \in L_2^\sigma$  do
7         $Cand_{len} := Cand_{len} \cup \{\mathbf{a}_1 \dots \mathbf{a}_j \mathbf{a}\};$ 
8     $L_{len}^\sigma := \{s \mid s \in Cand_{len} \wedge s \text{ is } \sigma\text{-frequent in } \mathcal{L}_P\};$ 
9     $L_{len}^\gamma := \{s \mid s \in Cand_{len} \wedge s \text{ is } \gamma\text{-frequent in } \mathcal{L}_P\};$ 
10   for each sequence  $\mathbf{a}_1 \dots \mathbf{a}_j \mathbf{a} \in (Cand_{len} - L_{len}^\gamma)$  do           //update features
11     if  $\exists \mathbf{a}_1 \dots \mathbf{a}_j \mathbf{b} \in (Cand_{len} - L_{len}^\gamma)$  such that  $\mathbf{ab} \in L_2^\sigma$  and
12        $\exists [c_1 \dots c_k] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{a}$  in  $\mathcal{F}$ , such that  $tasks(c_1 \dots c_k) \subseteq tasks(\mathbf{a}_1 \dots \mathbf{a}_j)$  then
13          $F_{len} := F_{len} \cup \{[\mathbf{a}_1 \dots \mathbf{a}_j] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{a}\};$ 
14     end for
15    $\mathcal{F} := \mathcal{F} \cup F_{len}; \quad len := len + 1;$ 
16 end while
17 return  $mostDiscriminantFeatures(\mathcal{F}, maxFeatures);$ 

```

Fig. 7. **Algorithm IdentifyRelevantFeatures**

or all the features up to the length ℓ are found.

It is worth noting that the algorithm does not directly output the set \mathcal{F} ; rather, it invokes the function *mostRelevantFeatures*, whose aim is to select (if possible) the *maxFeatures* most representative in \mathcal{F} . This latter task is carried out for reducing the dimensionality of the vectorial space. Usually, feature selection problem is a very complex activity, and some general-purpose techniques have been proposed in the literature (see, e.g., [31], [35]). In this case things are simpler; indeed, the specific semantics of discriminant rules induces quite a natural ordering among elements in \mathcal{F} , consisting in preferring features having the lowest value for the threshold γ . These are, in fact, the most unexpected rules. Therefore, *mostRelevantFeatures* simply returns the top-*maxFeatures* elements in \mathcal{F} w.r.t. this ordering.

After that the algorithm has been described, we can now state its main properties.

Theorem V.6: *The algorithm IdentifyRelevantFeatures is such that: (1) The main loop is repeated $\ell - 2$ times at most, and $\ell - 2$ is in fact the maximum number of log scans. (2) At each iteration, L_{len}^σ contains the set of all the σ -frequent sequences of length len . (3) At the end of the computation, the set \mathcal{F} contains the set of all the minimal features of length bounded by ℓ .*

Proof: (1) By conditions in step 3, there are at most $\ell - 2$ iterations. Moreover, for each iteration, only Steps 8 and 9 requires to access to the log \mathcal{L}_P , which can be done by one scan only.

(2) L_{len}^σ trivially contains σ -frequent sequences only, because of Step 8. Moreover, all the sequences in L_{len}^σ have length len ; indeed, sequences in L_2^σ have, by construction, length 2, and each sequence in L_{len}^σ is the result of the merging of a sequence $\mathbf{a}_1 \dots \mathbf{a}_j$ (of length $len - 1$) and $\mathbf{a}_j \mathbf{a}$ — formally, one can prove such property by induction. Thus, we have to show that L_{len}^σ , in fact, contains all the σ -frequent sequences of length len . The proof is by structural induction on the length len . **Base:** for $len = 2$, the property holds by construction. **Induction:** Assume that L_{len-1}^σ contains all the σ -frequent sequences of length $len - 1$. Clearly, sequences in L_{len}^σ can be written by adding to any sequence in L_{len-1}^σ exactly one activity, because of the fact that the space of such sequences forms a lower semi-lattice w.r.t. the ‘ \prec ’ relation. This is precisely what is done in Step 7.

(3) We show that F_{len} contains exactly the minimal features of size len , for $len \geq 3$. By construction in Steps 10–12, F_{len} contains only minimal features of size len . Therefore, we have to show that every feature ϕ of size len of the form $[\mathbf{a}_1 \dots \mathbf{a}_j] \not\rightarrow_{\langle \sigma, \gamma \rangle} \mathbf{a}$ is in F_{len} . To this aim, it is sufficient to exploit the completeness result in point (2) above, and notice that $\mathbf{a}_1 \dots \mathbf{a}_j$ is in L_{len-1}^σ and $\mathbf{a}_j \mathbf{a}$ is in L_2^σ . ■

Now that all the procedures in the `ProcessDiscovery` algorithm have been discussed, we can explicitly note that it requires a number of scans on the log which linearly depends on the parameters ℓ , k and m . This is a nice property for scaling in real applications. To see why this is the case, recall that, by Theorem V.1, the algorithm makes m iterations at most. At each iteration, the dominant operation is the procedure `refineWorkflow` in which there is one invocation of `IdentifyRelevantFeatures` (for computing features of length ℓ) for each of the k different invocations of the algorithm `MineWorkflow`. Then, by combining results in Theorem V.6 and Theorem IV.3, the following is obtained.

Corollary V.7: Algorithm `ProcessDiscovery` requires $O(m \times k \times \ell)$ scans of the log \mathcal{L}_P .

VI. EXPERIMENTAL RESULTS

All the algorithms proposed in the paper have been implemented and, recently, integrated in the *ProM* process mining framework [36] as an *analysis* plug-in,² where the user

²Available at <http://www.icar.cnr.it/workflowmining>.

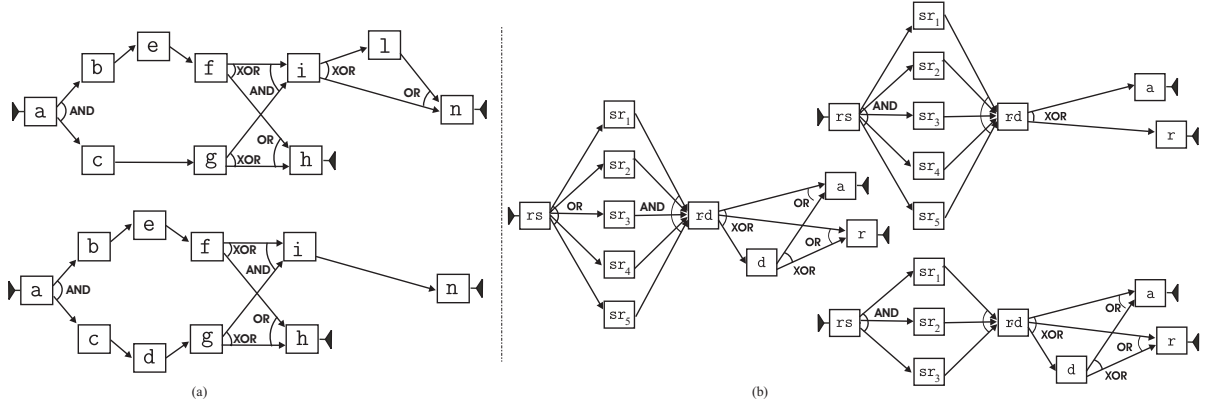


Fig. 8. Results for the ProcessDiscovery algorithm on example processes.

is allowed to exploit any available mining algorithm to equip each cluster with a schema.

In this section, we illustrate the results of experimental activity aimed at assessing the practical effectiveness of the proposed approach, the theoretical guarantee on the scaling (cf. Corollary V.7) and the effectiveness in deriving conformant models (cf. Theorem V.1).

A. Qualitative Results

We start our analysis by discussing the results of ProcessDiscovery on some example scenarios. We firstly considered the *OrderManagement* process, and we randomly generated 5,000 traces for the workflow schema in Figure 1. Notably, in the generation of the log, we also required that task *m* could not occur in any trace containing *e*, and that task *l* could not appear in any trace containing *d*, thereby modelling the restriction that a fidelity discount is never applied to a new customer, and that a fast dispatching procedure cannot be performed whenever some external supplies were asked for.

As a matter of facts, these constraints determine different usage scenarios, which are mixed up in the schema in Figure 1, thereby leading to a low value of soundness. Conversely, by running ProcessDiscovery with $maxFeature = 5$, $maxLevels = 1$, $k = 4$, $\sigma = 0.05$, $\gamma = 0.01$ and $\ell = 5$, the four schemas associated with the discovered clusters are the two of Figure 2 plus the two reported in Figure 8.(a). Each of these schemas captures, in fact, exactly one of the possible usage scenarios for the process and models one of the clusters discovered in Example V.5. This mined model is both 1-sound and 1-complete, and yields a clearer picture of the latent behavior of the process.

As a further example, let us consider the *ReviewPaper* process of handling the revisions for a paper submitted to a scientific conference. The process consists of the tasks: (*rs*)

receiving the submission, (\mathbf{sr}_i , $1 \leq i \leq 5$), sending the paper to the reviewers, (\mathbf{rd}) receiving the revisions and taking a decision, (\mathbf{d}) discussing on the paper in the case revisions are not uniform, (\mathbf{a}) accepting the paper, and (\mathbf{r}) rejecting the paper. Actually, in the case where the paper is authored by a program committee member, it has to be reviewed by 5 reviewers and it is immediately rejected in the case some reviewer does not want it to be accepted for publication. Otherwise, only 3 reviewers are assigned to the paper. A possible workflow schema for the process is reported on the left of Figure 8.(b) — notice, e.g., that \mathbf{rs} is an *or-split*. According to it and to the above specified rules, a log (of 5,000 traces) was randomly generated and `ProcessDiscovery` was invoked with $\mathit{maxFeature} = 2$, $\mathit{maxLevels} = 1$, $k = 2$, $\sigma = 0.05$, $\gamma = 0.01$ and $\ell = 5$. The two resulting schemas are shown on the right of the same figure. It should be clear that one schema is, in fact, the 1-sound model for handling the revision of a paper written by a program committee member, while the other is the 1-sound model for handling the revision of all the other papers — notice, e.g., that for both schemas, \mathbf{rs} is an *and-split* task, now.

The examples above have put into evidence that the clustering technique is very effective in providing insights into a process especially in the case where the enactments are constrained by some kinds of behavioral rules, possibly involving information which is beyond the pure execution of activities (e.g., stored in some database). This is quite a common situation in practical applications. And, in fact, research in modelling languages already evidenced the importance of these properties, that cannot be captured by a graph model, and that, in the current workflow management systems, are either left unexpressed or modelled by means of some form of logics.

We conclude by noticing that even in the case where no behavioral rules are defined and, hence, there is only one usage scenario, the *ProcessDiscovery* algorithm is still useful in order to identify some (hidden) variants, which correspond to anomalies and malfunctioning in the system. In these cases, the effect of the algorithm is to identify the “normal” behavior of the process and to single out the instances that are deviant w.r.t. it. This intuition has been confirmed by several experiments on real data sets available at <http://www.processmining.org>. For the sake of completeness, in Figure 9 we report the hierarchy built for an example log file (cf. *a12f0n20.xml*, $\mathit{maxFeature} = 3$, $k = 3$, $\sigma = 0.01$, $\gamma = 0.4$ and $\ell = 5$) and the models associated with each node in the first level of

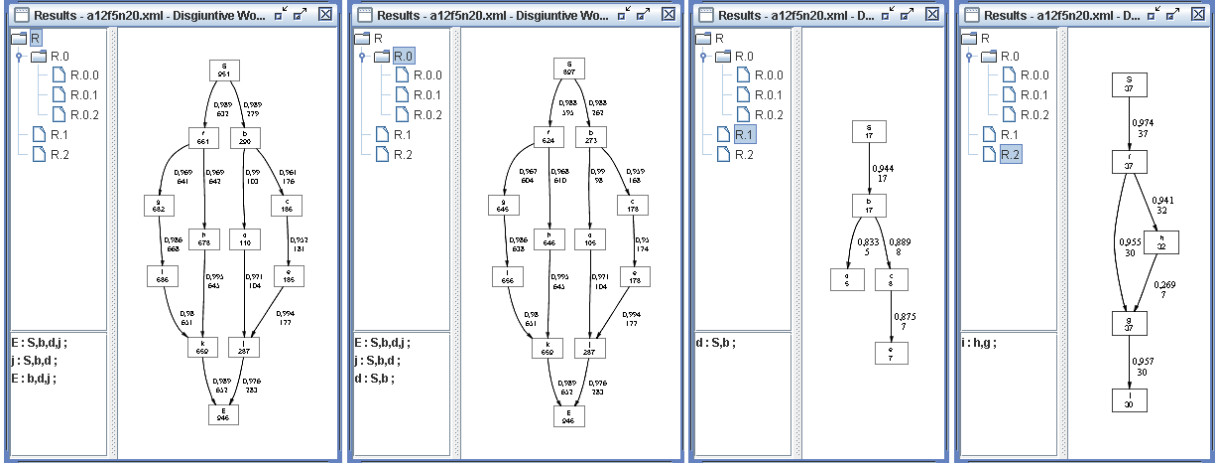


Fig. 9. Results on log data available at <http://www.processmining.org>.

the hierarchy. The algorithm discovered one large cluster $R0$ whose schema coincides with the one the root R is equipped with. Moreover, clusters $R1$ and $R2$ (containing 17 and 37 traces, respectively) may, indeed, be perceived as outliers w.r.t. the discovered main behavior.

B. Quantitative Results

Besides qualitative tests, we performed an extensive experimentation on synthetic logs, produced by means of a generator which takes advantage of ideas exposed in [33] and extends the one described in [23]. Notably, the generator can be used to produce a log of traces which are compliant with a given workflow schema, as well as to generate random workflow schemas.

Text Procedure. In order to assess the effectiveness of the technique, we defined a simple test procedure for comparing available workflow models with the output of the `ProcessDiscover` algorithm. Let \mathcal{WS} be a given workflow schema and \mathcal{L}_P be a log of traces compliant with \mathcal{WS} , produced by means of the generator. Then, the result of each test is a disjunctive workflow schema \mathcal{WS}^\vee extracted by providing the `ProcessDiscovery` algorithm with \mathcal{L}_P as input. Notice that for evaluating the quality of a mined schema we are only interested in computing its soundness, since it was proven, by Theorem V.1, that any schema discovered by means of our technique has maximal completeness. In particular, the soundness of \mathcal{WS}^\vee is estimated by computing the percentage of the traces in a log L_{test} (randomly generated from \mathcal{WS}^\vee) that are also compliant with the original

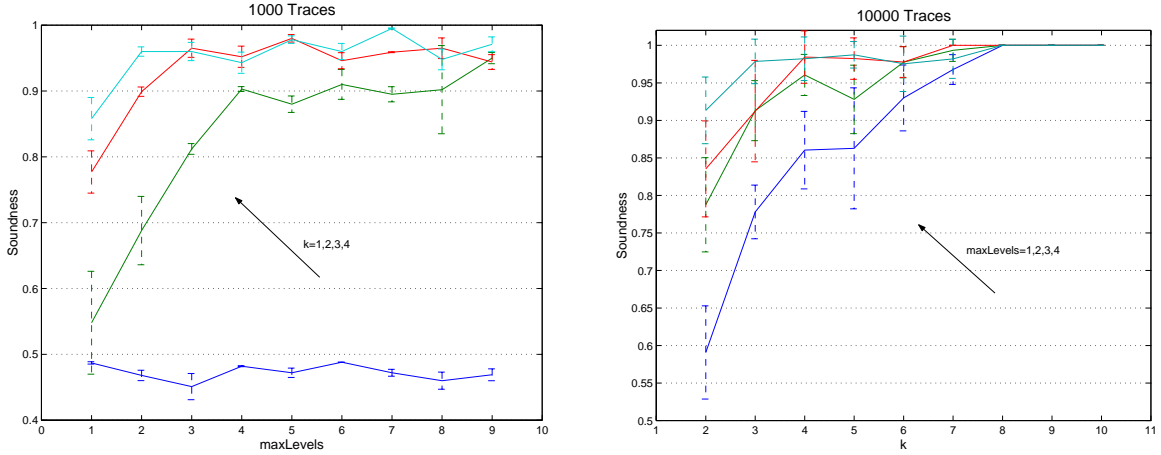


Fig. 10. Fixed Workflow: Soundness w.r.t. nr. of levels (left) and k (right)

schema \mathcal{WS} . Ideally, when L_{test} contains all the possible traces of \mathcal{WS}^V , the estimate and the actual value coincide.

In the light of Corollary V.7, the experiments mainly focus on the influence of the branching factor k and on the number of levels ($maxLevels$) in the hierarchy of clusters — notice that $maxLevels$ and k in fact determine the size of the mined disjunctive schema (m). The parameter ℓ , i.e., the length of the features, is instead kept fixed to 5, since after several experiments this appeared to be a good compromise between running time and quality of results. And, in fact, for $\ell > 5$ the soundness of the mined schema rarely improves. All the tests were conducted on a 1600MHz/256MB Pentium IV machine running *Windows XP Professional*.

Results. A first set of experiments was conducted to assess the soundness of the mined models. To this aim, we considered a fixed workflow schema at time and some log traces randomly generated according to it. This first set of experiments was repeated for different synthetic and real schemas, and the algorithm performed quite similarly under all the circumstances. Moreover, for each schema, we generated a number of different training logs. Here, we report the average values and their associated standard deviation.

For a process with 40 tasks, the soundness values of the mined model are graphically illustrated in Figure 10, which reports the mean and standard deviation of the soundness obtained with different values of k and $maxLevels$, over input logs consisting of 1,000 traces (on the left) or 10,000 traces (on the right). Both charts in the figure show that the quality of the mined schema generally gets better when either of these parameters increases, up to achieving the maximum value of soundness. It is interesting to express a

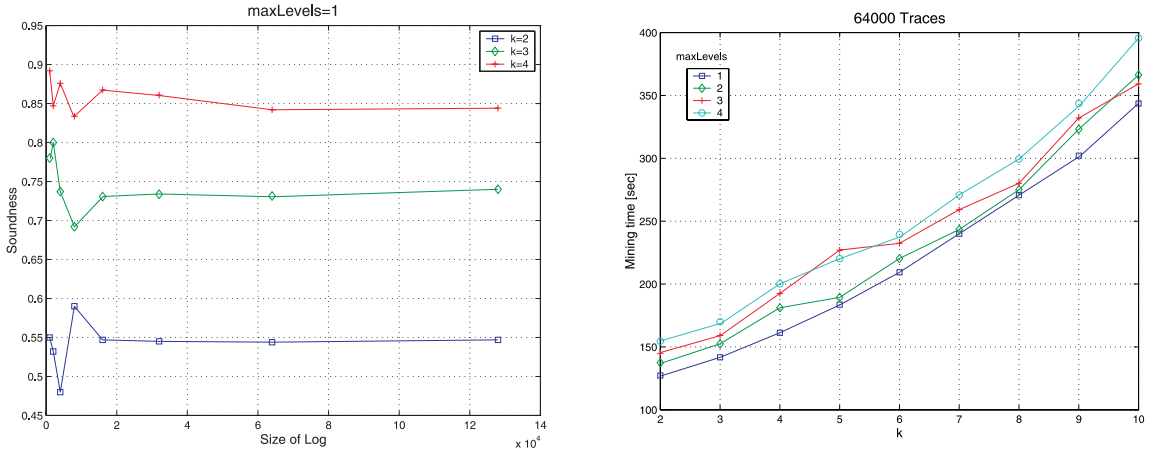


Fig. 11. Fixed Workflow: Soundness (left) and Mining Time (right) w.r.t. k .

further remark about the way the branching factor k impacts on soundness. As a matter of fact, observe that for $k = 1$ the algorithm degenerates in computing a unique schema and, hence, the soundness is not affected by the parameter $maxLevels$. On the contrary, for $k > 1$ the algorithm is able to rediscover the original schema, after performing a suitable number of iterations — see, in particular, the right side of Figure 10.

Notice that the case where $k = 1$ is that of all classical process mining algorithms (see next section), because no clustering is performed. However, since the aim of the paper is to demonstrate the effectiveness of the clustering approach (independently of the way `MineWorkflow` is carried out), we are not interested in a thorough comparison here. We only evidence that, while we often get a low value of soundness in the case where $k = 1$, mined schemas are always guaranteed to be 1-complete no matter of the log, whereas most of previous algorithms either assume that the log itself is complete or accept incompleteness as manifestation of noise. Thus, these approaches in process mining are orthogonal and enjoy specific advantages, and their comparison is left as a subject for further research.

A second set of experiments was aimed at providing more insight on the impact of k on effectiveness and efficiency. Input data are the same of those used for above experiments. Figure 11 shows the results. In particular, the left side of the figure confirms that k strongly impacts the soundness, almost independently of the log size $|\mathcal{L}_P|$, provided that it is big enough to reduce the effects of statistical fluctuations in the log composition. Moreover, the right side of the same figure shows that also the total time needed for building a schema increases when we use higher values for either k or $maxLevels$. However, as expected, the scaling is linear in both parameters. In particular, higher values for $maxLevels$ only mildly

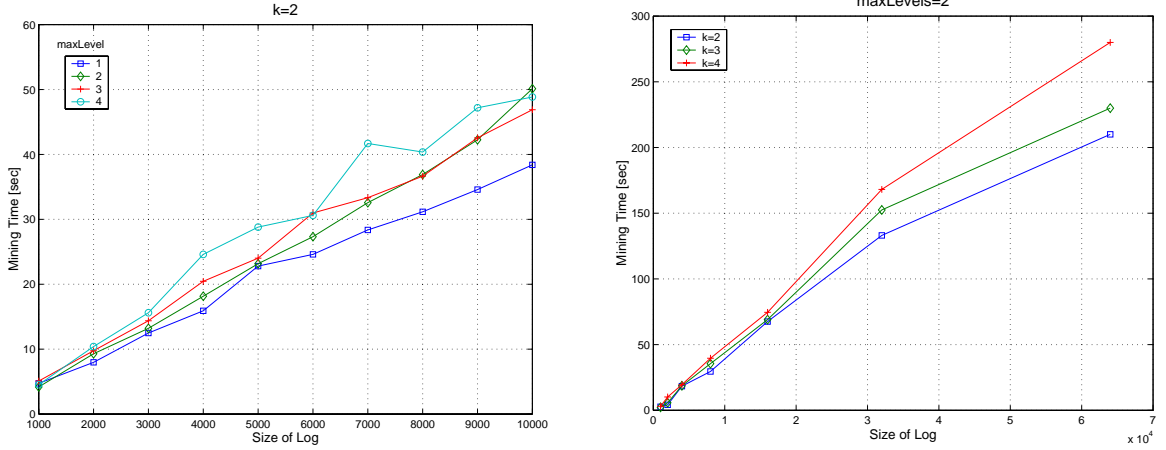


Fig. 12. Fixed Workflow: Scalability w.r.t. log size, varying the nr. of levels (left) and k (right).

increase the running time; this observation and the behavior shown on the left of Figure 10 are arguments in favor of using small values for k (see the discussion in Section V-A).

In Figure 12, we report the trend of the mining time with respect to the log size. In particular, in the left side several curves are plotted which correspond to different values of $maxLevels$, with a fixed branching factor ($k = 2$); analogously, in the right side it is k to be varied, while $maxLevels$ is always set to 2. All these curves substantiate the good scalability of the approach, which takes a time that linearly depends on the number of traces used as input.

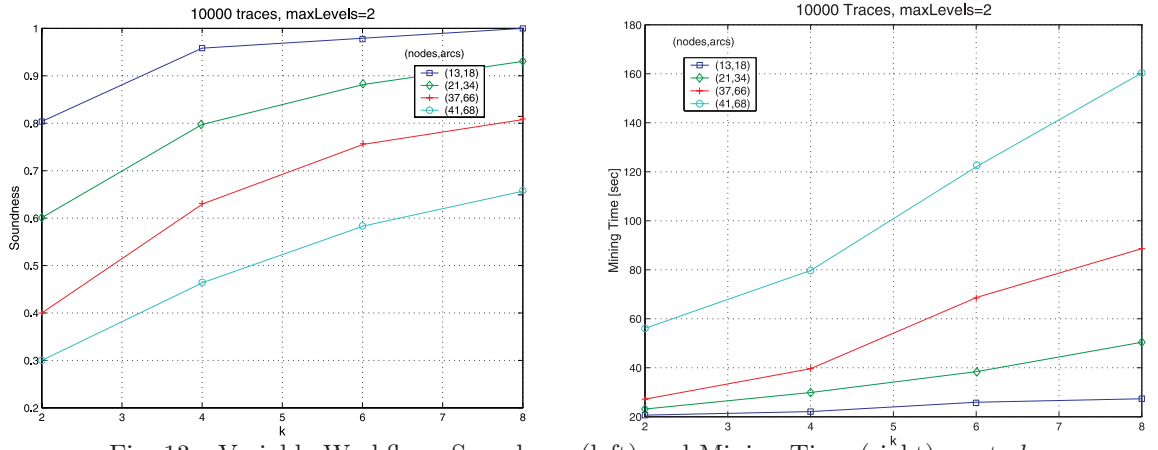


Fig. 13. Variable Workflow: Soundness (left) and Mining Time (right) w.r.t. k .

In the last series of experiments, we randomly generated several workflow schemas with different number of tasks and precedence relationships. In Figure 13, we report the results obtained with four workflow schemas. Observe (on the left) that for a fixed value of k , the soundness of the mined schema tends to become lower as the schema complexity augments, i.e., the number of activities, links and constraints increases. This witnesses the fact that

in order to have an effective reconstruction of the process it is necessary not only to fix $k > 1$, but also to deal with several levels of refinements. Obviously, for complex schemas the algorithm takes more time, as shown in the same figure on the right.

VII. OVERVIEW OF PROCESS MINING ALGORITHMS

We next briefly review some previous works on process mining, which constituted a fundamental source of inspiration for our research. In fact, these approaches solve the $\text{MPD}(\mathcal{L}_P, 1)$ problem and may be even used in our algorithm in place of `MineWorkflow`. A broader, and up-to-date enough, overview on this topic can be found in [1].

Process mining has been firstly introduced in a Software Engineering setting by [37], and subsequently extended in [5], [38]. The paper proposes three techniques, namely (i) a statistical approach, based on neural networks, (ii) a purely algorithmic approach, and (iii) a hybrid approach, based on Markov models, for automatically deriving a formal model from execution's log. The model is a Finite State Machine (FSM) model, where the activities are associated with the edges and specify transitions between states.

In [8], processes are more naturally represented through pure directed graphs, which yet allow to express precedence relationships only, by disregarding richer control flow constructs, such as concurrency, synchronization and choice. Notably, techniques in [8] have been designed to work even in presence of cyclic dependencies between activities.

A special kind of Petri nets, named Workflow nets (WF-net), is adopted in [39], [40], [1], [21] for modelling and mining workflow processes. There, each transition represents a task, while the relationships between the tasks are modelled by arcs and places. Importantly, WF-nets allows for recursive schemas and iterated executions. A basic algorithm, called α -algorithm, is introduced in [21], which is able to derive a WF-net from a workflow log, under the assumption that the log is complete and free of noise. However, the α -algorithm can easily deal with cycles, which is a functionality currently missing in our approach. Actually, the capability of the algorithm to mine WF-net workflow models and its limitations are analyzed in [1], where the concept of *structured workflow (SWF) net* is introduced to capture a class of WF-nets which a process mining algorithm should be able to rediscover. Some extensions to the approach are presented which address these problems. In particular, in [40], simple statistics are exploited in the construction of the ordering relations in order to cope with noise in the logs, whereas specific preprocessing

and postprocessing strategies for capturing short loops are devised in [41].

A further approach to mining a process model from event logs is described in [7], [42], where a subset of the ADONIS language [43] is adopted to represent a process model. An important peculiarity of the approach mainly resides in its capability of recognizing duplicate tasks in the control flow graph, i.e., many nodes associated with the same task.

Yet another approach is adopted in [9], where a mining tool is presented which is able to discover hierarchically structured workflow processes. Such a model corresponds to an expression tree, where the leafs represent tasks (operands) while any other node is associated with a control flow operator. In this context, the mining algorithm mainly consists of a suitable set of term rewriting systems.

The possibility of grouping workflow traces has been previously explored by [11]. The basic idea is to exploit a multi-phase process mining approach, where individual models are firstly generated for each process instance by means of the technique in [10] and, subsequently, aggregated into *aggregation graphs*. These graphs may be eventually translated into EPCs or Petri nets. Interestingly, the first phase can be possibly avoided in the case where logs are registered by some tools such as ARIS PPM [44], where individual models are already available in terms of *instance graphs* and do not need to be preliminary mined.

As a matter of facts, the technique in [11] can be used to aggregate any predetermined set of instance graphs and, therefore, it does not fit our setting where the traces to be aggregated are not known in advance (in fact, clusters are the result of an automatic partitioning on the log based on some suitably extracted features). Yet, the grouping technique can be still used to implement some kinds of *agglomerative* clustering algorithm. Indeed, an algorithm may start with each instance graph associated with an aggregation graph and, at each step, it may select the most similar aggregation graphs (according to some suitably defined metric) whose instances are in turn fused into a more general aggregation graph. By this way, the hierarchy of schemas may be defined in a bottom-up fashion, rather than in a top-down one. While agglomerative approaches are often less efficient than divisive strategies, we believe that the above sketched approach is worth being formalized and analyzed in order to assess its practical effectiveness.

VIII. CONCLUSIONS

In this paper, we have continued on the way of investigating data mining techniques for discovering process models from event logs. We have devised a novel framework that substantially differs from previous approaches, for it performs a hierarchical clustering of the logs, in which each trace is seen as a point of a properly identified space of features. The resulting model is a disjunctive schema that explicitly takes care of variants of the process. The computational complexity of the different problems involved in our investigation has been thoroughly investigated. It turned out that even checking whether a given model at hand is conformant is a difficult task, so that any efficient algorithm for computing the best schema has to search into the space of the possible solutions by never checking for the soundness. Our solution to this problem was a smart algorithm that stepwise refine a starting schema by guaranteeing that each refinement leads to a sounder schema. The performances of the proposed approach have been analyzed over a number of data sets, thereby getting some appreciable evidence for its effectiveness and scalability.

Notably, the proposed approach is to large extent independent of the adopted workflow model and is, indeed, modular and able to benefit of other results available in the literature for dealing with more elaborate features, such as cyclicity in the log. And, in fact, our clustering algorithm has been made recently available as an *analysis* plug-in for the *ProM* process mining framework, which allows to exploit well-established process mining algorithms to equip clusters with models. An extensive investigation of this kind of integrated approach constitutes an avenue of further research.

We conclude by observing that, the whole framework proposed in the paper is essentially *propositional*, for it assumes a simplification of the schema and of the enactments in which many real-life details are omitted. This is a standard assumption in current research in process mining. Therefore, another interesting avenue for further research is to extend our techniques to take care of the environment, so that clusters may reflect not only structural similarities among traces, but also information about, e.g., users and data values.

REFERENCES

- [1] W.M.P. van der Aalst, B.F. van Dongen, J. Herbst, L. Maruster, G.Schimm, and A.J.M.M. Weijters, "Workflow mining: A survey of issues and approaches," *Data and Knowledge Engineering*, vol. 47, no. 2, pp. 237–267, 2003.

- [2] A. Rozinat and W.M.P. van der Aalst, “Conformance testing: Measuring the fit and appropriateness of event logs and process models,” in *Proc. Int. Workshop on Business Process Intelligence (BPI'05)*, 2004, pp. 1–12.
- [3] F. Casati, M. Castellanos, and M. Shan, “Enterprise cockpit for business operation management,” in *Proc. of the 23rd International Conference on Conceptual Modeling (ER2004)*, 2004, pp. 825–827.
- [4] D.-R. Liu and M. Shen, “Workflow modeling for virtual processes: an order-preserving process-view approach,” *Information Systems*, vol. 28, pp. 505–532, 2003.
- [5] J.E. Cook and A.L. Wolf, “Event-based detection of concurrency,” in *Proc. 6th Int. Symposium on the Foundations of Software Engineering (FSE'98)*, 1998, pp. 35–45.
- [6] A.J.M.M. Weijters and W.M.P. van der Aalst, “Rediscovering workflow models from event-based data using little thumb,” *Integrated Computer-Aided Engineering*, vol. 10, no. 2, pp. 151–162, 2003.
- [7] J. Herbst, “Dealing with concurrency in workflow induction,” in *Procs. European Concurrent Engineering Conference*, 2000.
- [8] R. Agrawal, D. Gunopulos, and F. Leymann, “Mining process models from workflow logs,” in *Proc. 6th Int. Conf. on Extending Database Technology (EDBT'98)*, 1998, pp. 469–483.
- [9] G. Schimm, “Mining most specific workflow models from event-based data,” in *Proc. of Int. Conf. on Business Process Management*, 2003, pp. 25–40.
- [10] B.F. van Dongen and W.M.P. van der Aalst, “Multi-phase process mining: Building instance graphs,” in *Proc. Int. Conf. on Conceptual Modeling (ER)*, 2004, pp. 362–376.
- [11] B.F. van Dongen and W.M.P. van der Aalst, “Multi-phase process mining: Aggregating instance graphs into EPCs and Petri Nets,” in *Proc. Int. Work. on Applications of Petri Nets to Coordination, Workflow and Business Process Management (PNCWB) at the ICATPN 2005*, 2005.
- [12] H. Davulcu, M. Kifer, C.R. Ramakrishnan, and I.V. Ramakrishnan, “Logic based modeling and analysis of workflows,” in *Proc. of the 17th ACM Symposium on Principles of Database Systems (PODS'98)*, 1998, pp. 25–33.
- [13] P. Muth, J. Weifenfels, M.Gillmann, and G. Weikum, “Integrating light-weight workflow management systems within existing business environments,” in *Proc. 15th IEEE Int. Conf. on Data Engineering (ICDE'99)*, 1999, pp. 286–293.
- [14] P. Senkul, M. Kifer, and I.H. Toroslu, “A logical framework for scheduling workflows under resource allocation constraints,” in *Proc. 28th Int. Conf. on Very Large Data Bases (VLDB'02)*, 2002, pp. 694–702.
- [15] H. Schuldt, G. Alonso, C. Beerli, and H. Schek, “Atomicity and isolation for transactional processes,” *ACM Trans. Database Syst.*, vol. 27, no. 1, pp. 63–116, 2002.
- [16] W.M.P. van der Aalst, A. Hirnschall, and H.M.W. Verbeek, “An alternative way to analyze workflow graphs,” in *Proc. 14th Int. Conf. on Advanced Information Systems Engineering*, 2002, pp. 534–552.
- [17] Mohan Kamath and Krithi Ramamritham, “Correctness issues in workflow management,” *Distributed Systems Engineering*, vol. 3, no. 4, pp. 213–221, 1996.
- [18] W.M.P. van der Aalst, “The application of petri nets to workflow management,” *Journal of Circuits, Systems, and Computers*, vol. 8, no. 1, pp. 21–66, 1998.
- [19] W.M.P. van der Aalst, J. Desel, and E. Kindler, “On the semantics of EPCs: A vicious circle,” in *Proc. EPK 2002: Business Process Management using EPCs*, 2002, pp. 71–80.
- [20] G. Keller, M. Nüttgens, and A.W. Scheer, *Semantische Prozessmodellierung auf der Grundlage Ereignisgesteuerter Prozessketten (EPK)*, University of Saarland, Saarbrücken, 1992.
- [21] W.M.P. van der Aalst, A.J.M.M. Weijters, and L. Maruster, “Workflow mining: Discovering process models from event logs,” *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 16, no. 9, pp. 1128–1142, 2004.
- [22] G.Greco, A.Guzzo, L.Pontieri, and D. Saccà, “Discovering expressive process models by clustering log traces [Appendix],” available at <http://www.icar.cnr.it/wfmining>, 2005.
- [23] G. Greco, A. Guzzo, G. Manco, and D. Saccà, “Mining and reasoning on workflows,” *IEEE Trans. on Data and Knowledge Engineering*, vol. 17, no. 4, pp. 519–534, 2005.

- [24] G. Greco, A. Guzzo, and L. Pontieri, "Mining hierarchies of models: From abstract views to concrete specifications," in *Proc. of Int. Conf. on Business Process Management*, 2005, pp. 32–47.
- [25] Rob J. van Gabbeek and W. Peter Weijland, "Branching time and abstraction in bisimulation semantics," *Journal of ACM*, vol. 43, no. 3, pp. 555–600, 1996.
- [26] T. Basten and W. van der Aalst, "Inheritance of workflows: an approach to tackling problems related to change," *Theoretical Computer Science*, vol. 270, no. 1–2, pp. 125–203, 2002.
- [27] V. Guralnik and G. Karypis, "A scalable algorithm for clustering sequential data," in *Proc. IEEE Int. Conf. on Data Mining (ICDM'2001)*, 2001, pp. 179–186.
- [28] J. Han, J. Pei, B. Mortazavi-Asl, U. Dayal, and M. Hsu, "Freespan: frequent pattern-projected sequential pattern mining," in *Proc. Int. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD'00)*, 2000, pp. 355–359.
- [29] Y.S. Kim, W.N. Street, and F. Menczer, "Feature selection in unsupervised learning via evolutionary search," in *Proc. 6th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD'00)*, 2000, pp. 365–369.
- [30] N.Lesh, M.J. Zaki, and M.Ogihara, "Mining features for sequence classification," in *Proc. 6th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD'00)*, 1999, pp. 342–346.
- [31] H. Motoda and H. Liu, "Data reduction: feature selection," *Handbook of data mining and knowledge discovery*, pp. 208–213, 2002.
- [32] R. Agrawal, T. Imielinski, and A. Swami, "Mining association rules between sets of items in large databases," in *Proc. of SIGMOD'93*, 1993, pp. 207–216.
- [33] R. Agrawal and R. Srikant, "Fast algorithms for mining association rules," in *Proc. of the 20th Int'l Conference on Very Large Databases*, 1994, pp. 487–499.
- [34] R. Agrawal and R. Srikant, "Mining sequential patterns," in *Proc. 11th Int. Conf. on Data Engineering (ICDE95)*, 1995, pp. 3–14.
- [35] B. Padmanabhan and A. Tuzhilin, "Small is beautiful: discovering the minimal set of unexpected patterns," in *Proc. 6th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD'00)*, 2000, pp. 54–63.
- [36] B.F. van Dongen, A.K.A. de Medeiros, H.M.W. Verbeek, A.J.M.M. Weijters, and W.M.P. van der Aalst, "The prom framework: A new era in process mining tool support," in *Proc. of 26th International Conference on Applications and Theory of Petri Nets (ICATPN'05)*, 2005, pp. 444–454.
- [37] J.E. Cook and A.L. Wolf, "Automating process discovery through event-data analysis," in *Proc. 17th Int. Conf. on Software Engineering (ICSE'95)*, 1995, pp. 73–82.
- [38] J.E. Cook and A.L. Wolf, "Software process validation: quantitatively measuring the correspondence of a process to a model," *ACM Trans. Softw. Eng. Methodol.*, vol. 8, no. 2, pp. 147–176, 1999.
- [39] W.M.P. van der Aalst and K.M. van Hee, *Workflow Management: Models, Methods, and Systems*, MIT Press, 2002.
- [40] W.M.P. van der Aalst and B.F. van Dongen, "Discovering workflow performance models from timed logs," in *Proc. Int. Conf. on Engineering and Deployment of Cooperative Information Systems (EDCIS'02)*, 2002, pp. 45–63.
- [41] A.K.A de Medeiros, B.F. van Dongen, W.M.P. van der Aalst, and A.J.M.M. Weijters, "Process mining: Extending the a-algorithm to mine short loops," Tech. Rep., University of Technology, Eindhoven, 2004, BETA Working Paper Series, WP 113.
- [42] J. Herbst and D. Karagiannis, "Integrating machine learning and workflow management to support acquisition and adaptation of workflow models," *Journal of Intelligent Systems in Accounting, Finance and Management*, vol. 9, pp. 67–92, 2000.
- [43] S. Junginger, H. Kuhn, R. Strobl, and D. Karagiannis, "Ein geschäfts-prozessmanagement-werkzeug der nächsten generation - adonis: Konzeption und anwendungen," *Wirtschaftsinformatik*, vol. 42, no. 3, pp. 392–401, 2000.
- [44] IDS Scheer, "Aris process performance manager (aris ppm): Measure, analyze and optimize your business process performance (whitepaper)," <http://www.ids-scheer.com>, 2002.