

# Comparative Process Mining in Education: An Approach Based on Process Cubes

Wil M.P. van der Aalst<sup>1,2(✉)</sup>, Shengnan Guo<sup>1</sup>, and Pierre Gorissen<sup>3</sup>

<sup>1</sup> Eindhoven University of Technology, P.O. Box 513,  
5600 MB Eindhoven, The Netherlands  
{w.m.p.v.d.aalst,s.guo}@tue.nl

<sup>2</sup> International Laboratory of Process-Aware Information Systems,  
National Research University Higher School of Economics (HSE),  
33 Kirpichnaya Str., Moscow, Russia

<sup>3</sup> Fontys Hogescholen, P.O. Box 347, 5600 AH Eindhoven, The Netherlands  
p.gorissen@fontys.nl

**Abstract.** Process mining techniques enable the analysis of a wide variety of processes using event data. For example, event logs can be used to automatically learn a process model (e.g., a Petri net or BPMN model). Next to the automated discovery of the real underlying process, there are process mining techniques to analyze bottlenecks, to uncover hidden inefficiencies, to check compliance, to explain deviations, to predict performance, and to guide users towards “better” processes. Dozens (if not hundreds) of process mining techniques are available and their value has been proven in many case studies. However, existing techniques focus on the analysis of a single process rather than the comparison of different processes. In this paper, we propose *comparative process mining using process cubes*. An event has attributes referring to the dimensions of the process cube. Through slicing, dicing, rolling-up, and drilling-down we can view event data from different angles and produce process mining results that can be compared. To illustrate the process cube concept, we focus on educational data. In particular, we analyze data of students watching video lectures given by the first author. The dimensions of the process cube allow us to compare the process of students that passed the course versus the process of students that failed. We can also analyze differences between male and female students, between different parts of the course, and between Dutch students and international students. The initial analysis provided in this paper is used to elicit requirements for better tool support facilitating comparative process mining.

**Keywords:** Process mining · Online analytical processing · Learning analytics · Comparative process mining

## 1 Introduction

*Process mining* can be seen as the missing link between model-based process analysis (e.g., simulation and verification) and data-oriented analysis techniques

such as machine learning and data mining [1]. It seeks the “confrontation” between real event data and process models (automatically discovered or hand-made). As process mining techniques mature, more ambitious types of analysis come into reach. Whereas classical process mining techniques focus on a single process, this paper focuses on *comparing different processes using event data*.

In [3], we proposed the notion of *process cubes* where events and process models are organized using different dimensions. Each cell in the process cube corresponds to a set of events that can be used to discover a process model, to check conformance, or to discover bottlenecks. The idea is related to the well-known OLAP (Online Analytical Processing) data cubes and associated operations such as slice, dice, roll-up, and drill-down [20]. However, there are also significant differences because of the process-related nature of event data. For example, process discovery based on events is incomparable to computing the average or sum over a set of numerical values. Moreover, dimensions related to process instances (e.g. male versus female students), subprocesses (e.g. group assignments versus individual assignments), organizational entities (e.g. students versus lecturers), and time (e.g. years or semesters) are semantically different and it is challenging to slice, dice, roll-up, and drill-down process mining results efficiently.

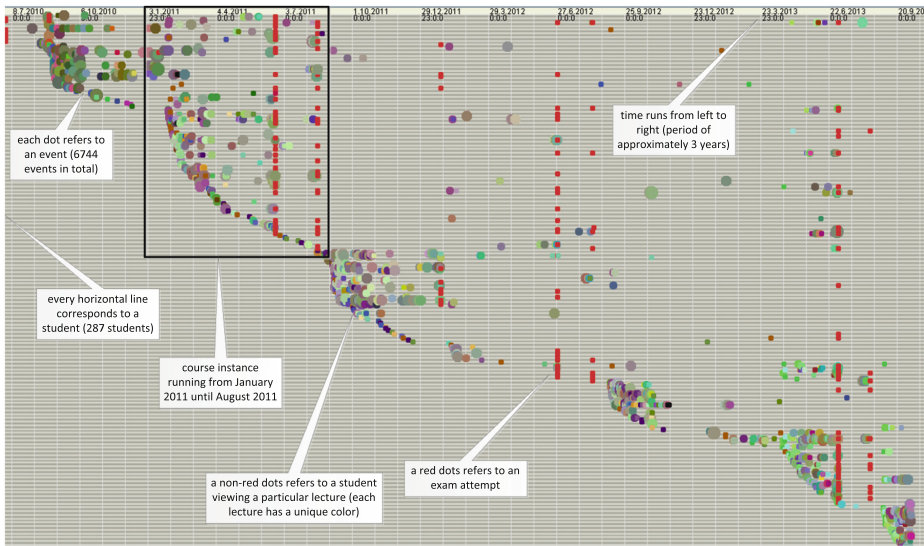
This paper focuses on *comparative process mining using process cubes*. We discuss the main challenges related to comparative process mining. To do this, we use a data set describing behavior of students taking the “Business Information Systems” (2II05) course given at Eindhoven University of Technology. The data set contains two types of events: (1) events generated by students watching recorded video lectures and (2) events generated by students making exams. To understand differences in behavior among different student groups we apply comparative process mining and sketch the possible dimensions of the process cube.

The remainder is organized as follows. In Sect. 2 we briefly introduce the process mining spectrum. Section 3 defines the process cube notion as a means to view event data from different angles. Then we discuss a concrete case study analyzing the way students watch recorded lectures and correlate this with exam results (Sect. 4). Section 5 lists the main requirements and open challenges. Related work is briefly discussed in Sect. 6. Section 7 concludes the paper.

## 2 Process Mining

Process mining provides a powerful way to analyze operational processes based on event data. Unlike classical purely model-based approaches, process mining is driven by “raw” observed behavior instead of assumptions or aggregate data. Unlike classical data-driven approaches, it is truly process-oriented and relates events to high-level end-to-end process models.

Normally, *event logs* serve as the starting point for process mining. An event log can be viewed as a multiset of *traces*. Each trace describes the life-cycle of a particular *case* (i.e., a *process instance*) in terms of the *activities* executed.



**Fig. 1.** Dotted chart showing all events related to the course Business Information Systems (2II05). The dotted chart was created using one of the over 600 *ProM* plugins (cf. [www.processmining.org](http://www.processmining.org)) (Color figure online).

Often event logs store additional information about events, e.g., the *resource* (i.e., person or device) executing or initiating the activity, the *timestamp* of the event, or *data elements* recorded with the event.

Process mining has been applied in a wide variety of organizations (e.g., banks, hospitals, municipalities, governmental agencies, webshops, and high-tech system manufacturers). Moreover, there are dozens of process mining techniques answering a wide variety of questions [1]. Due to space restrictions we can only illustrate a fraction of the available process mining techniques. To do so, we use a concrete data set involving two types of events recorded for students that took the Business Information Systems (2II05) course at Eindhoven University of Technology from 2009 until 2014. A *view* event refers to a student watching a particular lecture. An *exam attempt* event refers to a student taking an exam. Since the course is quite challenging, it is not uncommon that students need to resit the 2II05 exam multiple times. There are at least two exams per year. The initial log contains 6744 events generated by 287 students.

Figure 1 shows a so-called *dotted chart* taking the viewpoint that each student corresponds to a case (i.e., process instance). The dotted chart has 287 horizontal lines each corresponding to a student. The dots along such a line define the corresponding trace. The color of the dot refers to the corresponding activity, e.g., viewing a particular lecture. The red dots in Fig. 1 refer to exam attempts. It can be seen that some students need to take multiple exams and that students tend to watch irregularly. Note that the video lectures are an additional service to the students (i.e., next to regular lectures).

Let us zoom in on the group of 47 students taking the course in the period January 2011–August 2011. These students took the exam on 21-6-2011 and/or the retake exam on 16-8-2011. There were 24 lectures (two lectures per week) in the period from January until May 2011. Students could watch the lectures via an internet connection soon after recording. For example, activity “2II05 College 11b” refers to the second lecture in the 11th week. Figure 2 shows a process model discovered for this event log. Indeed students tend to watch the videos in chronological order. However, Fig. 2 only considers the most frequent paths. The actual process is more “Spaghetti-like”.

We have made a process model having 24 viewing activities in sequence followed by an exam. Hence, the model only allows for the trace ⟨“2II05 College 1a”, “2II05 College 1b”, “2II05 College 2a”, “2II05 College 2b”, . . . , “Exam”⟩. Using *ProM* we can check the conformance of such a model in various ways. For example, we can compute alignments that map the traces in the event log to valid paths in the model [5,10]. Figure 3 shows four such alignments. Overall, the fitness is low (0.33) showing that despite the tendency to watch lectures in the expected order, few students do so consistently. Also the times at which students watch the videos show a lot of variation. For example, it can be noted that, other than just before the exam, students rarely watch video lectures in the second half of the course.

### 3 Process Cubes

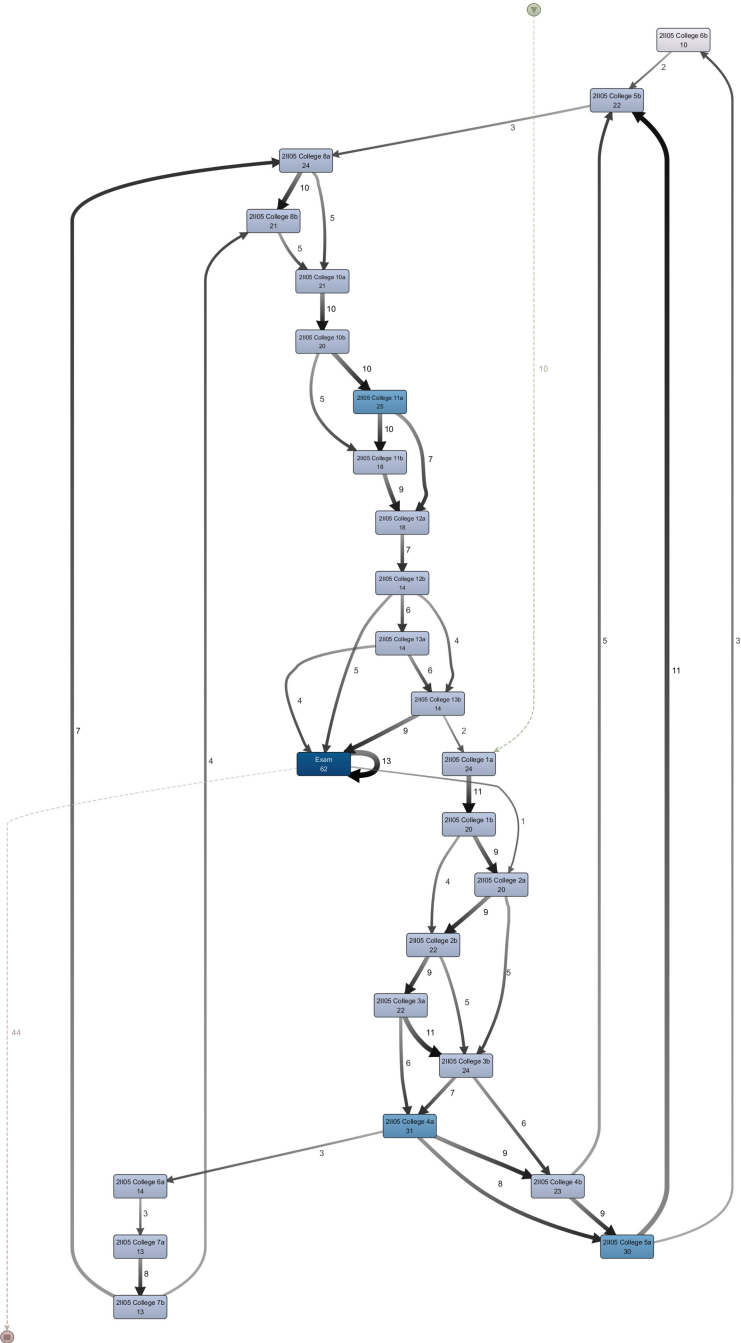
Figures 1, 2 and 3 show some example results computed for a given event log, i.e., a collection of events. However, often the goal is to compare different variants of the same process or to zoom in on particular parts or aspects of the process. OLAP (Online Analytical Processing) tools aim to support this by organizing data in a cube having multiple dimensions [20]. However, OLAP is only used for numerical comparisons, e.g., the average transaction amount in different shops on different days in the week. In a *process cube* [3] we organize events using different dimensions and compute a (process) model per sublog associated to a cell. This way we can slice, dice, roll-up, and drill-down process mining results easily.

Throughout the paper we assume the following universes [3].

**Definition 1 (Universes).**  $\mathcal{U}_V$  is the universe of possible attribute values (e.g., strings, numbers, etc.).  $\perp \in \mathcal{U}_V$  denotes a missing (“null”) value.  $\mathcal{U}_S = \mathcal{P}(\mathcal{U}_V)$  is the universe of value sets.  $\mathcal{U}_H = \mathcal{P}(\mathcal{U}_S)$  is the universe of value set collections (set of sets).

Note that  $v \in \mathcal{U}_V$  is a single value (e.g.,  $v = 300$ ),  $V \in \mathcal{U}_S$  is a set of values (e.g.,  $V = \{male, female\}$ ), and  $H \in \mathcal{U}_H$  is a collection of sets. For example,  $H = \{\{x \in \mathbb{N} \mid x < 50\}, \{x \in \mathbb{N} \mid 40 \leq x < 80\}, \{x \in \mathbb{N} \mid x \geq 70\}\}$ .

An *event base* is a “raw” collection of *events* having *properties*.



**Fig. 2.** Process model created using *Disco* while abstracting from infrequent paths. There are 24 activities corresponding to the video lectures and one activity corresponding to an exam attempt. The coloring indicates the frequency of the activities.



**Fig. 3.** Alignments for four students showing that few students actually watch the lectures sequentially. The “synchronous moves” show where model and log agree. The “moves on model” show where the student skipped a video lecture. The “moves on log” show where the student watched a video lecture that was not next in line.

**Definition 2 (Event Base, [3]).** An event base  $EB = (E, P, \pi)$  defines a set of events  $E$ , a set of event properties  $P$ , and a function  $\pi \in P \rightarrow (E \rightarrow \mathcal{U}_V)$ . For any property  $p \in P$ ,  $\pi(p)$  (denoted  $\pi_p$ ) is a function mapping an event onto a value for property  $p$ . If  $\pi_p(e) = v$ , then event  $e \in E$  has a property  $p \in P$  and the value of this property is  $v \in \mathcal{U}_V$ . We write  $\pi_p(e) = \perp$  for missing values.


Independent of the event base  $EB$  we define the *structure* of the process cube. The structure is fully characterized by the *dimensions* of the cube.

**Definition 3 (Process Cube Structure, [3]).** A process cube structure is a triplet  $PCS = (D, type, hier)$  where:

- $D$  is a set of dimensions,
- $type \in D \rightarrow \mathcal{U}_S$  is a function defining the possible set of values for each dimension, e.g.,  $type(age) = \{0, 1, 2, \dots, 120\}$  for  $age \in D$ , and
- $hier \in D \rightarrow \mathcal{U}_H$  defines a hierarchy for each dimension such that for any  $d \in D$ :  $type(d) = \bigcup hier(d)$ .

Note that a hierarchy is merely a collection of sets of values. To relate an event base and a process cube structure, both need to be *compatible*, i.e., dimensions should correspond to properties ( $D \subseteq P$ ) and concrete event property values need to be of the right type ( $\pi_d(e) \in type(d)$  for any  $d \in D$  and  $e \in E$ ). Moreover, for process mining we often assume that  $\{case, activity, time, resource\} \subseteq D \subseteq P$ , i.e., each event refers to a case and an activity, occurred at a particular time, and was executed by a particular resource. These properties do not need to be

decided upfront and can be changed during analysis. For example, the notion of *case* may be changed to create another viewpoint (see Chap. 4 in [1]).



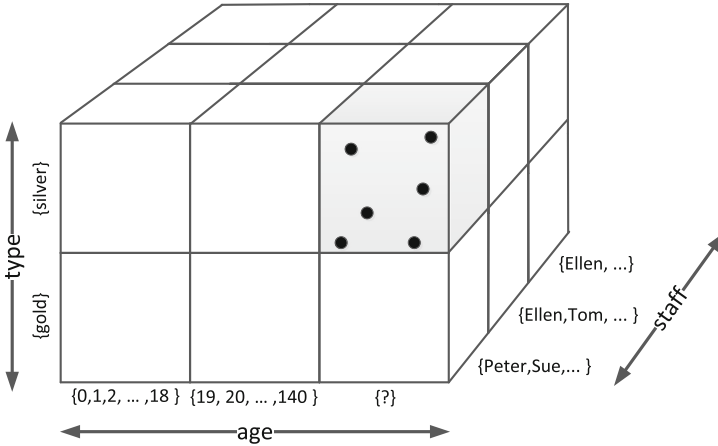
**Fig. 4.** Two example hierarchies, both having three levels. The left-hand-side hierarchy can be used to group events according to the patient’s age. The right-hand-side hierarchy can be used to group events according to the role of the resource performing the activity.

To clarify the notion of a *process cube structure*, let us consider an example  $PCS = (D, type, hier)$ .

- $D = \{patient\_id, type, age, activity, staff, time, \dots\}$  defines the set of dimensions.
- *type* is a function mapping each dimension onto a set of possible values:
  - $type(patient\_id) = \{99000, 99001, \dots, 99999\}$  is the set of patient identifiers (this value *can* be used as a case identifier),
  - $type(type) = \{gold, silver\}$  is the set of patient types (patients of type *gold* have a better insurance allowing for extra privileges),
  - $type(age) = \{0, 1, 2, \dots, 140, ?\}$  is the set of possible ages (value “?” denotes that the age of the patient is unknown),
  - $type(activity) = \{blood\_test, doctor\_visit, X\_ray, handle\_payment, \dots\}$  is the set of activities,
  - $type(staff) = \{Peter, Sue, Ellen, Tom, \dots\}$  is the set of resources (doctors, nurses), and
  - $type(time)$  is the set of possible timestamps.
- *hier* is a function defining a hierarchy for each dimension:
  - $hier(opatient\_id) = \{\{99000, 99001, \dots, 99999\}, \{99000\}, \{99001\}, \{99002\}, \dots, \{99999\}\}$ . The element  $\{99000, 99001, \dots, 99999\}$  can be seen as the root of the hierarchy, i.e., all possible values. The other singleton elements can be seen as the leaves of the hierarchy, e.g., the set  $\{99023\}$  represents one individual patient.
  - $hier(type) = \{\{gold, silver\}, \{gold\}, \{silver\}\}$ . Set  $\{gold, silver\}$  can be seen as the root of the hierarchy, i.e., all patient types. The singleton  $\{gold\}$  refers to patients having special privileges and the singleton  $\{silver\}$  refers to “normal” patients.

- $hier(age) = \{\{0, 1, 2, \dots, 140, ?\}, \{0, 1, 2, \dots, 18\}, \{19, 20, \dots, 140\}, \{0\}, \{1\}, \{2\}, \dots, \{140\}, \{?\}\}$  defines an age hierarchy with three levels (see Fig. 4), e.g., the element  $\{0, 1, 2, \dots, 18\}$  represents the subset “young”.
- $hier(activity) = \{\{blood\_test, doctor\_visit, X\_ray, handle\_payment, \dots\}, \{blood\_test\}, \{doctor\_visit\}, \dots\}$  groups all activities in a simple hierarchy,<sup>1</sup>
- $hier(staff) = \{\{Peter, Sue, Ellen, Tom, \dots\}, \{Peter, Sue, \dots\}, \{Ellen, Tom, \dots\}, \{Ellen, \dots\}, \{Peter\}, \{Sue\}, \{Ellen\}, \{Tom\}, \dots\}$  defines a role hierarchy as shown in Fig. 4. Staff member *Ellen* is both a doctor and a surgeon and appears in four subsets of  $hier(staff)$ .
- $hier(time)$  can be used to group timestamps into years, months, weekdays, days, etc.

Figure 4 shows two hierarchies. The arcs are based on inclusion and can be interpreted as “is a”. In this paper we use a rather simplistic, but also general, notion of hierarchy. Only the possible elements of a dimension are mentioned and no further constraints are given. We do not specify which elements can be used at the same time (see also [3]). Normally, one will make the different levels explicit and only select elements of a given level. We also do not specify navigation rules and do not explicitly name the sets in a hierarchy. The names used Fig. 4 (“all”, “young”, and “nurse”) are not formalized, but should of course be supported by software tools.



**Fig. 5.** Process cube view having three selected dimensions:  $D_{sel} = \{type, age, staff\}$ .

A *process cube view* defines which dimensions are visible and which events are selected.

<sup>1</sup> The  $hier(activity)$  hierarchy could be further refined based on some taxonomy. Here we just consider individual activities (singleton sets) and the set of all activities.



**Definition 4 (Process Cube View, [3]).** Let  $PCS = (D, type, hier)$  be a process cube structure. A process cube view is a pair  $PCV = (D_{sel}, sel)$  such that:

- $D_{sel} \subseteq D$  are the selected dimensions,
- $sel \in D \rightarrow \mathcal{U}_H$  is a function selecting the part of the hierarchy considered per dimension. Function  $sel$  is such that for any  $d \in D$ :  $sel(d) \subseteq hier(d)$ .

Figure 5 shows an example of a process cube view.  $D_{sel} = \{type, age, staff\}$  are the selected dimensions.  $sel(type) = \{\{gold\}, \{silver\}\}$ ,  $sel(age) = \{\{0, 1, 2, \dots, 18\}, \{19, 20, \dots, 140\}, \{?\}\}$ , and  $sel(staff) = \{\{Peter, Sue, \dots\}, \{Ellen, Tom, \dots\}, \{Ellen, \dots\}\}$ . Note that function  $sel$  returns a set of sets for each dimension. For the non-selected dimensions  $D \setminus D_{sel}$  we cannot see the value of  $sel$  in Fig. 5, but still these dimensions may have been used for filtering. For example, it may be that  $sel(activity) = \{\{blood\_test\}\}$ , i.e., only events related to blood tests are considered. Moreover, it could be that  $sel(type)$  is used to select events that happened in 2013 and 2014.

By removing elements from  $D_{sel}$ , the number of dimensions is reduced. This is orthogonal to function  $sel$  which decides on the granularity and filtering of each dimension. Given an event base and a process cube view, we can compute an event log for every cell in the process cube.

**Definition 5 (Materialized Process Cube View).** Let process cube structure  $PCS = (D, type, hier)$  and event base  $EB = (E, P, \pi)$  be compatible. The materialized process cube for some view  $PCV = (D_{sel}, sel)$  of  $PCS$  is  $M_{EB, PCV} = \{(c, events(c)) \mid c \in cells\}$  with  $cells = \{c \in D_{sel} \rightarrow \mathcal{U}_S \mid \forall d \in D_{sel} c(d) \in sel(d)\}$  being the cells of the cube and  $events(c) = \{e \in E \mid \forall d \in D_{sel} \pi_d(e) \in c(d) \wedge \forall d \in D \pi_d(e) \in \bigcup sel(d)\}$  the set of events per cell.

The term “materialized” may be misleading, it is just used to express that, in order to apply standard process mining techniques, we need to create an event log (e.g., in XES format [28]) for every cell in the process cube view. Figure 5 highlights one of the 18 cells: This cell contains all events corresponding to a “normal” patient having a unknown age, and performed by a nurse. The events in this cell can be transformed into an event log and analyzed using process mining techniques. Moreover, results for different cells can be compared.

Using the above definition we can formalize notions such as slice, dice, roll-up, and drill-down for process cubes.

**Definition 6 (Slice, [3]).** Let  $PCS = (D, type, hier)$  be a process cube structure and  $PCV = (D_{sel}, sel)$  a view of  $PCS$ . For any  $d \in D_{sel}$  and  $V \in sel(d)$ :  $slice_{d,V}(PCV) = (D'_{sel}, sel')$  with  $D'_{sel} = D_{sel} \setminus \{d\}$ ,  $sel'(d) = \{V\}$ , and  $sel'(d') = sel(d')$  for  $d' \in D \setminus \{d\}$ .

Through slicing, a dimension  $d$  is removed. At the same time one value set  $V$  is chosen for the removed dimension.

For example, Fig. 6 shows  $slice_{staff, \{Peter, Sue, \dots\}}(PCV)$ , starting from the view in Fig. 5. This new process cube view is the result of slicing using dimension




Fig. 6. The process cube view after slicing based on dimension *staff*.

*staff* and value set  $\{Peter, Sue, \dots\}$ , i.e., the *staff* dimension is no longer selected and only events performed by nurses are considered in the resulting view.

**Definition 7 (Dice, [3]).** Let  $PCS = (D, type, hier)$  be a process cube structure and  $PCV = (D_{sel}, sel)$  a view of  $PCS$ . Let  $res \in D_{sel} \neq \mathcal{U}_H$  be a restriction such for any  $d \in dom(res)$ :  $res(d) \subseteq sel(d)$ .  $dice_{res}(PCV) = (D_{sel}, sel')$  with  $sel'(d) = res(d)$  for  $d \in dom(res)$  and  $sel'(d) = sel(d)$  for  $d \in D \setminus dom(res)$ .

Dicing does not remove a dimension, but restricts the values sets for one or more dimensions.





Fig. 7. The result after dicing based on dimensions *staff* and *age*.

Consider Fig. 5 again. Suppose that we would like to remove events referring to patients whose age is unknown and that we only consider events performed

by nurses and doctors. Hence,  $dom(res) = \{staff, age\}$  because these are the two dimensions we would like to dice. Moreover,  $sel(age) = \{\{0, 1, 2, \dots, 18\}, \{19, 20, \dots, 140\}\}$  (note that  $\{?\}$  was removed) and  $sel(staff) = \{\{Peter, Sue, \dots\}, \{Ellen, Tom, \dots\}\}$  (note that the surgeon role was removed). The result is shown in Fig. 7.

**Definition 8 (Change Granularity, [3]).** Let  $PCS = (D, type, hier)$  be a process cube structure and  $PCV = (D_{sel}, sel)$  a view of  $PCS$ . Let  $d \in D_{sel}$  and  $H \in \mathcal{U}_H$  such that:  $H \subseteq hier(d)$  and  $\bigcup H = \bigcup sel(d)$ .  $chgr_{d,H}(PCV) = (D_{sel}, sel')$  with  $sel'(d) = H$ , and  $sel'(d') = sel(d')$  for  $d' \in D \setminus \{d\}$ .

Drilling down is an example of an operator for changing the granularity of the cube, e.g., refining a year into months. However, it is also possible to make the view more coarse (roll-up).



**Fig. 8.** The process cube view after changing the granularity of dimensions *type* and *staff*. The *type* dimension was rolled up (made coarser) and the *staff* dimension was refined (drill-down to the level of individual staff members).

Figure 8 shows a process cube view that was obtained by changing granularity of the original view depicted in Fig. 5. Function  $chgr_{d,H}$  was applied twice: the *type* dimension was coarsened and *staff* dimension was refined. Events related to silver and gold patients have been merged, but still the *type* dimension is visible. Moreover, events are now related to individual staff members rather than roles. Figure 8 now uses the leaves of the *staff* hierarchy.

Through slicing (cf. Fig. 6), dicing (cf. Fig. 7), and changing the granularity (cf. Fig. 8), we can change the process cube view in Fig. 5. At any point in time

we can generate an event log per cell and compare the process mining results. To be able to apply process mining per cell, the classical requirements need to be satisfied, i.e., events need to be (partially) ordered (e.g., based on some timestamp), one needs to select a case identifier to correlate events and an event classifier to determine the activities. See [1, 3, 28] for more information on process mining, process cubes, and event logs.

Based on the definitions in this section we have developed an initial prototype (called *ProCube*) using the process mining framework *ProM* and the *Palo OLAP* toolset [31]. *ProCube* runs as a plugin in *ProM*. The plug-in creates sublogs per cell on-the-fly and visualizes process models discovered using the fuzzy miner [26] and the heuristic miner [45], social networks derived using *ProM*'s social network miner [7], and dotted charts [42] computed per cell. The prototype has many limitations (too slow for high-dimensional process cubes, poor visualization of the results, and limited support for hierarchies), but nicely illustrates the concepts. Currently, we are working on more mature process cube support through software.

## 4 Video Lectures: A Case Study

The primary goal of the process cube notion is to facilitate comparison. To illustrate this, we return to the data set described in Sect. 2. At Eindhoven University of Technology various lectures are recorded and made available online. The *Lecture Capturing System* (LCS) used is *Mediasite* developed by SonicFoundry. This system is able to provide an audit trail, but does not provide any form of process mining. We also used a database with exam results to relate student performance to viewing behavior. Student names were replaced by anonymous identifiers before analysis. Our analysis builds on the PhD research of Pierre Gorissen who analyzed the viewing behavior of students using Mediasite by means of more traditional methods rather than process mining [25].

### 4.1 Data Available on Video Lectures and Exams

To understand the data available consider the class diagram shown in Fig. 9. A *course* has a unique code and a name. The same course may be given multiple times, e.g., once or twice per year. Such a *course instance* has a start date and an end date, in-between these dates *video lectures* are recorded. Per course there are *exams*. Each exam refers to the *last* course instance given. Per course instance there are one or more corresponding exams (typically a regular exam and an additional exam for the students that failed). A *student* may use multiple *exam attempts*. Such an attempt refers to an exam on a particular day and a student. Students may view lectures, also of earlier course instances. An *atomic view* refers to one student and one lecture. Per atomic view the interval watched is recorded, e.g., a student watches the first 14 min of “2II05 College 11b”. Also the time and date at which the student watches the fragment is recorded.

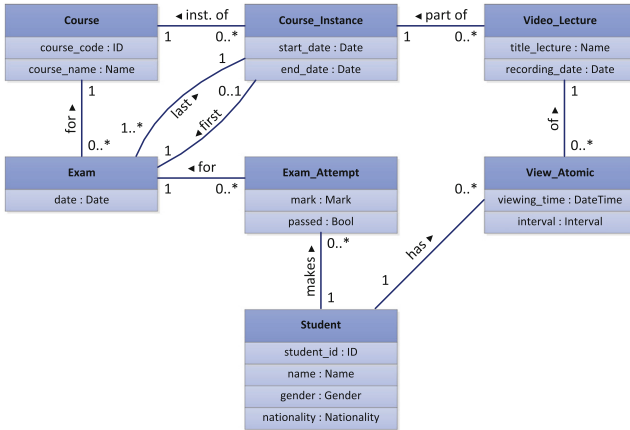


Fig. 9. Overview of the raw data available.

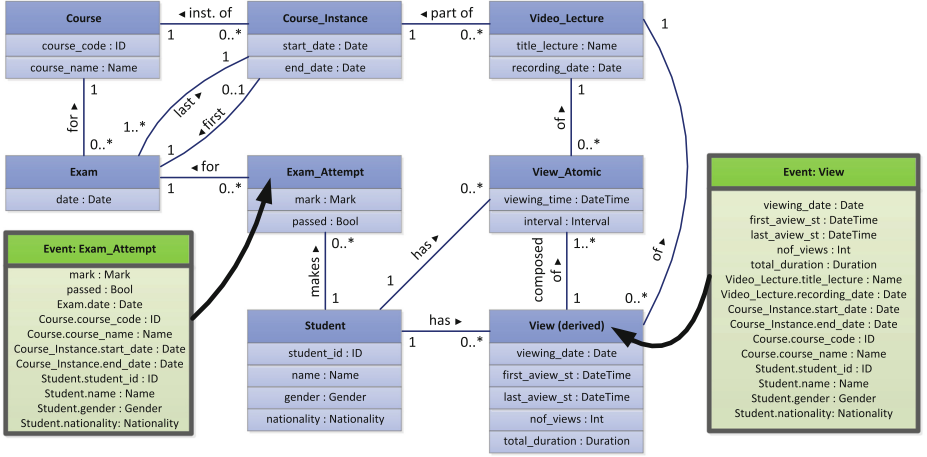
### 4.2 Identifying Events

Students can view a lecture in smaller chunks, e.g., a student can fast forward 20 min, then watch 5 min, go back to the beginning and watch 10 min, etc. Therefore, the same student may generate hundreds of atomic views of the same lecture on the same day. Since we would like to have a model at the level of lectures, we add a new class named “View (derived)” in Fig. 10. Entities of this class correspond to compositions of atomic views. Attribute “nof\_views” in Fig. 10 refers to the number of atomic events and “total\_duration” is the sum of all the durations of the corresponding atomic events. We also record the start time of the first and last atomic view of the lecture by that student on that day.

Based on the available data we propose two types of events: *exam attempts* (entities of class “Exam\_Attempt”) and *views* (entities of class “View (derived)”). Figure 10 shows all the properties of these events. These properties follow directly from the class model. For example, a view refers to a student and her properties (id, name, gender, and nationality), a view refers to a video lecture and its properties (name and recording date). Because a video lecture belongs to a course instance and therefore also to a course, additional properties can be derived.

### 4.3 Defining the Process Cube Structure and Selecting Views

The two event types in Fig. 10 list many properties. When merging both event types in an event base  $EB = (E, P, \pi)$  we take the union of these. Each of the properties may serve as a dimension in the process cube structure  $PCS = (D, type, hier)$ . If a course is composed of different parts, function *hier* can be used to reflect this. Also there is a natural hierarchy for the time domain (years, months, weeks, days, etc.). It is also possible to group exams in a hierarchical manner: a course instance or course can also be viewed as a collection of exams. This way we obtain a hierarchy consisting of three levels: individual exams,

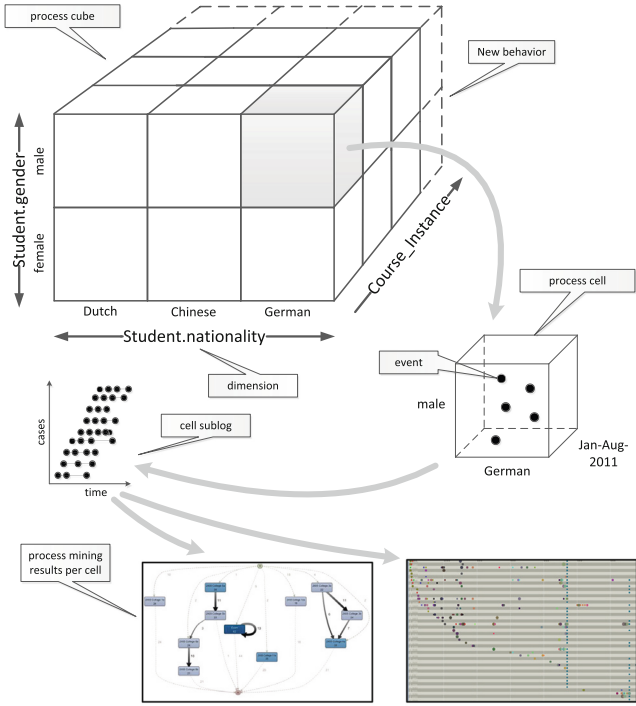


**Fig. 10.** The class diagram with the two types of events considered: *views* and *exam attempts*.

exams belonging to a course instance, and exams for a course. The video lectures can also be grouped in a hierarchical manner.

The process cube view  $PCV = (D_{sel}, sel)$  defines the cells that we would like to consider. We can slice the cube to focus on a particular course. For example we can select only events related to course “2II05” and remove the dimensions “Course.course\_code” and “Course.course\_name” from  $D_{sel}$ . The process cube operators defined in Sect. 3 can be used to create the desired view  $PCV = (D_{sel}, sel)$ . At any time the view  $PCV$  can be materialized resulting in an event log per cell (sublogs). Figure 11 illustrates the overall process. *Once we have an event log per cell, we can apply any process mining technique to all cells and compare the results.* This facilitates comparative process mining.

Recall that for process mining we often assume that  $\{case, activity, time, resource\} \subseteq D$ .  $\pi_{case}(e)$  is the case associated to event  $e$ . There are two obvious choices: a case is an exam attempt or a case is a student. If a student used three exam attempts and we assume the first notion, then the same student will generate three cases. Also other choices are possible, e.g., cases may also correspond to video lectures or course instances. To simplify the interpretation of the results in this paper, we assume that  $\pi_{case}(e)$  equals the “Student.Student\_id” event attribute in Fig. 10. For exam events we choose  $\pi_{activity}(e)$  to be the string “Exam”. For view events we choose  $\pi_{activity}(e)$  to be the title of the lecture, e.g., “2II05 College 11b”. For exam events we choose  $\pi_{time}(e)$  to be the date and time of the exam. For view events we choose  $\pi_{time}(e)$  to be the date and time of viewing. We actually have start and complete events for all activities and can thus measure the duration of activities. Resources seem less relevant here because the case is a student and hence  $\pi_{resource}(e)$  refers to the case itself.



**Fig. 11.** Given a process cube view, the cells can be materialized. Subsequently, process mining techniques can be applied to the cell sublogs.

#### 4.4 Analyzing Process Cube Views: Some Examples

Assume we have “sliced and diced” the process cube in such a way that we only consider the course instance of 2II05 running from January 2011 until August 2011. Moreover,  $D_{sel}$  contains only the dimensions *gender* (male or female) and *nationality* (Dutch or international). This results in four cells. Table 1 shows some results for this process cube view. Dutch students tend to watch the video lectures more frequently and are more likely to pass. Students that pass tend to watch the video lectures more frequently than the ones that do not pass. Note that Table 1 is based on just the 47 students that followed this particular course instance. Hence, based on these findings we cannot (and should not) generalize.

Table 1 also shows the average trace fitness for each of the four cells. Conformance checking based on alignments and a sequential idealized process model are used to compute these numbers (see Sect. 2). The average trace fitness is 1 if all the students watch all video lectures in the default order and conclude with an exam. The Dutch students that passed have the highest average trace fitness (0.39). International students and students that failed have a lower average trace fitness. The average trace fitness for all students that passed is 0.37.


**Table 1.** Some numerical results based on a simple process cube view having only two selected dimensions.

	Dutch student	International student
Student passed	Number of students: 20	Number of students: 2
	Events per student: 28.5	Events per student: 12
	Views per student: 13.1	Views per student: 4
	Exams per student: 1.15	Exams per student: 2
	Trace fitness: 0.39	Trace fitness: 0.25
Student did not pass	Number of students: 17	Number of students: 8
	Events per student: 20.2	Events per student: 17.5
	Views per student: 8.7	Views per student: 7.1
	Exams per student: 1.4	Exams per student: 1.6
	Trace fitness: 0.29	Trace fitness: 0.26

This is significantly higher than the average trace fitness for all students that did not pass (which is 0.28).

Any process mining algorithm can be applied to the materialized cells of a process cube view. Figure 12 shows dotted charts for each of the four cells also used in Table 1. It shows that the Dutch students that passed often took the first exam and passed immediately. They also watched the lectures right from the start. The students that did not pass often skipped the first exam or even made two unsuccessful attempts. It can also be seen that some students systematically watched the videos to prepare for the exam.

For each cell we can also discover process models using process mining techniques. Here, we compare the students that passed (Fig. 13) with the students

**Fig. 12.** Four dotted charts based on the simple process cube view having two dimensions.



that did not pass (Fig. 14), i.e., we only use one dimension for comparison. Figure 13 is based on 22 cases and Fig. 14 is based on 25 cases. Hence, it is not easy to generalize the results. However, there are obvious differences between both process models that confirm our earlier findings. Students that pass, tend to watch the lectures more regularly. For example, students that pass the course tend to start by watching the first lecture (see connection from the start node to “2II05 College 1a”), whereas students that fail tend to start by making the exam (see connection from the start node to “Exam”) rather than watching any video lectures.

Although the small of number students in the selected course instance does not allow for reliable generalizations, the case study nicely illustrates the applicability of process cubes for comparative process mining. Due to space restrictions, we could only present a fraction of the results and touched only a few of the available dimensions. For example, we also compared different courses instances, investigated differences in study behavior between male and female students, etc.

## 5 Requirements and Challenges

The case study just presented nicely illustrates the usefulness of process cubes for comparative process mining. However, the case study also reveals severe limitations of the current approach and implementation.

### 5.1 Performance

The process cube notion is most useful if the users can *interactively* “play with the cube” (slice, dice, drill-down, roll-up, etc.). However, this necessitates a good *performance*. There are two potential performance problems: (1) it takes too long to materialize the cells selected for analysis (i.e., create the event logs) and (2) it takes too long to compute the process mining results for all selected cells. Most process mining algorithms are linear in the number of cases and exponential in the number of activities or average trace length. Hence, it may be worthwhile to precompute results. However, if there are many dimensions each having many possible values, then this is infeasible. In the latter case one also needs to deal with the *sparsity* problem [36]. Suppose that there are just 10 dimensions each having just 10 possible values. Then, there are  $10^{10}$  cells at the lowest level of granularity. This means that even if we have one million of events, at least 99.99% of cells will be empty. This creates an enormous overhead if sparsity is not handled well by the implementation [31].

### 5.2 Interpreting the Results: Comparing Graphs

Another difficulty is the *problematic interpretation* of the results. The goal of showing multiple process mining results is to facilitate comparison. However, this is far from trivial as is illustrated by the four dotted charts in Fig. 12. How to highlight differences and commonalities among multiple dotted charts?

Process mining results are often presented as *graphs*. In earlier sections we presented a few discovered process models showing only the control-flow. Most discovery algorithms return process models that are graphs (Petri nets, BPMN models, transition systems, Markov chains, etc.). Moreover, many other types of mining algorithms create graphs, e.g., organizational graphs, social networks, or richer process maps also showing data flow and work distribution [1]. The layout of such graphs is typically not tailored for comparison. See for example Fig. 13 where the “Exam” activity appears in the upper half of the diagram and Fig. 14 where the same activity appears at the bottom of the process map. To visualize such networks the nodes in the different graphs need to be aligned. Assuming that there is a mapping between the nodes of the graphs, e.g., based on activity labels, there are different ways of visualizing such graph alignments. One can use a “side-by-side” approach where the individual graphs are shown next to each other. To highlight the aligned nodes one can try to give related nodes the same relative position. One can also create a “all-in-one” graph that is the union of all individual graphs and use coloring or annotations to facilitate the reconstruction of the individual graphs. Brasch et al. [15] propose a mixture of the above two approaches (2.5D layouts). Here, the correspondence of aligned nodes is implied by drawing all 2D layouts simultaneously. The third dimension is used to place corresponding nodes on top of each other. The all-in-one approach seems to be most promising for a binary comparison of two models. Here one can take one of the two graphs as a reference and then highlight the differences in a so-called comparison graph. For example, if the arcs in the individual graphs are annotated with durations, then the comparison graph can show the differences highlighting parts that are faster or slower than the reference process.

### 5.3 Refinements

Next to challenges in performance and visualization, also conceptual refinements of the process cube notion are needed. We use the data set presented in Sect. 4 to illustrate these.

As Fig. 10 indicates, we have two very distinct types of events (exam attempts and views) having different properties. By taking the union over these properties, we get many missing values (i.e.,  $\pi_p(e) = \perp$  for property  $p$  and event  $e$ ). For example, an exam attempt event does not have a “nof\_views” property and a view event does not have a “mark”. Moreover, one should make sure that the properties having the same name across different event types actually have the same semantics. For process mining, it makes no sense to only focus on events of one particular type. End-to-end processes may refer to a wide variety of activities. To address this problem one could extend the process cube notion into an *array of process cubes* (one for each type of event) with explicit relations between them. Alternatively, one can also define dimensions to be irrelevant for certain event types. For example, when slicing the process cube using the “mark” dimension (e.g.  $slice_{mark, \{8\}}(PCV)$ ) one would like to retain all view events and only remove exam attempts having a different mark (e.g. not equal to 8). In the current formalization this can be mimicked by including events having a  $\perp$  value when doing a slice or dice action.

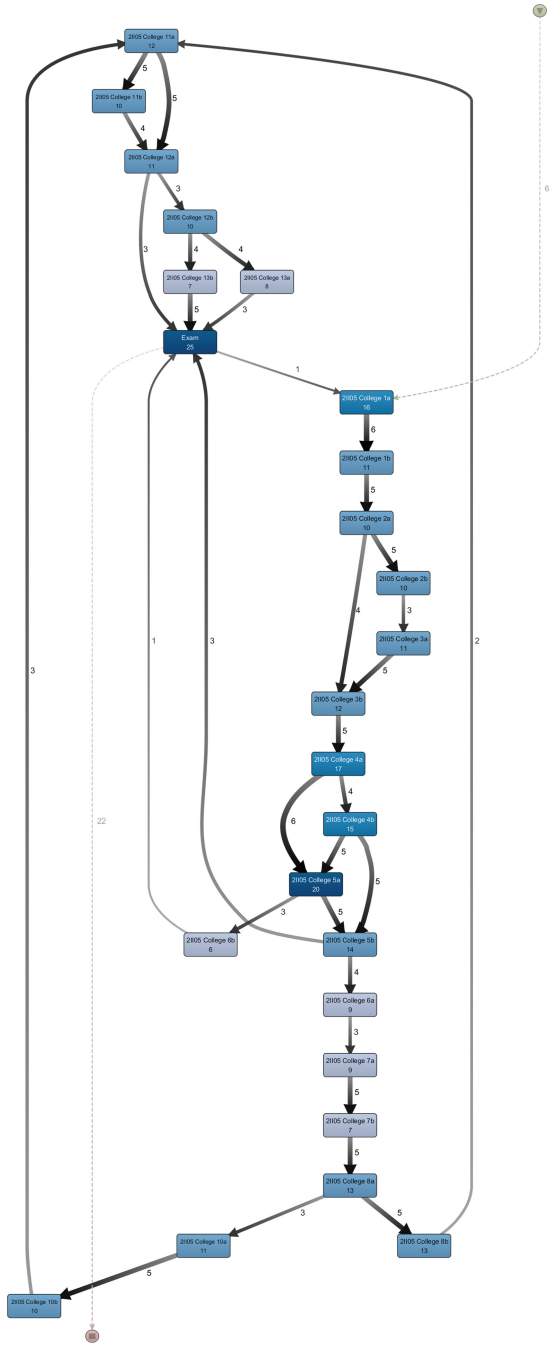


Fig. 13. Process model for the students that passed.

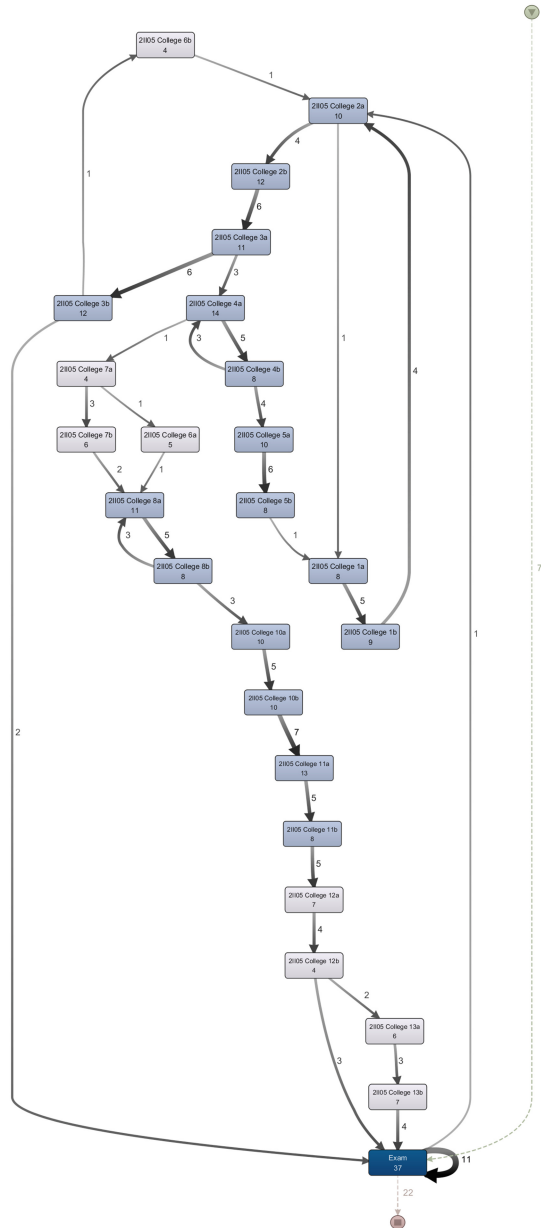


Fig. 14. Process model for the students that did *not* pass.

When applying process mining, events need to refer to a case and an activity, and have a timestamp. Hence,  $\{case, activity, time\} \subseteq D$ . However, these are *not* ordinary dimensions. For example, we may alter the notion of a case during analysis. For instance, using the event data described in Fig. 10, we can first consider students to be the cases and later change the case notion into exam attempts. We can investigate the study progress of a group of students across different courses (i.e., a case should refer to a student). We can also investigate the results for a particular exam (i.e., a case refers to an exam attempt). Similarly, we can change the activity notion. For example, in our analysis we did not distinguish between different exams when naming activities ( $\pi_{activity}(e) = \text{“Exam”}$ ). We could also have opted for different activity names depending on the exam (date and/or course). The process cube should support various case and activity notions and not consider these to be fixed upfront.

The same event may appear in *multiple* cells, e.g., organizational units may partially overlap and subprocesses may have common interface activities [3]. The notion of hierarchy should be as general as described in Sect. 3 to allow for this.

A case is defined as a collection of events. Hence, a case cannot have zero events, because it will simply not appear in the process cube. This may lead to incorrect or misleading interpretations. On the one hand, some process mining algorithms cannot handle empty traces. On the other hand, also empty traces contain information. See for example the importance of empty traces in decomposed process mining [2]. Now all dimensions in *PCV* are event properties. However, *some properties clearly reside at the case level*. For example, a view event in Fig. 10 has as property the gender of the student that viewed the lecture. However, this is more a property of the case (i.e., student) rather than the event (i.e., view). As a result, information is replicated across events. Hence, it may be useful to distinguish between case dimensions and event dimensions in a process cube. Moreover, it is often useful to exploit *derived* case information. For example, we may want to use the number of failed exam attempts or flow time. Generally speaking processes are executed within a particular context [6], and events and cases may need to be enriched with derived context information. For an event  $e$ , we may want to know the utilization level of the resources involved in the execution of  $e$ , and insert this utilization level as another dimension. For a case, we may want to know the relative duration compared to other cases and insert it as a dimension at the case level.

## 6 Related Work

See [1] for an introduction to process mining and the Process Mining Manifesto [27] for the main challenges in process mining. For example, dozens of process discovery [1, 8, 9, 13, 14, 18, 19, 21, 23, 24, 29, 30, 33, 41, 45, 46] and conformance checking [5, 10–12, 16, 22, 24, 34, 35, 39, 44] approaches have been proposed in literature. This paper is not about new process mining techniques, but builds on the notion of process cubes introduced in [3]. Process cubes are related to the well-known OLAP (Online Analytical Processing) cubes [20] and large process model repositories [38].

## 7 Conclusion

As process mining techniques are maturing and more event data become available, we no longer want to restrict analysis to a single process. We would like to compare different variants of the process or different groups of cases. Organizations are interested in *comparative process mining* to understand how processes can be improved. We propose to use *process cubes* as a way to organize event data in a *multi-dimensional data structure tailored towards process mining*.

To illustrate the process cube concept, we used a data set containing partial information about the study behavior of students. When do they view video lectures? What is the effect of viewing these lectures on the exam results? Learning analytics is the broader field that aims to answer such questions. It is defined as “the gathering and analysis of and reporting on data relating to students and their environment for the purpose of gaining a better understanding of and improving education and the environment in which it is provided” [40]. The terms Learning Analytics (LA) and Educational Data Mining (EDM) are used interchangeably and both advocate the intelligent use of event data [17]. Although Pechenizkiy et al. explored the use of process mining in the context of LA and EDM, most of the existing analysis approaches are not process-centric [17, 37, 43]. We consider the comparison of the learning processes inside and between courses as an essential ingredient for a better understanding of the study behavior of students.

The results presented in this paper are only the starting point for a more comprehensive analysis of learning behavior. First of all, we would like to address the foundational challenges described in Sect. 5 (i.e., refining the process cube concept and improving performance in higher-dimensional cubes). Currently, major implementation efforts are ongoing to provide better support for process cubes. We are developing a solution embedded in ProM (work of Shengnan Guo) and a solution based on database technology calling ProM plug-ins from outside ProM (work of Alfredo Bolt). The latter solution is related to calling ProM plug-ins from other tools like RapidMiner [32] and KNIME. Next, we would like to apply the process cube notion to all video lectures recorded at Eindhoven University of Technology and offer such an in-depth analysis as a service for teachers and students. Finally, we would like to apply comparative process mining based on process cubes to new data sets. For example, we will use process cubes to analyze the Massive Open Online Course (MOOC) on “Process Mining: Data science in Action” [4]. Moreover, we envision that more and more data on learning behavior will become available from very different sources. This supports our quest to store event data in such a way that interactive analysis becomes possible. Of course there are also serious privacy concerns. Students should be aware of what is recorded and be able to use it to their advantage. For teachers there is often no need to know the progress of individuals. Often it is sufficient to understand the effectiveness of lectures and teaching material at a group level. Only when students give their consent, one should use the information at the level of individual students.

## References

1. van der Aalst, W.M.P.: *Process Mining: Discovery, Conformance and Enhancement of Business Processes*. Springer, Heidelberg (2011)
2. van der Aalst, W.M.P.: Decomposing Petri nets for process mining: a generic approach. *Distrib. Parallel Databases* **31**(4), 471–507 (2013)
3. van der Aalst, W.M.P.: Process cubes: slicing, dicing, rolling up and drilling down event data for process mining. In: Song, M., Wynn, M.T., Liu, J. (eds.) *AP-BPM 2013*. LNCS, vol. 159, pp. 1–22. Springer, Heidelberg (2013)
4. van der Aalst, W.M.P.: *Process Mining: Data science in Action*. Coursera Course, November 2014. <https://www.coursera.org/course/procmin>
5. van der Aalst, W.M.P., Adriansyah, A., van Dongen, B.: Replaying history on process models for conformance checking and performance analysis. *WIREs Data Min. Knowl. Discov.* **2**(2), 182–192 (2012)
6. van der Aalst, W.M.P., Dustdar, S.: Process mining put into context. *IEEE Internet Comput.* **16**(1), 82–86 (2012)
7. van der Aalst, W.M.P., Reijers, H.A., Song, M.: Discovering social networks from event logs. *Comput. Support. Coop. Work* **14**(6), 549–593 (2005)
8. van der Aalst, W.M.P., Rubin, V., Verbeek, H.M.W., van Dongen, B.F., Kindler, E., Günther, C.W.: Process mining: a two-step approach to balance between underfitting and overfitting. *Softw. Syst. Model.* **9**(1), 87–111 (2010)
9. van der Aalst, W.M.P., Weijters, A.J.M.M., Maruster, L.: Workflow mining: discovering process models from event logs. *IEEE Trans. Knowl. Data Eng.* **16**(9), 1128–1142 (2004)
10. Adriansyah, A., van Dongen, B., van der Aalst, W.M.P.: Conformance checking using cost-based fitness analysis. In: Chi, C.H., Johnson, P. (eds.) *IEEE International Enterprise Computing Conference (EDOC 2011)*, pp. 55–64. IEEE Computer Society (2011)
11. Adriansyah, A., van Dongen, B.F., van der Aalst, W.M.P.: Towards robust conformance checking. In: Muehlen, M., Su, J. (eds.) *BPM 2010 Workshops*. LNBP, vol. 66, pp. 122–133. Springer, Heidelberg (2011)
12. Adriansyah, A., Sidorova, N., van Dongen, B.F.: Cost-based fitness in conformance checking. In: *International Conference on Application of Concurrency to System Design (ACSD 2011)*, pp. 57–66. IEEE Computer Society (2011)
13. Agrawal, R., Gunopulos, D., Leymann, F.: Mining process models from workflow logs. In: Schek, H.-J., Saltor, F., Ramos, I., Alonso, G. (eds.) *EDBT 1998*. LNCS, vol. 1377, pp. 469–483. Springer, Heidelberg (1998)
14. Bergenthum, R., Desel, J., Lorenz, R., Mauser, S.: Process mining based on regions of languages. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) *BPM 2007*. LNCS, vol. 4714, pp. 375–383. Springer, Heidelberg (2007)
15. Brasch, S., Fuellen, G., Linsen, L.: VENLO: interactive visual exploration of aligned biological networks and their evolution. In: Linsen, L., Hagen, H., Hamann, B., Hege, H. (eds.) *Visualization in Medicine and Life Sciences II. Mathematics and Visualization*, pp. 231–249. Springer, Heidelberg (2012)
16. Calders, T., Guenther, C., Pechenizkiy, M., Rozinat, A.: Using minimum description length for process mining. In: *ACM Symposium on Applied Computing (SAC 2009)*, pp. 1451–1455. ACM Press (2009)
17. Calders, T., Pechenizkiy, M.: Introduction to the special section on educational data mining. *SIGKDD Explor. Newslett.* **13**(2), 3–6 (2012)

18. Carmona, J., Cortadella, J.: Process mining meets abstract interpretation. In: Balcázar, J.L., Bonchi, F., Gionis, A., Sebag, M. (eds.) ECML PKDD 2010, Part I. LNCS, vol. 6321, pp. 184–199. Springer, Heidelberg (2010)
19. Carmona, J., Cortadella, J., Kishinevsky, M.: A region-based algorithm for discovering Petri nets from event logs. In: Business Process Management (BPM2008), pp. 358–373 (2008)
20. Chaudhuri, S., Dayal, U.: An overview of data warehousing and OLAP technology. *ACM sigmod Rec.* **26**(1), 65–74 (1997)
21. Cook, J.E., Wolf, A.L.: Discovering models of software processes from event-based data. *ACM Trans. Softw. Eng. Methodol.* **7**(3), 215–249 (1998)
22. Cook, J.E., Wolf, A.L.: Software process validation: quantitatively measuring the correspondence of a process to a model. *ACM Trans. Softw. Eng. Methodol.* **8**(2), 147–176 (1999)
23. Gaaloul, W., Gaaloul, K., Bhiri, S., Haller, A., Hauswirth, M.: Log-based transactional workflow mining. *Distribut. Parallel Databases* **25**(3), 193–240 (2009)
24. Goedertier, S., Martens, D., Vanthienen, J., Baesens, B.: Robust process discovery with artificial negative events. *J. Mach. Learn. Res.* **10**, 1305–1340 (2009)
25. Gorissen, P.: Facilitating the use of recorded lectures: analysing students’ interactions to understand their navigational needs. Ph.D. thesis, Eindhoven University of Technology, June 2013
26. Günther, C.W., van der Aalst, W.M.P.: Fuzzy Mining – adaptive process simplification based on multi-perspective metrics. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 328–343. Springer, Heidelberg (2007)
27. van der Aalst, W., et al.: Process mining manifesto. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) BPM Workshops 2011, Part I. LNBIP, vol. 99, pp. 169–194. Springer, Heidelberg (2012)
28. IEEE Task Force on Process Mining. XES Standard Definition (2013). [www.xes-standard.org](http://www.xes-standard.org)
29. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs - a constructive approach. In: Colom, J.-M., Desel, J. (eds.) PETRI NETS 2013. LNCS, vol. 7927, pp. 311–329. Springer, Heidelberg (2013)
30. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs containing infrequent behaviour. In: Lohmann, N., Song, M., Wohed, P. (eds.) BPM 2013 Workshops. LNBIP, vol. 171, pp. 66–78. Springer, Heidelberg (2014)
31. Mamaliga, T.: Realizing a process cube allowing for the comparison of event data. Master’s thesis, Eindhoven University of Technology, Eindhoven (2013)
32. Mans, R., van der Aalst, W.M.P., Verbeek, E.: Supporting process mining workflows with RapidProM. In: Limonad, L., Weber, B. (eds.) Business Process Management Demo Sessions (BPMD 2014), vol. 1295, CEUR Workshop Proceedings, pp. 56–60 (2014). [CEUR-WS.org](http://CEUR-WS.org)
33. Alves de Medeiros, A.K., Weijters, A.J.M.M., van der Aalst, W.M.P.: Genetic process mining: an experimental evaluation. *Data Min. Knowl. Discov.* **14**(2), 245–304 (2007)
34. Muñoz-Gama, J., Carmona, J.: A fresh look at precision in process conformance. In: Hull, R., Mendling, J., Tai, S. (eds.) BPM 2010. LNCS, vol. 6336, pp. 211–226. Springer, Heidelberg (2010)



35. Munoz-Gama, J., Carmona, J.: Enhancing precision in process conformance: stability, confidence and severity. In: Chawla, N., King, I., Sperduti, A. (eds.) IEEE Symposium on Computational Intelligence and Data Mining (CIDM 2011), pp. 184–191. IEEE, Paris (2011)
36. Naydenova, I., Kaloyanova, K.: Sparsity handling and data explosion in OLAP systems. In: Mediterranean Conference on Information Systems (MCIS 2010), p. 62. AIS Electronic Library (2010)
37. Pechenizkiy, M., Trcka, N., Vasilyeva, E., van der Aalst, W.M.P., De Bra, P.: Process mining online assessment data. In: Romero, C., Ventura, S., Pechenizkiy, M., Baker, R. (eds.) Educational Data Mining (EDM 2009), pp. 279–288 (2009). [www.educationaldatamining.org](http://www.educationaldatamining.org)
38. La Rosa, M., Reijers, H.A., van der Aalst, W.M.P., Dijkman, R.M., Mendling, J., Dumas, M., Garcia-Banuelos, L.: APROMORE: an advanced process model repository. *Expert Syst. Appl.* **38**(6), 7029–7040 (2011)
39. Rozinat, A., van der Aalst, W.M.P.: Conformance checking of processes based on monitoring real behavior. *Inf. Syst.* **33**(1), 64–95 (2008)
40. Siemens, G., Baker, R.: Learning analytics and educational data mining: towards communication and collaboration. In: Proceedings of the International Conference on Learning Analytics and Knowledge (LAK 2012), pp. 252–254. ACM, New York (2012)
41. Solé, M., Carmona, J.: Process mining from a basis of state regions. In: Lilius, J., Penczek, W. (eds.) PETRI NETS 2010. LNCS, vol. 6128, pp. 226–245. Springer, Heidelberg (2010)
42. Song, M., van der Aalst, W.M.P.: Supporting process mining by showing events at a glance. In: Chari, K., Kumar, A. (eds.) Proceedings of 17th Annual Workshop on Information Technologies and Systems (WITS 2007), pp. 139–145. Montreal, Canada (2007)
43. Trcka, N., Pechenizkiy, M., van der Aalst, W.M.P.: Chapter 9: Process mining from educational data. In: Romero, C., Ventura, S., Pechenizkiy, M., Baker, R. (eds.) Handbook of Educational Data Mining, Data Mining and Knowledge Discovery Series, pp. 123–142. Taylor and Francis (2010)
44. De Weerd, J., De Backer, M., Vanthienen, J., Baesens, B.: A robust F-measure for evaluating discovered process models. In: Chawla, N., King, I., Sperduti, A. (eds.) IEEE Symposium on Computational Intelligence and Data Mining (CIDM 2011), pp. 148–155. IEEE, Paris (2011)
45. Weijters, A.J.M.M., van der Aalst, W.M.P.: Rediscovering workflow models from event-based data using little thumb. *Integr. Comput.-Aided Eng.* **10**(2), 151–162 (2003)
46. van der Werf, J.M.E.M., van Dongen, B.F., Hurkens, C.A.J., Serebrenik, A.: Process discovery using integer linear programming. *fundamenta Informaticae* **94**, 387–412 (2010)