Crawl, walk, run - a look ahead down the path to Process Mining breakthroughs

Prof. Wil van der Aalst
January 9, 2024

diginomica

SUMMARY:

'Crawl, walk, run' goes the saying in business, but that's exactly how meaningful breakthroughs happen. Wil van der Aalst, Chief Scientist at Celonis, looks at Process Mining developments he's most excited to see coming.

There's a line in a Hemingway novel when someone is asked how they went bankrupt. "Gradually, then suddenly" is the answer. I think the same is true of most technological advances – even the ones that seem to come out of nowhere. The truth is that they happen very slowly, over years and decades, and then all at once you see something like ChatGPT appear.

Incremental gains add up to breakthroughs

I see the same pattern of gradual (but meaningful, significant) progress happening in the world of process mining. While the science of process mining has been around since the nineties, and gaining traction in real-world business applications in the past two decades, it's still nowhere near reaching its full potential. There are milestone achievements along the way, some of which I'll touch on, but the real story is one of incremental gains. (And the most important one is the very first step, when you load up the data you have today, and start building your first model.)

However, the new year is often a time of looking ahead and making predictions, and there are a few developments I'm excited to explore – even if they don't pick up significant momentum for a few years.

Looking ahead – knowledge-based process mining and federated process mining

One development is what you might call knowledge-based process mining, where algorithms are essentially injected with human knowledge, for example about specialist domains or completely new situations like COVID. Working together in this way, hand in hand with the model, we create a kind of hybrid intelligence. Recent

milestones like the development of Object-Centric Process Mining (OCPM) have made improvements in Process Intelligence, but human knowledge will still play an important part during this evolution of responsibilities.

One powerful example of this knowledge-based approach is what I call Interactive Process Mining. Traditionally we approach process mining in one of two ways: by mining the data to discover what's happening, or by modeling it out and then conformance checking to see where the deviations lie. With interactive process mining, you will be able to do both at the same time, with haptic feedback to guide you along the way. It could even help you distinguish between the process problems you have the power to fix and those that are very difficult to influence. That's an extremely useful piece of intelligence that is often missing.

When you combine OCPM with a gen AI companion like Celonis Process Copilot, you make it much easier for people to interact with software, and with their own data. The ease and speed this enables will lead to new applications we could never have seen before. It's realistic that we will see gradual steps in this direction in 2024.

Another exciting development on my radar is what's called Federated Process Mining, which makes collaboration easier and more secure across organizational lines when there is difficulty or perhaps reluctance to share too much data. So you might see companies with multiple units that want to learn from one another. Or you might see independent companies that want to collaborate, for example, a car manufacturer and its suppliers.

Stepping stones to domain-specific models

The expansion of process mining to more and more processes, and across organizations, is encouraging, but there's a catch: often the processes we're optimizing are not the business's core function. We might start with order management or duplicate invoice checking – very important jobs to be done, but one step removed from the core service or product.

There are promising exceptions such as Lufthansa, who are applying process mining to the core task of moving people around the world – optimizing processes to reduce delays caused by gate changes, airplane turnaround times, and so on.

My wish is to see more of this, and key to achieving it will be knowledge-based process mining and solutions like the Celonis Process Intelligence Graph (PIG). You can imagine the domains where it could work really well: I would personally like to see more healthcare applications, like patient flows in hospitals. Manufacturing processes are another good one – there is no standard for manufacturing in Germany, but there are certain shared operating principles that exist in engineers' heads and could be captured in data models. Service provision in local government, too.

For common processes like P2P or O2C, we provide pre-cooked solutions. However, these are often not the core processes for an organization. An airline's core

processes have to do with flight management, and a car manufacturer's core processes have to do with production. You cannot expect a pre-cooked solution for such organization- or domain-specific processes. You need humans to inject knowledge, via interactive process mining for example. The PIG can capture this knowledge in such a way that it can be reused by others even when processes are not omnipresent. I think the path to more of these domain-specific models will be to identify what you might call 'the processes in the middle'. Once we see more of these half-generic models fully developed, end to end using OCPM, we will have a stepping stone towards greater domain specialization.

Crawl, walk, run

For all of these technological developments, the old strategy of 'crawl, walk, run' applies – but perhaps it would be more realistic to call it 'crawl, crawl, crawl…walk, run.' Real breakthroughs take time and patience. And turning them into real-world solutions takes even longer. What's important is the direction of travel.

https://diginomica.com/crawl-walk-run-look-ahead-down-path-process-mining-breakthroughs