

CASE STUDY: TERADATA THINK BIG ANALYTICS AND DANSKE BANK

Danske Bank Fights Fraud with Machine Learning and AI

**Analytics Platform Delivers Double-Digit Reductions
in False Positives**



Doug Henschen
Vice President and Principal Analyst

Copy Editor: Jim Donahue
Layout Editor: Aubrey Coggins

TABLE OF CONTENTS

AT A GLANCE 3

THE COMPANY 4

THE CHALLENGES 5

THE SOLUTION 6

THE IMPACT 10

THE TAKEAWAYS 13

THE RECOMMENDATIONS 14

ANALYST BIO 17

ABOUT CONSTELLATION RESEARCH 18



AT A GLANCE

Danske Bank is a Nordic financial services company that operates in 16 countries, serving more than 1,800 corporations and institutions, 236,000 small and mid-sized companies and 2.7 million personal customers. Starting in 2016, the bank saw an increase in fraud and, with that, an increase in transactions flagged as potentially fraudulent by the bank's rules engine. To reduce number of false positives—meaning legitimate payments flagged for verification—Danske Bank decided to enhance fraud prevention by applying machine learning (ML). The bank is also investigating applying deep learning (DL) to the problem. Executives wished to avoid commercial products and a siloed approach. Their vision was to use open-source technology and to gain expertise that could be applied in many use cases.

Danske Bank

- **Headquarters:** Copenhagen, Denmark
- **2016 Revenue:** \$8 billion (47.96 billion Danish krone)
- **No. of Employees:** 19,700
- **Industry:** Financial institution serving personal, business and corporate customers with banking, wealth management, life insurance, mortgage, real estate and leasing services
- **Other:** Founded October 1871

Business Themes



Data to Decisions



Technology Optimization

At a Glance	
Problems	<ul style="list-style-type: none">• Starting in 2016, Danske Bank saw increases in fraud accompanied by increases in false positives, meaning legitimate payments were flagged for verification.• False positives increased the burden on Danske Bank's fraud team and potentially delayed legitimate transactions.
Solutions	<ul style="list-style-type: none">• The bank sought ML and DL expertise that could address fraud and other business challenges.• Teradata Think Big Analytics offered experience building enterprise-grade ML and DL systems offering human-explainable results.
Benefits	<ul style="list-style-type: none">• ML-based models now in use reduce false positives 20 percent to 30 percent.• DL-based models now in R&D look promising, demonstrating double-digit improvements in fraud detection and further reductions in false positives in a test environment.

Danske Bank was successful with ML proof-of-concept projects, but it lacked experience in bringing these technologies into production. The bank turned to Teradata Think Big Analytics, the big data strategy, consulting and implementation unit of Teradata, which offered expertise and experience with similar high-scale deployments at major financial services companies and other organizations.

Working with Think Big Analytics consultants in 12-week development and operations (DevOps) sprints, Danske Bank's team developed and deployed ML ensemble models that reduced false positives by 20 percent to 30 percent. DL models developed by the team also look promising. Based on TensorFlow, these DL models have yet to be put into production, but they're demonstrating double-digit improvements in fraud detection and further reductions in false positives in a test environment.

THE COMPANY

Danske Bank is a 146-year-old financial institution that provides banking, wealth management, life insurance, mortgage, real estate and leasing services to personal, business, corporate and institutional customers. It operates in the Nordic region, the Baltic states and Ireland, and has offices in 16 countries throughout Europe and extending to Russia, North America and India.

Throughout its history, Danske Bank has embraced cutting-edge technology to deliver better, more-efficient customer service. In 1997, for instance, it introduced electronic banking for personal customers. In 1999 it added mobile e-banking through cellphones. In 2010 the bank added e-banking through smartphone apps, and in 2013 it introduced MobilePay on smartphones.

Danske Bank has embraced advanced, digital technologies internally to increase efficiency of streamline operations, thereby improving customer service. In recent years, the bank has developed what it calls the Danske Bank

Blueprint for Digitization, which provides reference architectures and sets forth explicit standards, instructions and approved tool sets for building information technology assets and components such as application programming interfaces.

A key objective at Danske Bank is to be data driven, according to Nadeem Gulzar, Danske Bank's head of global analytics. To that end, Gulzar's team has extended the Danske Bank Blueprint to analytics, developing a virtualized, private-cloud edge-server configuration that offers high-scale Hadoop and Spark capacity together with proven and hardened analytical tools.

"If a business unit wants to do high-scale analytics, they can order the required services through our internal web shop and they will get an edge server configured with the required resources within 15 minutes," Gulzar explains.

THE CHALLENGES

The volume of online and mobile payments processed by Danske Bank is significant, as this is the predominant form of payment in Europe—in contrast to the U.S., where checks and remittances still dominate.

The key challenge that Danske Bank faced was increasingly frequent and increasingly sophisticated fraud. By 2016, incidents of fraud were on the rise, and with that increase the bank experienced more false positives, meaning transactions flagged for verification that turned out to be legitimate. As the volume of flagged transactions mounted, Danske Bank started investigating ML as a technology that could supplement and boost the performance of its incumbent, rules-based fraud-detection system.

"Fraudsters constantly change their modus operandi, and we constantly need to be better at detecting potential fraud," says Gulzar. "At the same, we want to reduce the need for manual verification by reducing the number of false positives, so we can respond as fast as possible when we detect fraud."

From a business perspective, the bank wanted to avoid fraudulent transactions slipping through and too many legitimate transactions being flagged for manual validation.

Fraudulent transactions that slip through mean financial losses for the customers and the bank. False positives degrade customer experience both by slowing overall processing times and by requiring unnecessary customer interactions to verify the authenticity of payments requests.

THE SOLUTION

Payments aren't the only area where fraud crops up at Danske Bank or, indeed, at any bank. Credit card transaction processing and anti-money-laundering are two other areas that Danske Bank had in mind when considering a solution to the problem in the payment area. Gulzar and his team were interested in an open-source technology solution that could be applied to many areas of the business, and they were specifically considering ML and DL.

“We didn't want to go with a commercial, off-the-shelf solution from a major vendor because what we would get would be a point solution,” says Gulzar. “Our mission was to not only solve the fraud situation and the false positives, but at the same time to build generic components and capabilities that could be used for many other use cases. The only way we saw that could happen was if we built it ourselves.”

Gulzar's team completed proof-of-concept tests using ML in the summer of 2016, and early tests showed that the approach increased detection of fraud and reduced false positives with a performance edge on both counts of 20 percent to 25 percent over the rules engine. These early results confirmed that Danske Bank was on the right track, but Gulzar says the company knew it would need help deploying models in a production environment demanding 300-millisecond response times. At the conclusion of a vendor review process conducted during the summer of 2016, Danske Bank signed a contract with Teradata Think Big Analytics in September of that year.

“We considered many vendors, including both software companies and consultancies, but the main issue with some of the other players was that they relied at least partially on commercial products,” says Gulzar. “We turned to Think Big because they had handled this exact use case at a different bank. We also had talks with them about our thoughts about architecture, and they confirmed that they could work with open-source tools and deliver to our requirements.”

Founded in 2010 and focused on data engineering and data science consulting and support, Think Big Analytics has worked with top-10 banks, major insurers and retailers, tax authorities and health care organizations. The company was acquired by Teradata in 2014. Danske Bank was not a Teradata customer at the time of the engagement. It learned of the consulting services from Think Big Analytics by hearing about prior engagements with banks.

Think Big Analytics works with both open-source technologies and commercial, off-the-shelf products, according to Atif Kureishy, VP of global emerging practices. That choice depends on the client’s preferences. Regardless of the technology choices, the consultancy firm’s focus is on its “Analytic Ops” methodology. Kureishy describes it as a holistic, DevOps-style approach to analytics with a focus on repeatability and automation.

“Our mission was to not only solve the fraud situation... but to build generic components and capabilities that could be used for many other use cases.”

– Nadeem Gulzar,
Danske Bank, Head of
Global Analytics

The Technologies

- Data management platforms include the Hortonworks Hadoop/Spark distribution and Cassandra. Incumbent assets include an IBM mainframe and a Microsoft SQL Server data warehouse.
- Data analysis and model-development tech includes Hive, Spark, Python/PySpark on Anaconda and TensorFlow running on Nvidia GPU clusters.
- Development and deployment management tools include Jenkins, JIRA and Atlassian Bitbucket.
- Model visibility is supported by the open-source LIME (Locally Interpretable Model-agnostic Explanations) framework, which ensured compliance with EU GDPR requirements.

“What’s important is not necessarily the tools but the methodologies and approaches that tie everything together,” says Kureishy, one of lead consultants at Think Big Analytics on the Danske Bank project. “How do you deploy the models and handle inferencing and A/B testing and the entire project lifecycle, including training? That’s what we call Analytic Ops, and it’s what we worked closely with Danske Bank on mastering.”

The Deployment

The project got underway in October 2016. At its peak, the team included a diverse, interdisciplinary group of 40 people. That included more than 30 platform engineers, network engineers, data engineers, data scientists, payment and fraud experts, fraud investigators, and hardware and infrastructure experts from Danske Bank and six to seven consultants from Think Big Analytics. The team worked together in 12-week sprints pursuing two tracks: data science and data engineering.

By late January 2017 the team delivered the first production virtual machines capable

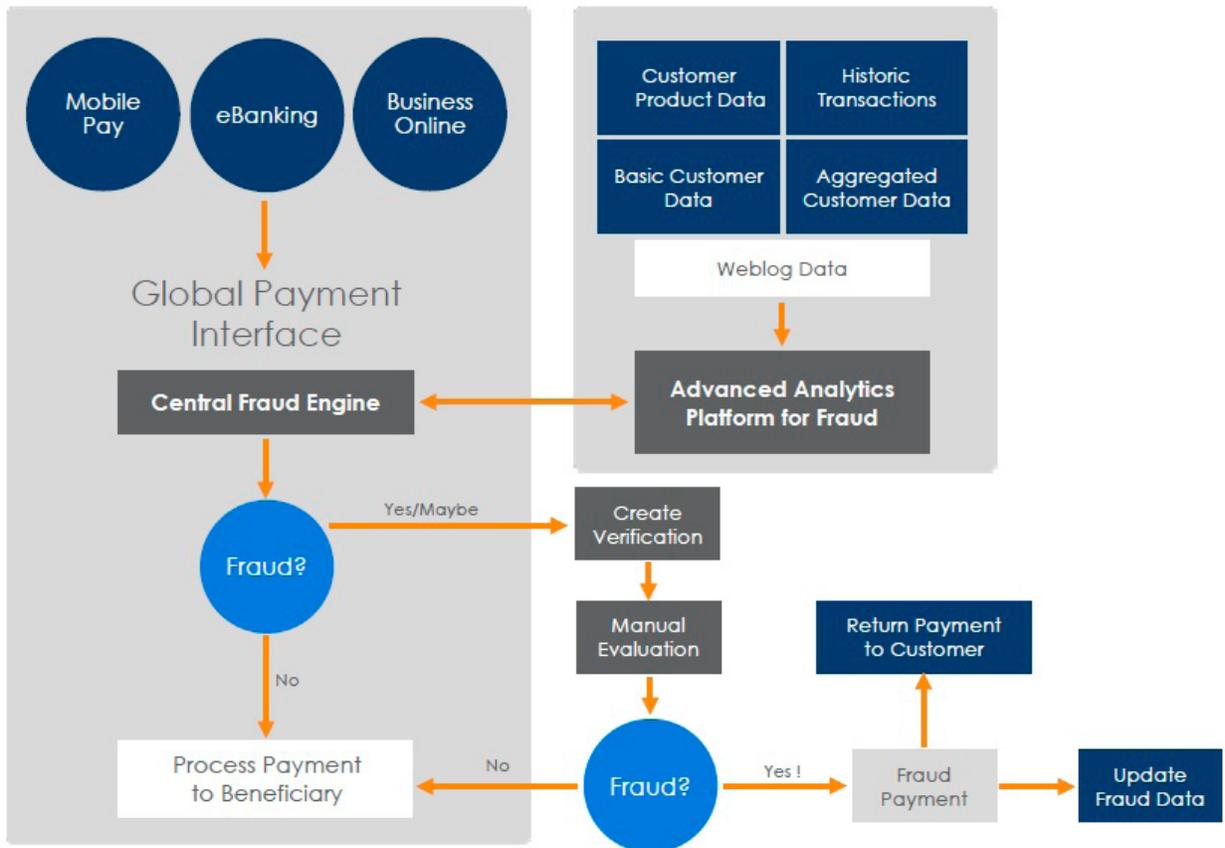
of meeting the data pipeline performance requirements. By mid-February 2017 the first round-trip test transactions were executed.

As shown in Figure 1, the new Advanced Analytics Platform for Fraud that the team developed was designed to work alongside the existing rules engine (a point we’ll return to in “The Takeaways” section on page 13). The deployment also created a closed-loop system in which ML and DL models could quickly learn from the outcomes of manual fraud investigations through updated fraud data. One of the key challenges was ensuring adequate performance.

“We had to ensure a total response time of no more than 300 milliseconds, from the mainframe transaction to data-enrichment to scoring and on to the response, but that meant we had to handle the model scoring within 100 milliseconds,” says Kureishy.

Once performance was proven, the Advanced Analytics Platform for Fraud was run in parallel with the incumbent rules engine from March until June. It was a period of testing, learning

Figure 1. Danske Bank’s New Advanced Analytics Platform for Fraud, Top Right, Complemented the Existing Fraud Engine



Source: Danske Bank

and optimization using live transactional data, but the results of the models were not yet put back into production. Engineers were production-hardening the new platform, including aspects of security and high availability.

Rules tend to be simple and human readable, as in this example: “if [amount] is larger than [X], queue transaction to manual verification.”

Model scores, in contrast, may be based on many data features and interactions among those features. Think Big Analytics and Danske Bank have chosen the LIME (Locally Interpretable Model-agnostic Explanations) framework to help system users understand and explain how scores are determined.

LIME helps explain the classifiers applied to flagged transactions as well as what features

of the data contributed to the prediction. It might be that compared to the customer's usual or recent behavior, if the location of the transaction was significantly different. If one moment transactions are coming in from Denmark and suddenly they're happening in China, or there is some other significant change in behavior, then the transactions showing a change in behavior are likely to be flagged for manual verification.

In mid-June 2017, Danske Bank brought the first wave of ensemble, ML models into production, having been satisfied on two fronts: that the new platform was ready for production performance and reliability and that the models were thoroughly trained and optimized.

The performance results of the ML models are discussed below in the next section, but this was just the first phase of the planned deployment. Following the ML model deployment, the team started work on DL-based modeling using the open-source TensorFlow framework running on high-performance Nvidia graphical processing

unit (GPU) servers. As discussed in the next section, these models have delivered even more promising performance than the ML models in spotting fraud and reducing false positives. For now, the DL research is on hold while Danske Bank focuses on supporting many more transaction types on the ML-based models.

THE IMPACT

Danske Bank was encouraged enough by the 20 percent to 25 percent improvement in performance it experienced in its proof-of-concept tests, but it was even more pleased with the performance of ML-based models put into production in June 2017. With help from Think Big Analytics, the team developed ensemble models combining boosted decision trees and regression. After training and optimization, these models reduced false positives by 20 percent to 30 percent.

Danske Bank declined to divulge the cost of the project or any monetary values tied to its benefits, but it did detail benefits on four fronts:

1. Increased efficiency in fraud investigations.

Danske Bank experienced an immediate 20 percent to 30 percent reduction in fraud false positives, with further reductions anticipated based on DL-based models.

2. **Faster response times.** The platform scores transactions and provides actionable insight within 300 milliseconds. The deployment has solved the first use case, fraud detection, but it's designed to support further solutions for Danske Bank business units.

3. **Improved customer experience.** Between faster performance and reductions in false positives, customers are experiencing improved customer service with fewer unnecessary delays in payments and fewer unnecessary payment-verification contacts.

The Promise of Deep Learning

After developing the Advanced Analytics Platform with ML-based models in June 2017, the team moved on to develop DL-based models using TensorFlow algorithms

running on two Nvidia GPU-based clusters. Experiments with three DL methods—convolutional neural networks (CNN), long short-term memory (LSTM) and auto-encoders—soon delivered double-digit reductions in false positives and further improvements in the identification of fraudulent transactions.

Because it is based on neural nets that mimic the brain and that are capable of self-learning, DL is often described as a form of artificial intelligence (AI). At this writing, the production deployment of DL models is on hold while Danske Bank works on identifying many other types of fraudulent transactions using ML. Though not yet planned, next steps toward putting DL models into production would include integrating the Nvidia GPU-based clusters into a low-latency production data pipeline and tuning and optimizing the DL-based models for 100-millisecond scoring performance.

It's not a question of if but when Danske Bank will need to draw on these more-advanced

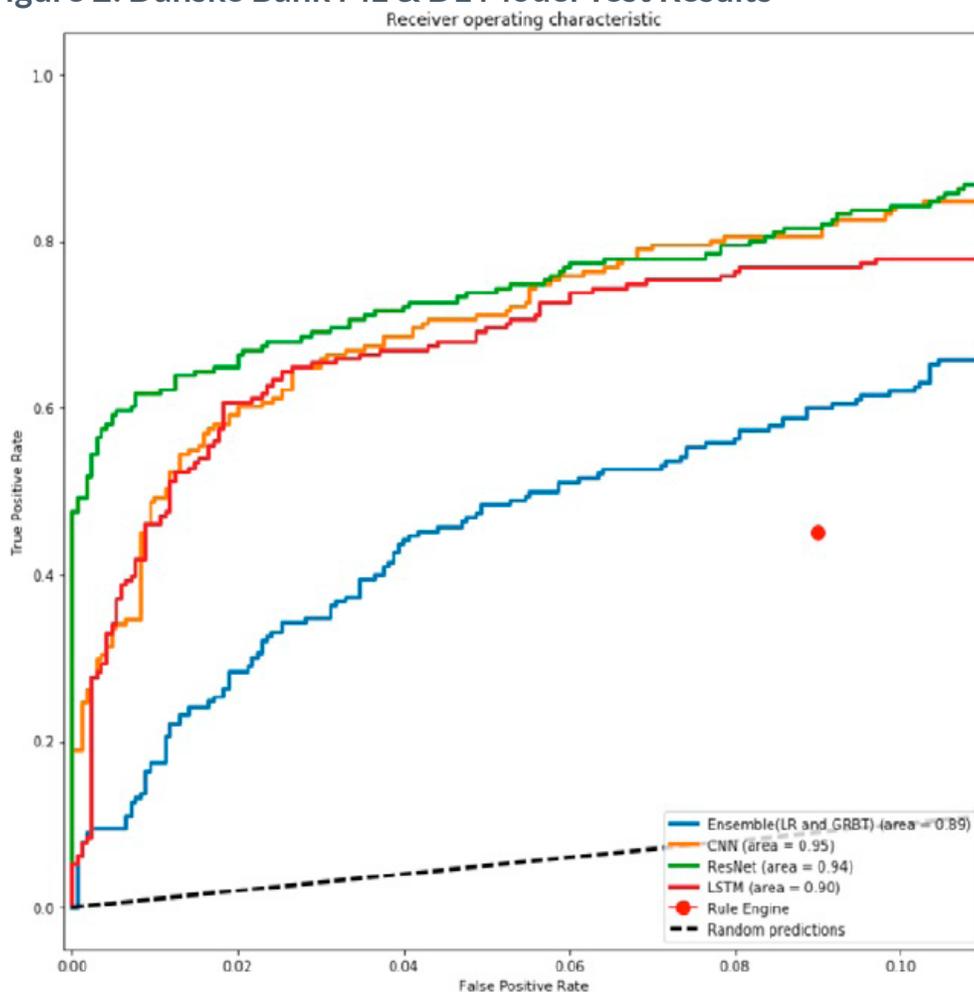
DL capabilities. “With machine learning we’ve achieved very good results, but it’s capped at a certain level,” says Gulzar.

Where ML models look at transactions atomically, for example, DL methods are better able to spot patterns across a series of transactions and to capture correlations

across many features in data. “This is where deep learning will come into the picture, and the early results have been really impressive,” Gulzar says.

As stated in Danske Bank’s winning entry (See Figure 2) in the Data to Decisions category of the Constellation Research 2017 SuperNova

Figure 2. Danske Bank ML & DL Model Test Results



A performance graph comparing random predictions (black dotted line), Danske Bank’s rules engine (red dot), ML-based models now in production (blue), and DL test results, including a CNN model (orange), a residual CNN model (green) and an LSTM model (red).

Source: Danske Bank

Awards, “Danske Bank recognizes the need to be ready with cutting-edge techniques to engage fraudsters not where they are today, but where they will be tomorrow.”

THE TAKEAWAYS

Both Danske Bank and Think Big Analytics shared lessons learned and best-practice advice that’s valuable for any organization looking to apply ML and DL techniques to fraud or other challenges.

- **Tap a range of subject matter experts.**

Danske Bank made a point of engaging a multidisciplinary team that included not just platform engineers, network engineers, data engineers and data scientists but also payment experts, fraud experts and fraud investigators. The fraud experts and investigators should not be brought in just for training on the new system; they can be crucial contributors to understanding fraud patterns and related anomalies in data.

- **Don’t rip and replace.** Danske Bank’s Advanced Analytics Platform for Fraud is

designed to support the fraud detection unit without disrupting or unnecessarily changing existing processes. The approach adds value to the systems already in place and supplements the work carried out by fraud analysts and investigators.

- **Don’t rely on ML or DL alone.** The bank is seeking to complement, rather than replace, its rules-based system. “The fraudsters are quite ambitious, and if they learn that you’re relying on machine learning alone they can launch a series of [non-fraudulent] transactions that would act as a precursor to an attack,” says Gulzar. “They’re training the models to behave in a certain way so that when the fraudulent transaction is triggered, the model would let it pass.” By combining approaches, you reduce the chances of fraudulent transactions slipping through.
- **Embrace open-source technology.** Danske Bank favored open-source technologies over commercial solutions because it sees open-source communities as a magnet for investment and innovation. “If you look

at [Apache] Spark, as an example, IBM, Microsoft and many other big players are putting hundreds of millions of dollars into development,” says Gulzar. “That’s obviously going to benefit [Spark] in terms of capabilities and functionality.”

- **Build flexible capabilities that can be applied to many use cases.** Rather than buying an off-the-shelf solution designed exclusively for payment fraud, Danske Bank wanted to develop infrastructure and skills that could be applied in many fraud scenarios. It’s an approach that Think Big Analytics wholeheartedly supports.

“Once you have the infrastructure and mechanisms in place, you can open up the aperture to different product lines, such as credit cards, loans, checking accounts, and savings accounts, and to other use cases, such as know-your-customer and anti-money-laundering compliance,” says Kureishy of Think Big Analytics. “That’s just on the risk, fraud and compliance side. Next you can start to determine how these factors are impacting

customer experience, customer journeys and customer satisfaction, and you can begin to extend your analytical reach.”

THE RECOMMENDATIONS

- **Determine your build-versus-buy strategy:** Think Big Analytics describes AI and ML as something that should be strategic for most companies. “You have to have skin in the game, which means you need to train models with your own data and learn how to use them,” says Ben MacKenzie, director of AI engineering at Think Big Analytics. “You can’t outsource that to vendors that are going to give you a packaged solution.”

Constellation’s POV: Constellation expects that organizations will increasingly face build-versus-buy decisions as ML- and AI-based products and services offerings proliferate and as adoption grows. Not all aspects of a company can be differentiating, so there will be a place for buying and building. Long-term, defensible differentiation tends to be most closely

tied to unique brand promises and the unique customer experiences that organizations can deliver.

- **Mature from proof-of-concept experimentation to repeatable production execution.** Many organizations have analytics competency centers and do-it-yourself attitudes, but there's a big difference between proof-of-concept projects and hardened production deployments. "Organizations see the power of these techniques, but the reality is that many POCs never see the light of day," says Kureishy. "You have to make an investment to get the talent, the infrastructure and access to the data in place to build a corporate intelligence function."

Constellation's POV: Many organizations talk a good game about being data-driven companies, but industry leaders are forging industrialized and highly automated approaches to building models and data pipelines. You have to mature from disparate, uncoordinated experimentation to adopt consistent standards and repeatable methodologies and processes.

- **Build trust with transparency and training.** ML- and DL-based systems can't be black boxes. That's a matter of regulation in banking, insurance and other industries where certain types of predictive decisions must be explainable and proven to be unbiased. But customers

"You have to invest to get the talent, the infrastructure and access to the data in place to build a corporate intelligence function."

– Atif Kureishy,
Teradata Think Big Analytics

and employees in any industry will naturally more readily trust decisions that are transparent and easily explainable.

and train and prepare employees on new systems and processes.

Some types of algorithms are more transparent than others. Random forest approaches, for example, are inherently interpretable, whereas the neural nets used in DL are more inscrutable. Think Big Analytics execs recommended the LIME framework in the Danske Bank deployment because it was one of the few options available for model-agnostic interpretability at the time. But this is an ongoing area for cutting-edge AI research, they say. The approach that Think Big Analytics took for LIME deployment has satisfied the requirements of auditors and regulatory officers.

Constellation's POV: Transparency, training and trust go together. Don't expect cutting-edge ML and DL deployments to succeed without transparency or training. Whether required by regulation or not, build interpretability into your systems

ANALYST BIO

Doug Henschen

Vice President and Principal Analyst

Doug Henschen is Vice President and Principal Analyst at Constellation Research, Inc., focusing on data-driven decision making. His Data-to-Decisions research examines how organizations employ data analysis to reimagine their business models and gain a deeper understanding of their customers. Data insights also figure into tech optimization and innovation in human-to-machine and machine-to-machine business processes in manufacturing, retailing and services industries.

Henschen's research acknowledges the fact that innovative applications of data analysis require a multi-disciplinary approach, starting with information and orchestration technologies, continuing through business intelligence, data visualization, and analytics, and moving into NoSQL and big data analysis, third-party data enrichment, and decision management technologies. Insight-driven business models and innovations are of interest to the entire C-suite.

Previously, Henschen led analytics, big data, business intelligence, optimization, and smart applications research and news coverage at InformationWeek. His experiences include leadership in analytics, business intelligence, database, data warehousing, and decision-support research and analysis for Intelligent Enterprise. Further, Henschen led business process management and enterprise content management research and analysis at Transform magazine. At DM News, he led the coverage of database marketing and digital marketing trends and news.

 [@DHenschen](https://twitter.com/DHenschen) |  www.constellationr.com/users/doug-henschen

 www.linkedin.com/in/doughenschen



ABOUT CONSTELLATION RESEARCH

Constellation Research is an award-winning, Silicon Valley-based research and advisory firm that helps organizations navigate the challenges of digital disruption through business models transformation and the judicious application of disruptive technologies. Unlike the legacy analyst firms, Constellation Research is disrupting how research is accessed, what topics are covered and how clients can partner with a research firm to achieve success. Over 350 clients have joined from an ecosystem of buyers, partners, solution providers, C-suite, boards of directors and vendor clients. Our mission is to identify, validate and share insights with our clients.

Organizational Highlights

- Named Institute of Industry Analyst Relations (IIAR) New Analyst Firm of the Year in 2011 and #1 Independent Analyst Firm for 2014 and 2015.
- Experienced research team with an average of 25 years of practitioner, management and industry experience.
- Organizers of the Constellation Connected Enterprise—an innovation summit and best practices knowledge-sharing retreat for business leaders.
- Founders of Constellation Executive Network, a membership organization for digital leaders seeking to learn from market leaders and fast followers.



www.ConstellationR.com



[@ConstellationR](https://twitter.com/ConstellationR)



info@ConstellationR.com



sales@ConstellationR.com

Unauthorized reproduction or distribution in whole or in part in any form, including photocopying, faxing, image scanning, e-mailing, digitization, or making available for electronic downloading is prohibited without written permission from Constellation Research, Inc. Prior to photocopying, scanning, and digitizing items for internal or personal use, please contact Constellation Research, Inc. All trade names, trademarks, or registered trademarks are trade names, trademarks, or registered trademarks of their respective owners.

Information contained in this publication has been compiled from sources believed to be reliable, but the accuracy of this information is not guaranteed. Constellation Research, Inc. disclaims all warranties and conditions with regard to the content, express or implied, including warranties of merchantability and fitness for a particular purpose, nor assumes any legal liability for the accuracy, completeness, or usefulness of any information contained herein. Any reference to a commercial product, process, or service does not imply or constitute an endorsement of the same by Constellation Research, Inc.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold or distributed with the understanding that Constellation Research, Inc. is not engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. Constellation Research, Inc. assumes no liability for how this information is used or applied nor makes any express warranties on outcomes. (Modified from the Declaration of Principles jointly adopted by the American Bar Association and a Committee of Publishers and Associations.)

Your trust is important to us, and as such, we believe in being open and transparent about our financial relationships. With our clients' permission, we publish their names on our website.

San Francisco | Belfast | Boston | Colorado Springs | Cupertino | Denver | London | New York | Northern Virginia
Palo Alto | Pune | Sacramento | Santa Monica | Sydney | Toronto | Washington, D.C.

