Effect of Attribute Alignment on Action Sequence Variability: Evidence from Electronic Medical Records[1]

Inkyu Kim[1], Brian T. Pentland[1], Julie Ryan Wolf[2], Yunna Xie[3], Kenneth Frank[4], Alice P. Pentland[2]

[1]College of Business, Michigan State University, East Lansing, Michigan, USA.
[2]Dermatology, University of Rochester, Rochester, New York, USA.
[3]Public Health Sciences, University of Rochester, Rochester, New York, USA.
[4]College of Education, Michigan State University, East Lansing, Michigan, USA
pentlan2@msu.edu

Abstract. Business process mining algorithms discover processes from event logs that record sequences of events or actions. Typical event logs may or may not contain information about the *attributes* of the actions, such as the particular workstations used to carry out an action or the identity of the person performing the action. In this paper, we test the effect of action attributes on action sequence using data from electronic medical records at five dermatology clinics. We demonstrate that action sequence is influenced by attributes such as actors (who does what) and workstations (what is done where) that are not typically considered relevant to process flow control. We introduce a new metric -- *attribute alignment* -- that summarizes the extent to which actions are carried out with the same attributes throughout a process instance. If each action is always performed with the same attributes, attribute alignment is 100%. We discuss the implications and limitations of this finding for research and practice.

Keywords: Attribute Alignment, Action Sequence, Electronic Medical Records.

1 Introduction

In process mining [1,2,3], processes are discovered from event logs that contain a stream of time-stamped actions or events [1,4]. In a standard XES event log [5], actions may be associated with a set of *attributes*, such as an actor, machine, or location, but

[1] This material is based upon work supported by the National Science Foundation under Grant No. SES- 1734237. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This research was also supported in part by University of Rochester CTSA (UL1 TR002001) from the National Center for Advancing Translational Sciences (NCATS) of the National Institutes of Health (NIH). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We are grateful for comments from Jan Recker and Jan Mendling on an early version of the analysis.

the discovered process is represented in terms of the actions. This makes sense because a process is a coherent, chronological sequence of interdependent events or actions [1,4,6].

In this paper, we investigate the effects of attributes such as actor and location on observed sequences of action in dermatology clinics. Using data from five dermatology clinics at the University of Rochester Medical Center (URMC), we examine the effects of *attribute alignment* on the sequences of action in clinical record keeping. Attribute alignment is a new construct that indicates the extent to which particular actions are consistently performed by the same actor in the same location. Here, we use attribute alignment as an indicator of the extent to which organizational roles (who does what) are consistently defined and carried out. When alignment is low, anyone can do anything, anywhere.

Contrary to our expectations, we find that attribute alignment has a stronger effect on action sequence than the clinic organization or the service performed in the particular visit. We use this finding as a basis for theorizing about how attributes can influence organizational processes.

While the contribution of the findings we report here is primarily theoretical, this research has an important practical motivation: the increasing cost and complexity of healthcare. This paper is part of a three-year research project that seeks to identify the antecedents of complexity in healthcare routines (NSF SES-1734237). Among other things, the research examines managerial factors such as clinical roles and organization. Preliminary results indicate that differences in how clinical roles are defined has a significant effect on process complexity [24]. In particular, clinics with "team documenters" (nursing staff who are responsible for maintaining patient records) have lower complexity than clinics where that responsibility is shared). Here, we dive into the underlying mechanisms that may help explain this phenomenon.

2 Theory

Research on organizational processes and routines naturally tends to focus on actions and patterns of action, because processes are described and defined in this way [6]. The focus is on the actions. In this section, we consider how the business process management incorporates (or excludes) context and attributes.

2.1 Layers of context in business process management (BPM)

In the BPM literature, there is interest in the role of context and in context-aware processes [8-13]. To help sort out the effect of context, Rosemann, Recker, and Flender [10] offer the "onion model", which consists of four layers of context: immediate, internal, external, and environmental.

Layers of the contextual onion tend to have different time scales relative to the cycle time of the focal process. Inner layers vary more quickly, while outer layers vary more slowly. Action attributes from the XES event log, such as the specific actor or workstation associated with an action, can be thought of as part of the immediate context of process execution [10]. Because action attributes (such as actor or workstation) can potentially vary with each action, *immediate* context can vary the fastest. The internal context of a process might include the sequential relationship of actions within a sequence, as in [11]. Because internal context is relative to other parts of a process, it can also vary during each process instance (i.e., during each patient visit). In contrast, external context might vary by weekday/weekend, and environmental context might vary by time of year. External and environmental context tend to remain constant during any single process instance.

There is increasing interest in the analysis and design of context-aware processes [8,11]. The effects of external and environmental context on process execution can be conceptualized and modeled using flow-control variables. For example, the execution of a car rental process may be based on location (airport vs. city pickup), season (winter vs. summer), and other contextual factors. Generally speaking, however, immediate contextual attributes such as who is performing an action, or which workstation is being used, are not considered relevant to flow control. These contextual details may or may not be present in the event log, and may not be included in the process model. The result is an action-only model that conforms the conceptual definition of a process [1,4,6], but leaves out the immediate context of process execution.

2.2 Task Design: Task qua task

Action-only models are entirely consistent with research on task design, where tasks are defined separately from the actor performing the task [14,15,16]. The phrase *task qua task* refers to the abstract idea of the task, separate from the execution of the task [16]. Research in this tradition advocates separating task from context as an explicit methodological principle [16] to avoid conflating properties of the task with properties of the people performing the task.

2.3 Action in context

In contrast to the action-only approach, there are well established research traditions that emphasize the importance of context in the definition and interpretation of actions. For example, the pragmatic force of speech acts [17,18] always depends on context. Expressions such as "here" and "now" mean something different depending on where and when they are uttered. More recently, theories of situated action [19, 20] make a strong argument for the importance of understanding the immediate context of an action. This leads us to expect that action attributes should not be overlooked when analyzing patterns of action.

In summary, there are theoretical reasons to expect that immediate context matters, but established theory in task design and current process mining methods generally do not consider action attributes when describing a process or a task.

3 Investigating the effect of attributes on action sequence

Event log data from an electronic medical record system provides an opportunity to explore the effects of immediate context empirically. We are particularly interested in understanding how action attributes, such as actors and locations, might influence the sequence of action in clinical work. To address this question, we need simple indicators that can be computed and compared using event logs with millions of observations. In the analysis that follows, we operationalize each of these constructs at the level of the process instance (one patient visit to one of the dermatology clinics).

Operationalizing action sequence. The most basic unit of sequential information is the 2-gram. In our usage, 2-grams represent pairs of sequentially adjacent actions in an event log. The number of unique 2-grams in a corpus of sequential data, such as a patient visit, is an indicator of how much sequential variety is present. If there are more unique 2-grams, there is more variety. Note that in principle, the number of unique 2-grams is independent of sequence length, because a sequence could consist of a single 2-gram repeated many times: a, a, a, a, a, a, a. In practice, we expect that longer sequences will have a larger number of unique 2-grams because greater length provides greater opportunity for variation. In the analysis that follows, the dependent variable is the number of unique 2-grams observed in an action sequence.

Operationalizing attributes. To operationalize the role of attributes, we introduce a new construct that we call attribute alignment. It expresses the extent to which attributes add information to the description of an action. If the same actor always does the same action at the same workstation, then attribute alignment = 1. In this idealized case, knowing the action (or the actor, or the location) would give perfect information about the other attributes. The other attributes would be irrelevant.

In contrast, if multiple actors can perform a given action in multiple locations, then attribute alignment is low. Knowing the action does not determine the actor or the location (or vice versa). Attribute alignment provides a single number that encompasses the diversity of attributes associated with each action. The more diversity of attributes observed, the lower the alignment. Attribute alignment can be computed as follows:

$$\textit{Attribute Alignment} = \frac{\textit{Number of unique actions}}{\textit{Number of unique action-attribute n-tuples}} \quad (1)$$

In the data we analyze here, we computed attribute alignment using action, role and workstation. To gain intuition for how this index works, consider a hypothetical example with three actions, three roles and three workstations. In the perfect alignment case, each unique action is performed by a single role at a specific workstation. The attribute alignment would be the maximum (1.0). In the low alignment case, there might be 27 distinct combinations of action-role-workstation (3 x 3 x 3). In that case, the alignment would be 3/27 = 0.11.

Using these constructs, we can state a simple null hypothesis that we test in the following sections. We state this hypothesis in terms of correlation, rather than causality, because the constructs are operationalized within each process instance, so we cannot establish a definitive causal direction.

H_0: Attribute alignment is not correlated with action sequence.

This hypothesis reflects the idea that action sequences should be independent of the immediate context of task performance [8] and the influence of non-control flow variables. Stated in more theoretical terms, it reflects the idea that the *task qua task* exists independently of the actors performing the task and other attributes in the immediate context.

4 Methodology

4.1 Source of data

Data was extracted from the EPIC Electronic Medical Record (EMR) audit trail at the University of Rochester Medical Center (URMC). This data traces actions in the EMR record keeping process. The data included two full years of patient visits from five dermatology clinics (over 7.7 million time-stamped records that provide a trace of actions for 57,836 patient visits, from January 2016 through December 2017). Descriptive features of the data are shown in Table 2 (below).

4.2 Example of data

Table 1 provides an example of the data from the first five minutes of one visit. In addition to the time-stamped action, it contains a number of contextual factors: the role (e.g., admin tech), the workstation, the diagnosis and clinic ID. The role and workstation can be interpreted as immediate context. Note that some actions (e.g., MR_REPORTS) can be performed by any role at any workstation, so the attribute alignment for this visit will be less than perfect. The rows in Table 1 are shaded to show how the immediate contextual factors change throughout a visit, even at the level of individual actions. In contrast, Diagnosis and Clinic ID could be interpreted as external contextual factors. They remain the same throughout the visit.

4.3 Measurement of variables

Unique 2-grams. The dependent variable in our analysis is the number of *unique* pairs of sequentially adjacent actions in a patient visit. To count unique 2-grams, we treated each patient visit as a sequence of actions. We identified 2-grams in each visit using the R package *n-gram*. We then counted the number of 2-grams that are unique. In any given visit, some 2-grams appear more than once, so the number of *unique* 2-grams is always lower than the length of the sequence.

Table 1. Example data

Time	Action	Role	Work-Station	Diagnosis	Clinic ID
2/2/15 8:53	CHECKIN TIME	Admin Tech	W1	Neoplasm	A
2/2/15 8:53	MR_SNAPSHOT	Admin Tech	W1	Neoplasm	A
2/2/15 8:53	MR_REPORTS	Admin Tech	W1	Neoplasm	A
2/2/15 8:53	MR_SNAPSHOT	Admin Tech	W1	Neoplasm	A
2/2/15 8:53	MR_REPORTS	Admin Tech	W1	Neoplasm	A
2/2/15 8:55	MR_SNAPSHOT	Admin Tech	W1	Neoplasm	A
2/2/15 8:55	MR_REPORTS	Admin Tech	W1	Neoplasm	A
2/2/15 8:56	MR_SNAPSHOT	Admin Tech	W1	Neoplasm	A
2/2/15 8:56	MR_REPORTS	Admin Tech	W1	Neoplasm	A
2/2/15 8:56	AC_VISIT_NAVIGATOR	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_HISTORIES	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_ENC_ENCOUNTER	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_VN_VITALS	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_REPORTS	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	FLOWSHEET	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_VN_CHIEF_COMPLAINT	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_REPORTS	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_SNAPSHOT	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:56	MR_REPORTS	Lic.Nurse	W3	Neoplasm	A
2/2/15 8:57	MR_REPORTS	Admin Tech	W1	Neoplasm	A
2/2/15 8:57	MR_SNAPSHOT	Admin Tech	W1	Neoplasm	A
2/2/15 8:58	MR_REPORTS	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	AC_VISIT_NAVIGATOR	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	MR_ENC_ENCOUNTER	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	MR_HISTORIES	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	MR_REPORTS	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	MR_VN_VITALS	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	FLOWSHEET	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	MR_REPORTS	Physician	W4	Neoplasm	A
2/2/15 8:58	MR_VN_VITALS	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	MR_HISTORIES	Lic.Nurse	W2	Neoplasm	A
2/2/15 8:58	MR_HISTORIES	Lic.Nurse	W2	Neoplasm	A
...	...	...	...	...	...

Attribute Alignment. Attribute alignment is the number of unique actions in a visit divided by the number of unique action-role-workstation 3-tuples in the same visit. Each quantity is counted for each visit, so each visit has a value for attribute alignment.

Control variables. We also control for a number of other factors that we expect to influence action sequences in the clinical record-keeping process.

- *Length of sequence.* Visits with more actions are likely to have more unique pairs of action, so we control for visit sequence length.
- *Clinic ID*. We know that each clinic has somewhat different procedures, so we include a dummy variable for Clinic.
- *Level of service.* The Level of Service is a measure of the complexity of the service provided to the patient. It is used for billing and insurance, so it is based on auditable, objective factors about the patient visit.
- *Number of procedures*. This is the number of medical procedures performed during the visit. A typical procedure in a dermatology clinic would be freezing a wart.
- *Number of actions.* Number of distinct actions during the visit. In the data set as a whole, there are 300 possible actions.
- *Number of roles.* Number of distinct actions during the visit. In the data set as a whole, there are 8 roles.
- *Number of workstations.* Number of distinct workstations during the visit. Across all four clinics, there were 118 workstations.

Table 2. Descriptive Statistics

Variables	N	Mean	St. Dev.
Unique 2-grams	57,784	89.27	23.75
Length of Sequence	57,784	133.75	45.00
Level of Service	55,294	3.05	0.38
Number of Actions	57,784	36.70	7.64
Number of Roles	57,784	3.33	1.01
Number of Workstations	57,784	4.30	1.45
Number Procedures	57,784	0.67	1.29
Attribute Alignment	57,784	0.65	0.11

5 Findings

Table 3 shows the results of four regression models. In each model, the number of unique 2-grams is the dependent variable. To correct for heteroskedasticity, we ran our analysis with robust standard error. Due to the large sample size, all of the effects are statistically significant. To facilitate interpretation of the results, we report standardized coefficients and introduce the variables in groups. Note that our findings do not depend on whether the incremental R^2 from one model to the next is significant. Rather, we are interested in the relative size of the effects in the full model, as indicated by the magnitude of the standardized coefficients in model (4). In particular, we are concerned with the effect of attribute alignment, after controlling for everything else.

Model (1) shows the effect of attribute alignment, controlling for the length of the sequence. As expected, the length of the visit sequence is the dominant effect on the number of distinct 2-grams. The length of the visit alone accounts for 88% of the variance (adjusted R^2 = 0.887). Longer visits have many more unique 2-grams than short visits. Together with attribute alignment, the length of the visit accounts for nearly 90% of the variance in the number of unique 2-grams.

Table 3. Number of Unique 2-grams per visit (standardized coefficients)

Variables	(1)	(2)	(3)	(4)
Intercept	0.000 *** (0.376)	0.000 *** (0.415)	0.000 *** (0.480)	0.000 *** (0.403)
Attribute Alignment	-0.101 *** (0.418)	-0.098 *** (0.444)	-0.098 *** (0.458)	-0.171 *** (0.388)
Length of Sequence	0.873 *** (0.001)	0.873 *** (0.001)	0.870 *** (0.001)	0.690 *** (0.001)
Clinic 1		0.011 *** (0.218)	0.014 *** (0.230)	0.006 *** (0.180)
Clinic 2		0.009 *** (0.116)	0.016 *** (0.121)	0.007 *** (0.096)
Clinic 3		0.010 *** (0.099)	0.019 *** (0.103)	0.021 *** (0.082)
Clinic 4		0.012 *** (0.103)	0.020 *** (0.107)	0.012 *** (0.089)
Level of Service			-0.006 *** (0.089)	-0.008 *** (0.069)
# of Procedures			0.008 *** (0.026)	0.001 *** (0.020)
# of Actions				0.282 *** (0.005)
# of Roles				-0.007 *** (0.032)
# of Workstations				-0.070 *** (0.026)
R^2	0.892	0.892	0.892	0.934
$Adjusted\ R^2$	0.892	0.892	0.892	0.934

* *p <0.05, ** p<0.01, *** p <0.001*

Model (2) controls for the effect of clinic organization and work practices by adding dummy variables for each clinic. We know that some clinics had dedicated staff that enter EMR data. In other clinics, residents and physicians do more of the recordkeeping work. As expected, clinic organization has a significant impact on action sequence.

Model (3) controls for the effect of the medical work as indicated by the level of service and the number of procedures. Some clinical visits are simple follow-ups to check if a condition is improving. Other clinical visits involve multiple procedures and tests. These effects are statistically significant, but the magnitude is quite small. This appears to be because the recordkeeping work is not directly proportional to the actual clinical work.

Model (4) controls for the number of actions, roles and workstations observed in each visit, in addition to all of the prior effects. We add these controls to check if the mere number of actions, roles or workstations can account for the effect of attribute alignment. As expected, visits with more actions have more unique pairs of actions. Interestingly, visits with more roles and workstations have slightly fewer unique pairs of actions.

Across all of these models, we find a common result: as the attribute alignment goes down, the number of unique 2-grams increases. In other words, in visits where the attribute alignment is lower, the variation in sequence is higher. We have checked this result in many different ways (adding and removing other control variables, and aggregating the data in various ways). The result is robust. This leads us to reject the null hypothesis that action attributes do not affect action sequences.

6 Discussion

Who does what, and where they do it, has a substantial effect on the sequence of actions in these dermatology clinics. When clinical record keeping is carried out with greater alignment, it has less sequence variety. When alignment is lower, there is more sequence variety. After controlling for sequence length, attribute alignment is the single largest influence on sequence variety. Its effect is larger than clinic organization or the complexity of the work.

This finding is interesting because the dependent variable -- the number of unique pairs of actions -- is based *only on the sequences of actions*, regardless of who performs them or where they are performed. This leads us to suspect that the idealized *task qua task* [21], independent of who and where it is performed, does not exist in these dermatology clinics.

The influence of workstation is interesting because when a user logs in to the EPIC system, the screen is configured for that user. From the point of view of the users, every workstation is identical. Thus, we interpret the workstation as indicating the location of the work (e.g., in the examination room, at the nurses' station in the hall, in the front office, etc.). However, although personalization of user interface makes the workstation digitally identical, the effect of workstation may not be surprising because the physical environment of hospital could determine its influence. A busy hallway is different than a private office. Of course, these contextual differences are not generally conceptualized as relevant to process execution, but our study suggests that they can be.

The implication is that taking a particular action (e.g., *check_meds*) takes on a different meaning depending on who performs it and where it is performed. The office staff can *check_meds* at the workstation in the front office. This might be in response to a patient question (e.g., can I refill this prescription?). This might occur as the patient is checking in or checking out. Alternatively, a nurse, resident or doctor might *check_meds* in the examination room, or outside the examination room, in order to confirm the dosage, look for conflicts, or write a new prescription. The point is obvious once we point it out: When the physician checks the patient's medication, it has a different significance than when the office staff does so. It looks like the same action in the event log, but it is not, because the immediate context is different.

6.1 Why do action attributes influence action sequence?

Intuitively, we did not expect action attributes and immediate context to influence action sequence. We expected that the structure of the work would determine the sequence of actions in the event log. In retrospect, we realize the error in our thinking: the "actions" in the event log are not fully defined by the action code. The logic here is simple. If *check_meds(physican, exam room)* is different than *check_meds(staff, front office),* then the lexicon of actions in the real work is larger than the lexicon of actions in the event log. If the lexicon is larger, the number of unique 2-grams could be larger, as well. By omitting aspects of the immediate context, we are masking valid signals about the nature of the work.

6.2 Including attributes in the definition of actions

Abstracting away contextual details can produce cleaner, more general models, but our findings suggest that this may be a mistake in some cases. Rather than suppressing the immediate context, perhaps we should find ways to include it in our models?

Towards that end, Pentland et al. introduced *ThreadNet*, a simple tool for visualizing and analyzing routines and processes [22,23]. *ThreadNet* is an *R* package that can be downloaded from GitHub (https://github.com/ThreadNet/ThreadNet). *ThreadNet* provides a convenient interface for defining nodes in the graph in terms of any number of attributes. Thus, action attributes become part of the model. When attribute alignment is low, this does tend to result in a proliferation of nodes. However, our results here indicate that the additional complexity may provide empirical insights that an action-only perspective would miss. For example, a process model that includes roles and workstations may help us understand the effect of clinic organization on outcomes such as process complexity and patient satisfaction [24].

6.3 Limitations

This is, in effect, a case study of EMR record-keeping in five dermatology clinics, all operating in the same hospital network. Thus, we should not over-generalize from these findings. In other contexts, the sequential structure of the *task qua task* may be impervious to who is doing the work, where it is performed, or other action attributes.

Using the number of unique 2-grams as the basis for comparison is simple, but two different process instances might have the same number of unique 2-grams. As a result, we believe this metric tends to understate the phenomenon it is intended to measure. It might be more informative to use optimal matching or some other methodology to measure sequential variety [11, 23].

Finally, our contribution at this stage is primarily theoretical. We have shown that attribute alignment can have an unexpected effect on process execution, but we have not yet connected this theoretical insight to practical outcomes, such as cost, quality, or satisfaction.

7 Conclusion

This research demonstrates that, at least in these dermatology clinics, action attributes that are not normally considered relevant to process execution can influence observed sequences of action. To demonstrate this effect, we have introduced a novel measure that we call attribute alignment. We suggest that future research should capture more detailed, event-level contextual information so that the managerial implications of these effects can be investigated more broadly.

References

1. Van Der Aalst, W.: Process mining: discovery, conformance and enhancement of business processes. Springer (2011)
2. Breuker, D., Matzner, M., Delfmann, P., Becker, J.: Comprehensible Predictive Models for Business Processes. MIS Quarterly 40, 1009-1034 (2016)
3. Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A.: Fundamentals of business process management. Springer (2013)
4. Van der Aalst, W., Weijters, T., Maruster, L.: Workflow mining: Discovering process models from event logs. IEEE Transactions on Knowledge and Data Engineering 16, 1128-1142 (2004)
5. Acampora, G., Vitiello, A., Di Stefano, B., van der Aalst, W., Gunther, C., Verbeek, E.: IEEE 1849: The XES Standard: The Second IEEE Standard Sponsored by IEEE Computational Intelligence Society [Society Briefs]. IEEE Computational Intelligence Magazine 12, 4-8 (2017)
6. Van Der Aalst, W., Adriansyah, A., De Medeiros, A.K.A., Arcieri, F., Baier, T., Blickle, T., Bose, J.C., Van Den Brand, P., Brandtjen, R., Buijs, J.: Process mining manifesto. In: International Conference on Business Process Management, pp. 169-194. Springer (2011).
7. Warren, E., Cooper, T.: Generalising the pattern rule for visual growth patterns: Actions that support 8 year olds' thinking. Educational Studies in Mathematics 67, 171-185 (2008)
8. Rosemann, M., Recker, J.C.: Context-aware process design: Exploring the extrinsic drivers for process flexibility. In: The 18th International Conference on Advanced Information Systems Engineering. Proceedings of Workshops and Doctoral Consortium, pp. 149-158. Namur University Press (2006)
9. Günther, C.W., Rinderle-Ma, S., Reichert, M., Van Der Aalst, W.M., Recker, J.: Using process mining to learn from process changes in evolutionary systems. Int'l Journal of Business Process Integration and Management, Special Issue on Business Process Flexibility 3, 61-78 (2008)
10. Rosemann, M., Recker, J.C., Flender, C.: Contextualisation of business processes. International Journal of Business Process Integration and Management 3, 47-60 (2008)
11. Bose, R.J.C., Van der Aalst, W.M.: Context aware trace clustering: Towards improving process mining results. In: Proceedings of the Ninth SIAM International Conference on Data Mining, pp. 401-412. SIAM (2009)

12. Van Der Aalst, W.M., Dustdar, S.: Process mining put into context. IEEE Internet Computing 16, 82-86 (2012)
13. Anastassiu, M., Santoro, F.M., Recker, J., Rosemann, M.: The quest for organizational flexibility: driving changes in business processes through the identification of relevant context. Business Process Management Journal 22, 763-790 (2016)
14. Wood, R.E.: Task complexity: Definition of the construct. Organizational behavior and human decision processes 37, 60-82 (1986)
15. Campbell, D.J.: Task complexity: A review and analysis. Academy of management review 13, 40-52 (1988)
16. Hackman, J.R.: Toward understanding the role of tasks in behavioral research. Acta psychologica 31, 97-128 (1969)
17. Austin, J.L.: How to do things with words. Oxford university press (1975)
18. Searle, J.R., Searle, J.R.: Speech acts: An essay in the philosophy of language. Cambridge university press (1969)
19. Suchman, L.A.: Plans and situated actions: The problem of human-machine communication. Cambridge university press (1987)
20. Feldman, M., Orlikowski, W.: Theorizing practice and practicing theory, organization science. Articles in Advance 1-14 (2011)
21. Hærem, T., Pentland, B.T., Miller, K.D.: Task complexity: Extending a core concept. Academy of Management Review 40, 446-460 (2015)
22. Pentland, B.T., Recker, J., Wyner, G.: Rediscovering handoffs. Academy of Management Discoveries 3, 284-301 (2017)
23. Pentland, B., Recker, J., Kim, I.: Capturing reality in flight? Empirical tools for strong process theory. In: Proceedings of Thirty Eighth Conference on Information Systems, pp.1-12 (2017)
24. Ryan, J. L., Xie, Y., Kim, I., Frank, K., Pentland, A. P., & Pentland, B. T.: Team documentation influences clinic complexity and patient satisfaction. Journal of Investigative Dermatology. (in press)