

Operational Analytics Aren't Really About Analytics

Thomas H. Davenport

Many times when I speak with analytics managers or business people interested in analytics, they tell me that performing some analytics on data is not the primary problem they have. “We have to get the analytics integrated with the process and the systems that support it,” they say. This issue is sometimes called “operational analytics” and it is the most important factors in delivering business value from analytics. It’s also critical to delivering value from cognitive technologies, which in my view are just an extension of analytics anyway.

A quick aside: someone who anticipated this issue early on was Bill Franks, the Chief Analytics Officer of Teradata. He published a book a couple of years ago called *The Analytics Revolution*, which is really about operational analytics. I wrote the foreword to the book, but the meat of the text is where the good advice is about integrating analytics with the core business processes of your organization.

Three things make operational analytics tough in my opinion. One is that to make them work you have to integrate them with transactional or workflow systems. Two is that you often have to pull data from a variety of difficult places. And problem three is that embedding analytics with operational processes means you have to change the behavior of people who perform that process. If you’re successful, you will eventually run into a fourth problem, which is that the embedded analytical models will have to be monitored over time to make sure they remain correct. But since that’s a second-order problem (you should be so lucky to have it), I won’t discuss it further here.

Unfortunately, to succeed with operational analytics a company has to combine transaction systems, workflow systems, analytical systems, databases, and display/user experience tools. Integrating with transactional systems takes a good deal of effort, although modern systems architectures make it a bit easier. Most transactional systems these days (including SAP and Oracle ERP systems) allow API-based connections. But there is usually a fair amount of effort involved in integrating an operational system—sucking out the data you need, doing the analytics somewhere (the cloud, in-database processing), and embedding the result into a transactional interface with the front-line user.

You might be able to accomplish much of the integration with a workflow-oriented overlay tool like case management, business process automation (BPA), or robotic process automation (RPA), although those types of systems don’t generally do any analytics. That means that human labor—your organization’s or some from an external services provider—will be required to combine workflow and analytics. For example, a Boston-based BPA company, Pegasystems (I don’t have a financial relationship with them), partners with professional services firms to combine analytics-based recommendation engines with Pega’s multi-channel marketing automation capabilities.

Problem two is getting all the needed data. That can be handled fairly easily if the data is in an information system and it's in some sort of accessible format. But in many cases, the data is in a variety of formats—from paper reports, PDF files, unstructured articles or medical records, etc. In order to get that kind of data into your operational analytics system, you need more than analytics—you need artificial intelligence.

One of the few vendors that combines AI capabilities with BPA is RAGE Frameworks, another Boston-based company headed by a former professor, Venkat Srinivasan, with a Ph.D. in computational linguistics (I don't have a financial relationship with them either, but I always like it when professors make good). The AI capabilities allow RAGE applications in, for example, financial asset management to extract and classify relevant content from analyst reports and drive investment recommendations. RAGE has also worked with audit firms to extract data from paper and PDF files for account reconciliations. You simply can't automate such processes if you can't automate the "data ingestion" process. In addition, RAGE employs a variety of other "engines"—21 of them, including a computational linguistics engine, a decision tree engine, a business rules engine, and so forth—to rapidly develop intelligent applications. This multiplicity of microservices is the only way I know of to quickly create operational systems that can analyze and think.

Finally, there is the need to persuade front-line users to change their behavior toward decisions and actions based on operational analytics. A "next best offer" system for bank tellers, for example, has to persuade the teller to actually use the recommendations in working with customers. They won't employ analytical recommendations if they don't trust them.

How do you build such trust? Transparency of analytical recommendations is a key factor. If the reason for the recommended product or action can't be described in understandable language, the user won't be able to assess whether it makes sense. That requires some sort of "natural language generation" capability to describe the decision logic. It doesn't favor many machine learning approaches to analytics, because most of the time there is simply no way to describe or interpret why a model prevails in a machine learning process.

What organizations embarking on operational analytics are learning is that analytics themselves are the easy part. There is no shortage of available vendors, both proprietary and open source, of analytical algorithms. But building an operational analytics system means integrating and changing existing architectures and behaviors, and that's always the hard part. It's worth the trouble, however, to build applications in which analytics and smart decision-making are embedded in a company's systems and processes.