What is the Problem?
Process Mining with a Control Theory Perspective

Andrea Burattin[0000−0002−0837−0183] and Ekkart Kindler[0000−0003−3895−6297]

Technical University of Denmark,
Kgs. Lyngby, Denmark
{andbur,ekki}@dtu.dk

Abstract. Process mining aims to discover and monitor business processes from event data, traditionally evaluated using criteria such as simplicity, generalization, fitness, and precision. But, these criteria mostly assess how well the discovered model reflects the observed event data; they do not assess how well the discovered model serves its purpose. Looking at business process management and process mining from a control theory point of view, led us to including business goals and key performance indicators into the feedback loop; which gives feedback on how well a discovered process model serves its purpose. In this paper, we propose a framework for process mining and business process management with a control theory perspective. This framework helps us identifying additional quality dimensions on the level of business goals, such as stability, robustness and responsiveness. In addition, this framework identifies the sources of disturbances on a more detailed level than the traditional notions of noise and incompleteness in process mining. This perspective not only enriches process mining evaluation but also offers a deeper understanding of the discipline.

Keywords: Process Mining · Control Theory · Stability · Robustness.

1 Introduction

Process Mining (PM) emerged from the vision of Wil van der Aalst [8], who not only conceived a new way of thinking about processes in combination with corresponding execution data, but also developed the theoretical, methodological, and technological foundations that enabled the discipline to flourish. Over the past few decades, what started as a discipline within Business Process Management (BPM) has matured into a very active area of research, combining conceptual rigor and advanced data-driven methodologies. Its impact has been expanding beyond academia, and Process Mining has achieved remarkable relevance in industry as well.

Process Mining is concerned with deriving process models from the observations of the execution of these processes in some IT system (event traces) and in monitoring such processes based on these observations [7,14]. The first is called process discovery, and the second is called conformance checking.

For assessing the quality of a discovered process model and the underlying process mining algorithm, there exist four well-established quality criteria: simplicity, generalization, fitness, and precision. These criteria determine how well a discovered process model captures the behaviour of the observations and the quality of the discovered model itself.

These well-established quality criteria, however, only indirectly assess the quality of the mining algorithm itself. Therefore, additional analysis includes the behaviour of the mining algorithm in the presence of noise, i.e., incomplete and inaccurate observations in event traces.

Therefore, we asked ourselves how we could capture some quality criteria like stability, robustness, speed, or responsiveness of process mining algorithms, and looked at how some of these notions are defined and formalized in other disciplines like mathematics, numerics, and control theory. And, in particular, control theory as a classical engineering discipline dealing with feedback loops and avoiding instability due to positive feedback loops seemed to provide a good starting point.

Our first attempts at adapting some of the concepts and frameworks of control theory to the area of business process management and process mining were failing. We could not figure out what the concepts of set points, input and output values, and disturbances and errors actually should be in the world of BPM and PM. Only after we added the concepts of *goals* and *key performance indicators* (KPI), things started to fall into place. In addition, goals and KPIs do not only allow us to adopt the control theory framework for process mining; they make sure that the process goals are a target value of the feedback loop, which will keep the process model *fit for its purpose*. Even though goals had been mentioned as a crucial part in BPM for a long time [6,12], they often did not get the attention they deserve. The control theoretic approach made them a necessity.

In this paper, we present a framework for business process management, process-aware information systems, and workflow-management inspired by control theory. And we show how it can help define criteria for the quality of such a system, which include stability, robustness, and performance. In addition, this framework allows us to identify different kinds of disturbances and their sources. Lastly, it emphasizes once again the goals of a process as a crucial part of defining and managing processes, and means to honestly measure whether or to which degree these goals are achieved.

We claim that having this framework in the back of our minds will help us to better understand what we are doing in process mining.

2 Background

To position our control-theoretic perspective on process mining, we first revisit the core ideas from these disciplines. This includes an overview of process mining concepts within BPM, a brief recap of control theory notions, and a discussion on how process mining projects are typically structured in practice. Together, these provide the conceptual basis for the rest of the paper.

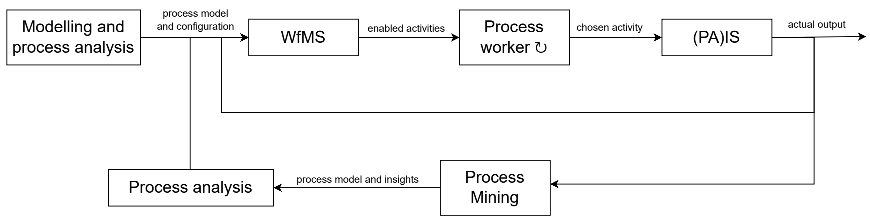

Fig. 1. Business Process Management and Process Mining

Business Process Management and Process Mining. Figure 1 shows the traditional and slightly simplified version of the life cycles of business processes in workflow management systems (WfMS) and process-aware information systems (PAIS) [5,16,9]. The diagram already includes some process mining activities for monitoring the processes and for updating and configuring the processes. The life cycle starts with the creation of a model of a process. This model is then used by a workflow management system for coordinating the work on these processes: showing which activities the agents can choose. The agents or process workers can be human actors or software agents, that pick activities, which, in the traditional view of the Workflow Management Coalition, would be done by the work item handler [6]. The agents (human actors or software agents) would pick and then perform the respective activity by adding and changing data, and eventually finishing the activity. The data would be stored and changed in some information system, but also trigger the workflow system to update the available activities, which the agents can choose from. And this would leave a trace of what has happened in the different processes, which is called the *event stream* or *event log*.

At the bottom of Fig. 1, it is shown how process mining techniques use these event logs for monitoring the process and also for discovering new and better fitting models of the process. These can be used to change the model or the configuration of the processes—either automatically or by human intervention—based on insights obtained by process mining. For this, the event log could either be processed offline as a section of the event log at certain points in time, or the event stream could be processed online while the system is running. We call the latter case “streaming process mining” [1].

Control Theory. For later reference, we briefly recapitulated some of the core concepts of control theory, picked from the introduction of Dorf’s and Bishop’s book on Modern Control Systems [4]. Control theory is about controlling a system, where the system is typically a *dynamic system* with continuous signals and variables. In control systems, there are several systems under consideration, the system that is to be controlled, the actual *controller*, and the overall system consisting of both together, which is called the *control system*. In order to make this distinction clear and to avoid confusion, the system to be controlled is

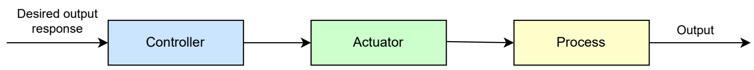

Fig. 2. Open-loop control system (after Fig. 1.2 from [4]). The different blocks are colored to simplify their identification.

often called the *process*, which nicely aligns with business processes management anyway.

Control theory distinguishes between two different types of control systems: *open-loop control systems* and *closed-loop control systems*. Figure 2 shows a block diagram of an open-loop control system. It consists of the process (i.e. the system to be controlled), for example, this could be a car, where the gas pedal or accelerator would be pressed down to a certain level to run the car at a certain speed, called the target or the set speed. Dorf and Bishop call this the "Desired output response" of the process. And there are a controller and an actuator that will affect the process to reach the desired output value. In our car example, this would be pressing the accelerator to a certain point for setting the desired speed; the controller would be all the mechanical (and digital) contraptions that open some valve in the combustor to a certain level to let in the right amount of gas to run the engine at certain rotation per minutes, and this would bring the car to run at a certain speed—hopefully close to the set target speed.

As you might have already guessed, this is not really how it works. Pressing the accelerator of a car to a certain level does not immediately correlate to the car's speed, even after the steady state is reached. There are many other factors that directly or indirectly affect the speed of the car, such as weight, wind, and, most importantly, the slope of the road. We would need to press the accelerator a bit more when the road goes up and a bit less when the road goes down. In general, these other factors are called *disturbances*.

In order to take care of these uncontrollable and unpredictable factors, control theory introduced control systems with a feedback loop or closed-loop control systems, as shown in Fig. 3. In a closed-loop control system, the actual output value (in our example, the speed of the car) is measured by a sensor, and the difference between the set target value and the actual output value (the error) is fed to the controller. This way, the speed can be controlled more accurately in spite of disturbances, noise in the measurement and even unknown factors.

In the real world, there are always disturbances. Therefore, the block diagram for closed-loop systems from Fig. 3, in addition to the feedback loop, shows the points where disturbances and noise might come into the system, or at least which sources of disturbances there are in the overall system.

Closed-loop control systems and the concept of feedback loops are at the core of control theory. They allow to take the actual output value into account for controlling the system, and this way are able to closely reach the desired output value even in the presence of disturbances and noise.

In spite of these advantages, the feedback loop introduces some additional problems: when not carefully designed, it could happen that the output value

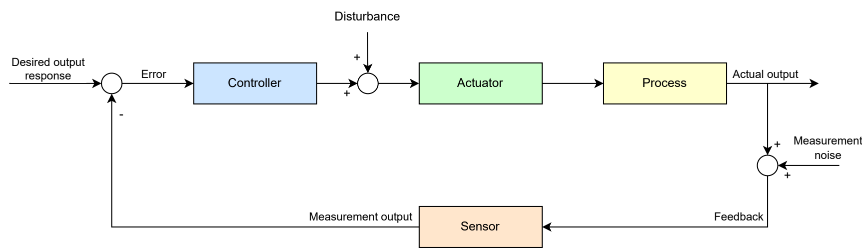

Fig. 3. Closed-loop control system with disturbances and noise (after Fig. 1.4 from [4]). The different blocks are colored to simplify their identification.

of such a system starts oscillating even when the desired output value is not changing and when there are no disturbances and noise at all; and the amplitude of the oscillation might steadily increase over time. The best-known example from daily life is the acoustic feedback loop, which results in howling loudspeakers when a microphone picks up the sound from the loudspeakers of a sound system again and is fed back to the sound system and its loudspeakers. This effect is also known as the Larsen effect.[1] Such systems are called *unstable* in control theory.

Control theory has developed different quality criteria for such closed-loop systems. Most prominently, stability, which would avoid the above-mentioned acoustic feedback, i.e., the Larsen effect. Even though the precise definitions vary depending on the context, in a nutshell, stability means that if the target value is not changed and there are no disturbances, the system will reach an equilibrium again; or if the input value is changed by a finite value, the output value will only change by a finite amount (bounded-input-bounded-output). Control theory provides a whole body of theories for analysing and guaranteeing the stability of closed-loop control systems under certain assumptions. See the book of Dorf and Bishop [4] for examples.

In addition to stability, control theory researchers came up with several other important quality criteria: Concerning performance, there are *accuracy* and *speed*: how accurate and fast is the desired output value obtained; and there is *robustness*, which captures how well the desired properties are achieved even in the presence of other unplanned errors, which include disturbances (that were not made explicit), but also the presence of design errors.

Process Mining Project Methodologies. After a brief introduction to control theory, let's move back to process mining. In practice, process mining projects are typically organized according to established methodologies. A prominent exam-

[1] This effect is named after the Danish scientist Søren Absalon Larsen, professor of electrical engineering at Danmarks Tekniske Højskole, which now is the Technical University of Denmark (DTU). Examples of such effects can be listened to at `https://en.wikipedia.org/wiki/Larsen_effect`.

ple is the PM^2 methodology [14,15], which provides a phased approach covering project planning, data extraction and preparation, analysis and evaluation, and the implementation of improvements. PM^2 emphasizes iterative cycles and the involvement of stakeholders to ensure that process mining results are actionable and aligned with organizational needs.

Similarly, the L* family of life cycles models [13,14] conceptualizes process mining as a recurring set of activities comprising discovery, enhancement, and operational support. These models are closely tied to the business process management life cycle, highlighting the iterative nature of process monitoring and adjustment.

Although these approaches provide valuable guidance on structuring process mining projects, they are largely procedural in nature: they do not explicitly incorporate or foresee aspects such as stability or robustness of the results. For example, PM^2 includes an evaluation step, but it does not really emphasize the stability of the outcomes, but relates them to previous iterations. And there are no feedback loops assessing whether the processes are fit for purpose, meaning whether they are actually achieving their goals.

Recent literature [17] shows that the involvement of human analysts and the choices they make, as well as the interpretative steps, remain a critical aspect of process mining projects. Without explicitly talking about stability and robustness, process mining projects risk yielding insights that are fragile. A control-theoretic approach can complement existing methodologies by providing a structural lexicon to discuss new quality dimensions and feedback concepts. In this way, the strength of structured project management approaches can be combined with a more rigorous system-level evaluation of the outcomes.

Two important phenomena in process mining and the representativeness of event logs are *noise* and *incompleteness* [14, Sect. 6.4.2]: *noise* refers to rare and infrequent behaviour not representative of the typical behaviour of the process, and *incompleteness* refers to event logs with too few events for process discovery. In the control theory setting, these are two particular kinds of disturbances. Our framework enables us to pinpoint this notion of noise, as well as other disturbances.

3 Aligning Process Mining and Control Systems

In this section, we revisit the concepts of business process management and process mining in light of control theory, as shown in Fig. 1 and discussed in Sect. 2. First and foremost, in BPM, the values between the different components are not continuous signals as they would be in control theory, but they are sequences of discrete values; and the values themselves can be composite and more complex than a single signal, such as events with many attributes, and even complete process models or trace alignments. This renders most of the mathematics of dynamic continuous systems of control theory not immediately applicable in our BPM and PM setting.

Figure 1 shows that there are some feedback loops, so we can already guess that we are in a *closed loop* scenario. Some of the performance criteria of control theory, like accuracy, are also used in process mining—even though with a slightly different meaning. In particular, for process mining, we would expect a notion of stability similar to the one in control theory. However, there are only a few papers that mention it or even go into details of a definition in the context of process mining [11,2]. What stands out even more is that the input and output of this "loop" do not match on the type level: the input is the process model—ignoring the configuration information for a moment—and the output is event data in the form of an event log. We could, of course, cut up the loop after the process mining block. If process mining were in the discovery setting, where the output would be another process model, the input and output would be of the same type: process models. In the setting of workflow management systems and process mining, however, this would be, at large, a bit trivial since, by the setup without any disturbances, the output model would always be the same[2] of the input model. Only disturbances would ever lead to a significant change in the process model—even when the process model is far from meeting its goals (cf. "rediscovery problem"). In addition, from the control theory setting of closed-loop systems, we have the feeling that the input and output of the controlled system should be on a more abstract level than a detailed model, like on the level of "business goals" and "key performance indicators" (KPIs). Process models would just be a *way* to accomplish these goals, and KPIs would be a way to measure the extent to which the goals have been achieved by a certain model and its executions (in a WfMS and PAIS). Therefore, we would need to have a closed loop, which involves setting goals and measures to determine the extent to which these goals have been achieved as input and output. It is at this level that the performance, stability, and robustness criteria should be applied.

Most closed-loop control systems, in control theory, make explicit at which point disturbances are coming in, as shown in the block diagram of Fig. 3.

Figure 4 shows an extension of the diagram in Fig. 1 and provides a more coherent and comprehensive picture of how control theory could be used to express where process mining "lives" in an organizational setting. The first extension comprises the modelling of the input of the system: "goals and values". In business process management, goals are typically formulated as text, if at all; and, unfortunately, there are often no metrics for measuring how well a process actually fits its purpose. The goals could, however, be broken down to features and formulated as metrics that capture the extent to which the goal was achieved. These metrics would need to be formulated in a way that allows them to be applied to the event logs and produce a numeric value. Then, the values for these metrics would be the output value of the system (at some point in time); in contrast to classical control flow theory, however, they would not be continuous signals.

[2] Or, when the agents choose to use only a subset of available behaviour, a subset.

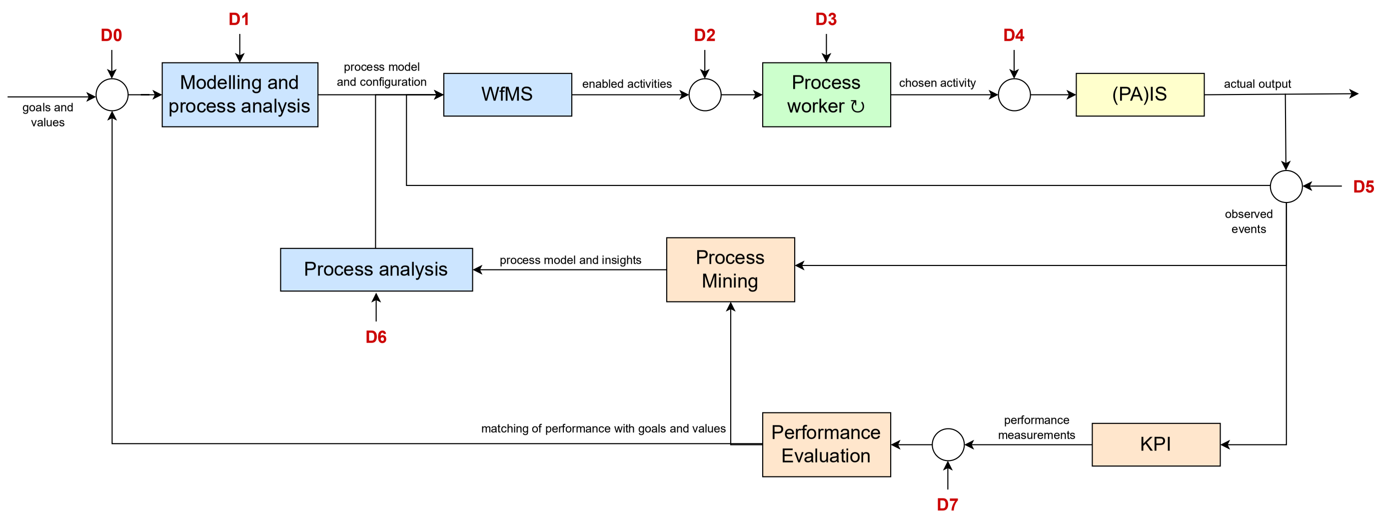

Fig. 4. Conceptual Framework with a Control Theory Angle. The background colours of the different components refer to the type reported in Fig. 3: blue are controllers, green is the actuator, yellow is the process, and orange are sensors.

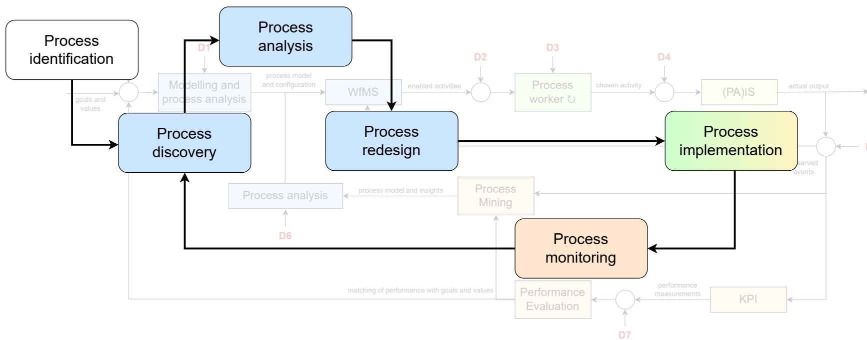

Fig. 5. The BPM life cycle overlaid on top of the proposed conceptual framework. The background colours of the different components refer to the type reported in Fig. 3: blue are controllers, green is the actuator, yellow is the process, and orange are sensors.

Typically, the overall goals would be broken down into some quality features and sub-features, which in turn could be formulated as metrics for measuring these features and aggregating them into the KPIs; these are the basis to formulate desired properties of such systems like stability, robustness, and performance, which will be discussed later in Sect. 5.

Altogether, the explicit formulation of the set goals and corresponding metrics form a closed-loop control system with set values and target values, and a feedback loop to ensure that the KPIs are actually achieved. Actually, Fig. 4 models a "multiloop feedback system" [4] where the outer-most loop verifies the achievement of the goals and values via the KPIs; and the innermost loop is focused on the process level, ensuring that the executions are happening according to the designed process model, regardless of whether this serves its purpose. The picture also highlights disturbances that could occur and disrupt the organization, which are be discussed in detail in Sect. 5. Please also note that the different blocks are coloured according to the block's meanings specified in Fig. 3: blue components refer to controllers, green to the actuator, yellow to the process, and orange to sensors.

It is also relevant to mention that the "feedback loop" aspect of the proposed control system had to some extent already been identified in the literature and is called the "BPM life cycle" [5]. We believe that the framework presented in Fig. 4 can be useful in making the BPM life cycle more tangible and, in particular, in establishing a more explicit link between process mining and such life cycle. Figure 5 presents the BPM life cycle on top of the proposed framework, where the different components of the BPM life cycle have been coloured according to the block meanings specified in Fig. 3.

4 The Need for Metrics in Process Mining

The evaluation of process mining algorithms is a central topic in the field. Traditionally, the quality of a discovered process model is assessed along four well-established dimensions: fitness, precision, generalization, and simplicity [14]. Fitness captures the extent to which the model reproduces the observed event traces, while precision measures whether the model allows for behaviour not present in the log. Generalization aims to avoid overfitting to a specific log by considering unseen but possible behaviour, and simplicity rewards models that are understandable and not overly complex. Together, these criteria provide a balanced view of how well a process model reflects the recorded events and whether it is usable for analysis and decision making.

Despite their widespread acceptance, these metrics capture only certain aspects of model quality. They are primarily focused on the outcome of the process mining and only indirectly assess the underlying algorithms and the more general insights. In addition, these metrics do not directly evaluate how a mining algorithm performs under realistic conditions, such as the presence of noise, incomplete logs, or changes in the process. For example, two algorithms might achieve high fitness and precision, but one may be highly sensitive to small perturbations in the input data, while the other produces more stable results. This difference remains unnoticeable under the traditional evaluation framework.

In addition, the existing evaluation criteria are inherently *static* in nature: they assume the availability of a fixed event log and do not easily extend to scenarios where data arrives in a continuous fashion. However, process mining outcomes, in order to be useful, should have consequences on the actual process being executed, thus creating a dynamic environment. In such dynamic contexts, additional qualities become relevant, such as stability of the mining results over time, robustness to missing or noisy events, and responsiveness to newly arriving information. Without appropriate metrics, these aspects remain unassessed.

Finally, current process mining metrics focus on the internal validity with respect to the *process model* rather than the general *process* and its *goals*. In practice, organizations apply process mining to improve precise metrics, such as process performance, compliance, and customer satisfaction [7,10]. Therefore, process mining evaluations should not only focus on structural model qualities, but also on whether the insights generated are reliable, actionable, and generally aligned with the organizational goals and key performance indicators.

5 Metric, Disturbances and Noise in Process Mining

The control theory framework presented in Sect. 2 provides a well-established vocabulary for analysing the behaviour of dynamic systems with feedback. As discussed, process mining can indeed be seen as such a type of system; therefore, when we transfer control-theory notions to process mining, it is also important to transfer the quality concepts and metrics into this new domain.

Specifically, concepts like *stability*, *robustness*, and *disturbance* offer a valuable perspective for evaluating process mining algorithms beyond the traditional dimensions.

Disturbances. As mentioned previously, in control theory, disturbances represent unexpected external inputs that affect the behaviour of the system [3]. Analogously, disturbances can happen in organizations too, and they can affect the smoothness of operations. Referring to Fig. 4, disturbances can happen essentially in any possible stage of the processing; some of these disturbances are well-known and investigated, others are less explored.

As mentioned at the end of Sect. 2, please remember that there is only a partial overlap between the concept of "disturbance" in control theory and "noise" in process mining. Let's explore the meaning of the disturbances in Fig. 4:

D0: Disturbances in this case can cause a misinterpretation of the goals and the values of the organization, for example, due to a cognitive bias. Examples:
  - In a healthcare scenario, the primary focus of the organization is on the social media presence instead of being on the treatment of the patients.
  - In an educational organization, there is little focus on the actual quality of education, compared to the recreational activities.

D1: Disturbances in the modelling phase cause the goals and the values of an organization to not be properly captured and described in the process model. Examples of such a situation are:
  - An emergency room models "patient admission" without considering a triage, treating patients on a first-come-first-served basis.
  - In a bank, the process model assumes that all loan requests follow a uniform approval flow, but in reality, a VIP customer may skip steps, leading to mismatches.

D2: Disturbances in this case are happening when the process worker does not perceive the correct set of activities to perform from the workflow management system. Examples of such a situation are:
  - A doctor's task list doesn't update after a nurse finishes a prerequisite activity due to system lag.
  - An agent follows a direct supervisor's ad-hoc instruction rather than the recommended protocol.

D3: The actual process worker, in certain scenarios, may decide to overrule the workflow management system and simply proceed in executing a different activity (e.g., knowledge-intensive situations). For example:
  - A doctor, during a medical procedure, makes the conscious decision to deviate from the clinical protocol due to specific circumstances not previously known.
  - Upon delivery of a package at an address, the delivery person realizes that the recipient of a second package, originally scheduled for a different address, is there, and they deliver the second package to the same place.

D4: In this setting, a disturbance is represented by the misalignment between what the process worker executed and what is reported into the (PA)IS. This could happen due to communication issues. Examples:

  - A nurse administers medication but forgets to log it in the system, leading to an incomplete event trace.
  - A worker forgets scanning a barcode on a package, so the system shows the item as “in storage” while it has already been shipped.

D5: Disturbances in this case indicate that the output of the (PA)IS is not observed correctly. Examples of such a situation are:
  - During a system crash, some of the appointment visits in a hospital are not logged.
  - In an online store, order confirmation emails are sent, but the event log fails to record the corresponding “order placed”.

D6: Disturbances in this setting typically refer to the wrong interpretation of process mining results and the effect that these have in implementing resolutions. Examples of such a situation are:
  - A department is flagged as “slow” but, in reality, it handles high-complexity cases (e.g., fraud detection team).
  - An analyst discards results referring to “rare cases” from the model, though they are critical to ensure compliance.

D7: This disturbance, similar to D5, can be caused by a faulty communication channel. This can be due to a delay in communication or just an incomplete view of the situation. Examples:
  - A hospital’s dashboard shows reduced waiting times, but it excludes emergency patients who were logged in a separate system.
  - Customer reports are not yet available from a key region (due to a different time zone), biasing the evaluation.

Please note that the characterization of disturbances described in this section is not complete, and new disturbances may be identified in specific scenarios.

Stability. In control theory, stability refers to the ability of a system to return to a desired equilibrium state after small perturbations. Translated to process mining, stability can be understood as the consistency of discovered models and analysis results under small changes in the input event data. For example, if a few traces are removed, perturbed, or slightly altered, the process mining/evaluation algorithms should still produce models that are structurally and behaviourally close to the original one. When tailoring this concept to the terminology depicted in Fig. 4, stability can actually refer to the results of process mining (i.e., the models being discovered) or, more importantly, to the “outer” feedback loop involving performance evaluation. The latter case refers, essentially, to the stability of the entire organization to disturbances that can occur throughout.

Robustness. Robustness in control theory refers to the ability of a system to maintain certain quality criteria under a variety of operating conditions, including noise and disturbances. Translated to process mining, robustness can be understood as the extent to which discovered models and analytical insights remain meaningful and useful when the data is incomplete, noisy, or, more generally, non-ideal. Unlike stability, which focuses on the sensitivity to small perturbations, robustness is about maintaining reliable analysis results under adverse conditions. Anchoring robustness in Fig. 4 means that an organization is capable of delivering good performance (i.e., resulting from the KPI measurements with respect to the goals) despite challenges in the modelling or in how agents are behaving.

Performance (accuracy and speed/latency). In addition to stability and robustness, our framework can be characterized by at least two performance dimensions: accuracy and speed/latency, similar to those in control theory. Accuracy here refers to how closely the results of the process mining and the general KPIs evaluation approximate the true values (resp. the original process and the goal/values of the organization). Speed or latency, on the other hand, becomes critical in a closed-loop setting as the feedback loop should allow the system to adjust itself in order to reach the desired goals. A process mining system that is highly accurate but suffers from delivering results very slowly may fail to provide timely insights for operational interventions, while a system that is fast but inaccurate risks producing misleading guidance. Therefore, evaluating the performance means to assess the balance of different aspects: algorithms should be benchmarked not only on how correct their results are, but also on how efficiently they can be deployed and how reactive their results can become actionable.

Considering all these aspects, we can conclude that following the analogy with control theory, process mining deployments that are stable (resistant to small perturbations), robust (able to operate under imperfect conditions), and performant (delivering accurate and fast results) will provide more reliable and actionable insights in real-world environments.

6 Conclusion and Outlook

Drawing from the ancient insights often attributed to Heraclitus that "*the only constant is change*" we argue that when deploying process mining techniques, we must consider that execution and behaviour evolve over time. In this paper, we propose a perspective on process mining that is enriched by concepts from control theory. While the traditional evaluation of process discovery has been dominated by quality criteria such as fitness, precision, generalization, and simplicity, our idea emphasizes additional qualities such as stability and robustness. By framing process mining as a feedback closed-loop system, we highlighted the importance of explicitly considering disturbances, their sources, and their impact on the reliability of results—including fitness for purpose. The alignment with control theory does not replace existing evaluation criteria but complements them, offering a more dynamic lens that is better suited for modern process mining applications where data arrives continuously and organizations evolve rapidly.

Looking ahead, we see several avenues for future research. First, the proposed mapping of stability, robustness, and performance needs to be formalized into metrics that can be systematically applied across different algorithms and scenarios. Second, empirical validation is required to test whether these metrics capture meaningful differences in practice and whether they correlate with business outcomes. Third, integrating these concepts into existing process mining methodologies (e.g., PM^2, L*) could help practitioners to pay adequate attention to performance, stability, and robustness of the results. Finally, making the disturbances more concrete can help in devising techniques and algorithms that will better cope with them.

By embracing these directions, process mining research and practice can move closer to an engineering discipline that is not only about understanding process models, but also about actively steering them towards organizational goals, thus properly answering the question: "*What is the problem?*"

References

1. Andrea Burattin. Streaming Process Mining. In Wil M.P. van der Aalst and Josep Carmona, editors, *Process Mining Handbook*, pages 349–372. Springer, 2022.
2. Pieter De Koninck and Jochen De Weerdt. A stability assessment framework for process discovery techniques. In *Proc. of BPM*, pages 57–72. Springer, 2016.
3. Richard C. Dorf and Robert H. Bishop. *Modern Control Systems.* Pearson, 14th edition edition, 2022.
4. Richard D. Dorf and Robert H. Bishop. *Modern Control Systems.* Pearson, 12^{th} edition, 2011.
5. Marlon Dumas, Marcello La Rosa, Jan Mendling, and Hajo A. Reijers. *Fundamentals of Business Process Management.* Springer, 2 edition, 2018.
6. David Hollingsworth. The workflow reference model. Technical Report TC00-1003, The Workflow Management Coalition (WfMC), January 1995.
7. IEEE Task Force on Process Mining. Process Mining Manifesto. In *Proc. of BPM Workshops*, pages 169–194. Springer, 2011.
8. Laura Maruster, Antal van den Bosch, Ton A. J. M. M. Weijters, and Wil M.P. van der Aalst. Process mining: discovering direct successors in process logs. In *Discovery Science*, pages 364–373. Springer, 2002.
9. Manfred Reichert and Barbara Weber. *Enabling Flexibility in Process-Aware Information Systems.* Springer Berlin Heidelberg, 2012.
10. Lars Reinkemeyer. *Process Mining in Action.* Springer, 2020.
11. Anne Rozinat, M Veloso, and Wil MP Van der Aalst. Evaluating the quality of discovered process models. In *conference; ECML/PKDD 2008 workshop on the induction of process models; 2008-09-15; 2008-09-15*, pages 45–52, 2008.
12. August-Wilhelm Scheer. *ARIS — Business Process Modeling.* Springer Science & Business Media, 2000.
13. Wil van der Aalst. Process mining: discovering and improving Spaghetti and Lasagna processes. In *Proc. of CIDM*, pages 1–7. IEEE, 4 2011.
14. Wil M.P. van der Aalst. *Process Mining.* Springer, second edition, 2016.
15. M. L. van Eck, X. Lu, S.J.J. Leemans, and W.M.P. van der Aalst. PM2: A Process Mining Project Methodology. In *Proc. of CAiSE*, pages 297–313, 2015.
16. Mathias Weske. *Business Process Management: Concepts, Languages, Architectures.* Springer Nature, 2024.
17. F. Zerbato, L. Zimmermann, K. Vrotsou, and B. Weber. From analysis to findings: How do process mining analysts discover results? *Information Systems*, 135, 1 2026.