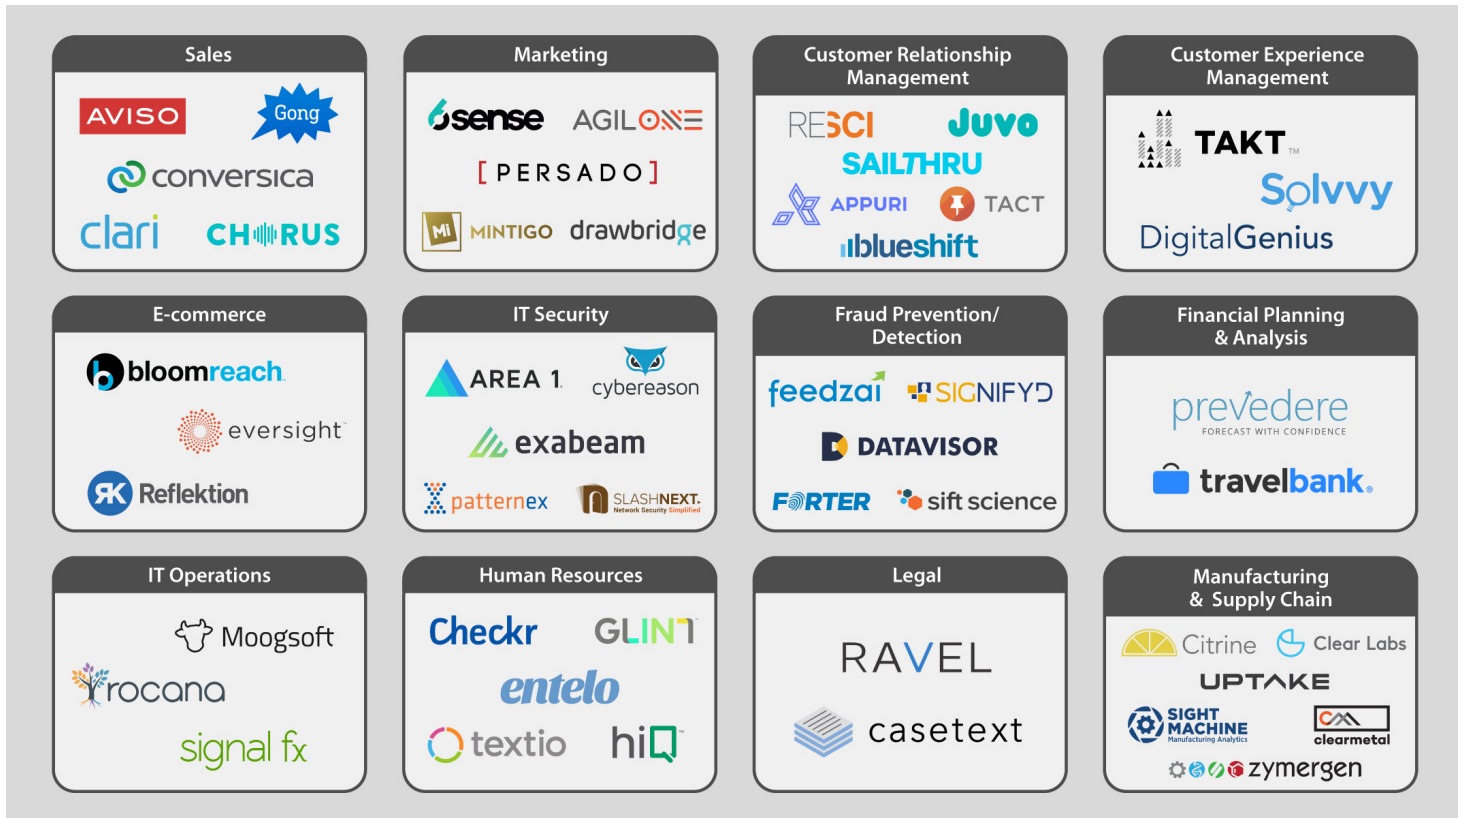


The Wing Data-First 50: AI-Powered Business Applications

MAY 11, 2017 MARTIN GILES DATA



Welcome to our inaugural list of AI-powered Data-First business applications. It highlights 50 innovative, venture-backed startups in America that have set out to reinvent the \$150BN-plus market for business applications using a powerful combination of data-centric strategies and artificial intelligence.

The original business application software paradigm, led by companies such as Oracle and SAP, produced a generation of applications that drove efficiency by mapping workflows and then codifying them. Data was treated as an afterthought. The subsequent SaaS revolution dramatically improved the software delivery and distribution model, but didn't change the underlying approach.

Now we're entering a new era in which an emerging generation of applications is turning this

approach on its head: data lies at their core, and business logic is applied to it rather than the other way around—hence the term “Data-First”. Artificial intelligence technologies embedded in the applications drive actionable insights from the outset.

We have seen this approach play out at scale in consumer markets, where companies like Netflix, Google, and Amazon pioneered it to drive things such as recommendation engines and personalization strategies. Inspired by their success, entrepreneurs are now applying it in business markets too, as we highlighted in [prior work \(http://wing.vc/blog/why-data-first-applications-will-come-to-rule-enterprise-software\)](http://wing.vc/blog/why-data-first-applications-will-come-to-rule-enterprise-software). The Data-First applications they are building may produce an even bigger collective shift in the business tech landscape than the one triggered by the SaaS movement before it. This shift will create many significant and valuable new companies. It’s too early to tell whether those on our list will be amongst the eventual winners, though they are off to an impressive start and have raised a total of over \$1.8BN in venture funding.

We plan to update the Wing 50 List to reflect future advances in the field, so we’d very much like to hear from founders who think their companies should be included on it. (To be considered, please contact us at research@wing.vc) We’ll also be holding an invitation-only event in July in San Francisco at which some founders of Data-First companies will talk about their approach to building next-generation applications. If you’d like to be considered for an invitation, please also use the email address above.

How We Built The Wing Data-First 50 List

To arrive at our list of companies, we began by developing an initial group of over 1500 US-based, venture-backed startups across multiple segments that appeared—from a review of their websites and marketing materials—to be using AI-powered Data-First strategies in their applications. The initial list was generated using various helpful databases, including Dow Jones VentureSource, Pitchbook, and Mattermark.

We then applied a filter, keeping only startups that had raised their first round of funding on or after January 1st 2010, and were in market with their products. This raised the likelihood that those left to analyze were born Data-First, given advancements in data architectures and AI technologies over the past few years—though we recognize that an arbitrary cut-off leaves out some otherwise worthy contenders.

Next, we looked for more detailed evidence that the remaining companies were using data and AI technologies in truly innovative ways. (We excluded AI platforms from the list unless they had a very strong vertical focus.) As we worked through the analysis, it quickly became clear that some categories were seeing a surge in Data-First applications. So these domains have more entries than others. But we were still able to include at least two examples in each of the 12 categories on our list. It includes four Wing portfolio companies: Juvo (Customer Relationship Management); SlashNext (IT Security); Moogsoft (IT Operations); and Clear Labs (Manufacturing & Supply Chain).

The applications that appear on our list exhibit some or all of the following characteristics:

- **Flexible and highly scalable data architectures**

Data-First applications process a wide variety of data types and structures to produce real-time, or near real-time, insights. Sailthru, a customer relationship management application, manages over 1BN unique customer profiles and processes a million emails a minute, while Drawbridge, a marketing application that tracks users across device platforms, can monitor over 80BN ad requests a day.

To cope with this kind of scale, Data-First applications are often built on fast and highly scalable data architectures that use open source processing frameworks like Apache Spark, data lakes like Hadoop and Cassandra, and next-generation data warehouses like Amazon Redshift and Snowflake.

DataVisor, a fraud prevention application, provides an example of this new architecture at work. It deploys its service in AWS, and data is parsed by a security analytics engine running on top of Spark and EC2 instances. The results are then fed into Apache HBase and Elasticsearch, from which they can be accessed by customers via DataVisor's intelligence console. Amazon S3 provides the underlying storage.

- **Embedded AI to recommend actions, predict outcomes, and/or automate responses**

The prior generation of business applications required specially trained analysts to query databases and then recommend actions to line-of-business users. Data-First applications eliminate the need for those jobs by leveraging embedded machine learning and other AI technologies, such as natural language processing and neural networks, to recommend actions or offer predictions to business

users as part of their workflows.

Prevedere, a financial planning application in the Wing 50, applies machine learning algorithms to customers' internal data and a wide range of other datasets to help them generate more accurate forecasts. ClearMetal, a supply chain management application, ingests multiple datasets covering things such as weather patterns, currency rates, and port activity in order to offer predictions about the availability and cost of shipping containers to carry goods around the globe.

- **Data-driving as well as data-driven capabilities**

AI-powered Data-First applications aren't just data-driven; they're "data-driving" too. Engineered to achieve broad deployment, they generate significant data exhaust of their own. These "synthetic" datasets are then used to drive additional algorithms delivering value-added functionality, which in turn generate even more refined data.

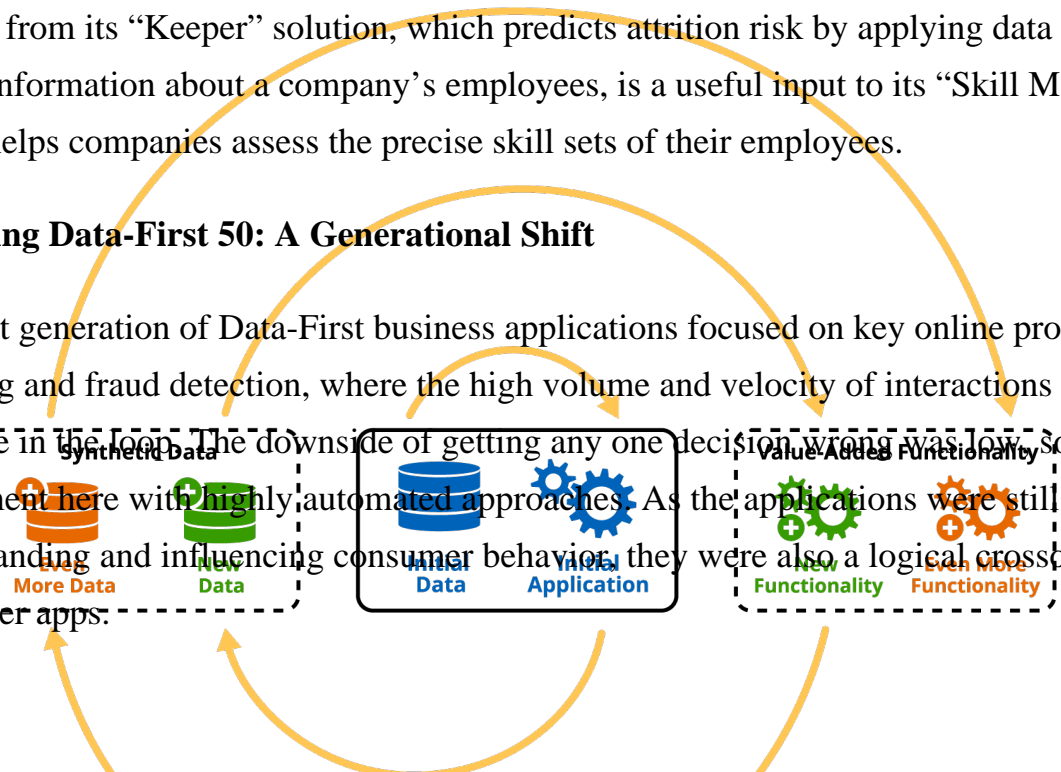
The result is a virtual breeder reactor of business-process optimization shown in the chart below, which gives companies using Data-First applications a significant advantage over rivals using legacy ones. The startups offering these next-generation applications benefit too, as their systems get smarter with each iteration of this virtuous data cycle. It is their proprietary synthetic datasets that give them such a powerful competitive edge.

Driving A Virtuous Data Cycle

HiQ, an HR application in the Wing 50, is a good example of a data-driving product. The data exhaust from its "Keeper" solution, which predicts attrition risk by applying data science models to public information about a company's employees, is a useful input to its "Skill Mapper" solution, which helps companies assess the precise skill sets of their employees.

The Wing Data-First 50: A Generational Shift

The first generation of Data-First business applications focused on key online processes, such as ad targeting and fraud detection, where the high volume and velocity of interactions meant no human could be in the loop. The downside of getting any one decision wrong was low, so it made sense to experiment here with highly automated approaches. As the applications were still all about understanding and influencing consumer behavior, they were also a logical crossover point from consumer apps.



Now a second generation of Data-First applications is emerging. These support skilled human operators and analysts, and are being deployed across multiple categories like IT Security and IT Operations where the consequences of errors are more serious (see chart below).



DigitalGenius

Solvvy

TAKT.

Handling interactions with customers via call centers is a costly and labor-intensive activity for many companies. It is also a very sensitive one, as poor customer experiences can lead to brand damage and lost sales.

Hence strong interest in applications that can assist, and in some cases completely automate, customer support tasks. Companies such as Solvvy and DigitalGenius ingest vast amounts of historical data about prior customer interactions via customer service transcripts, support forums, and other sources, and then use this to learn how best to deal with fresh requests for support.

DigitalGenius uses neural networks to analyze the content of incoming emails, social media messages, and texts. The application leverages this information to keep learning about the best ways to handle inbound queries and suggests responses to them to customer support agents. An agent can then decide whether to send the machine-generated reply to a customer or to personalize it themselves.

The scope for full automation here is significant, with “chatbots” resolving simpler queries and bringing a human operator into the loop only when more complex issues need to be resolved.

The use of Data-First applications to boost the efficiency of employees was also a theme in our Sales category, where another group of innovative startups has developed products that help sales teams drive more revenue. Chorus and Gong use a combination of speech recognition technology,

natural language processing and other approaches to capture, transcribe and analyze sales conversations. That analysis drives insights into how future calls can be more effective and provides feedback that can be used by managers of sales teams to coach employees.

Human Resources

The logo for Checkr, featuring the word "Checkr" in a bold, blue, sans-serif font.The logo for entelo, featuring the word "entelo" in a blue, lowercase, sans-serif font.The logo for GLINT, featuring the word "GLINT" in a green, uppercase, sans-serif font.The logo for hiQ, featuring the word "hiQ" in a blue, lowercase, sans-serif font.The logo for textio, featuring a colorful circular icon followed by the word "textio" in a lowercase, sans-serif font.

HR systems have long epitomized business-logic-first applications. Traditional products miss opportunities to capture relevant data, make little use of data they do have, and don't tap much external data at all. But that's starting to change thanks to the efforts of entrepreneurs inspired by the experience of companies such as Google, which pioneered a Data-First approach to "people analytics" and widely publicized its efforts.

Recruitment is one of the areas new entrants are focused on. It makes an attractive target for entrepreneurs because of the proven high monetization potential here. New approaches can have a swift and easily measurable impact, which encourages HR teams to adopt them.

Applications on our list leverage data science and AI to help HR executives write more effective job advertisements (Textio); filter potential candidates more efficiently (Entelo); and run more efficient background checks (Checkr). Textio's application adds more than 10M job postings a month and their outcomes, and applies natural language processing technology to understand what words and phrases are the most effective in helping companies to hire people into specific roles.

Having attracted talent, companies are also keen to retain it. That has created an opportunity for another kind of Data-First company, which applies algorithms to internal HR and/or external datasets in order to assess flight risk amongst a company's employees. HiQ is one example of this new breed and another is Glint, which uses natural language processing to analyze the results of employee surveys and combines the output with other internal HR data to assess attrition risk.

HR has a relatively low volume of data compared to certain other categories. At the other end of the spectrum are several areas, such as IT Operations and IT Security, where a data tsunami threatens to overwhelm human operators.

IT Operations



Managing IT Operations has become dramatically more complex over time, with event and alert volumes continually increasing as hackers become more sophisticated, and companies roll out initiatives such as cloud, mobile, and microservices as part of their digital transformation strategies. To deal with this complexity, businesses are turning to AI-powered Data-First applications to help them transform and automate incident management processes across their production stacks, including application, infrastructure ticketing, and monitoring tools.

“AIOps” is a recently-defined category by Gartner and includes algorithmically-oriented products from companies such as SignalFX and Moogsoft. These applications ingest massive amounts of event data of various types, and then apply machine learning models in real time to determine which events are routine and which are service-affecting situations that merit the attention of a human operator. As well as detecting anomalies rapidly, some applications also predict failures, and can automate remediation in certain areas in order to minimize downtime.

AIOps platforms are great examples of technologies that create very valuable synthetic data sets. Moogsoft, for example, has developed a “Situation Room” collaboration tool in its flagship AIOps product which uses pattern matching to group related events into “situations” and route them to the appropriate teams, complete with data on how previous incidents were resolved in order to build a historical data set of key people, symptoms and cures. This dataset is then used to derive recommendations for actions to address recurring incidents. As part of this process, the platform uses human operators’ reactions to past incidents to train a neural network so that it gets progressively smarter and can thus minimize the risk of downtime in future.

Moogsoft and some other AIOps applications also allow teams to automate responses to both pre-defined and learned conditions. These responses can range from opening a support ticket to provisioning additional capacity in a cloud service like AWS.

IT Security



Like IT operations teams, security ones are also faced with a rapidly expanding volume and variety of data. This trend is overwhelming traditional Security Information and Event Management products, and has created an opportunity for AI-powered Data-First applications to grab both mindshare and market share.

Exabeam leverages a flexible and scalable data architecture to collect and parse logs and other raw data sources at lightning speed. The company then applies advanced machine learning tools to these datasets to identify anomalous behaviors that indicate suspect activity. Exabeam's application can also automate responses to hacks via incident workflows, such as automatically resetting passwords or isolating infected devices.

PatternEx is an example of a Data-First application that sets out to anticipate threats, as well as detecting attacks that may already be underway. Its Threat Prediction Platform uses machine learning to flag behaviors that it thinks are malicious. Human analysts help train the system, which then produces refined models that are applied to yet more suspect behaviors. SlashNext also helps companies to anticipate and defend against attacks. It uses a unique form of machine learning that effectively translates expert security researchers' deep knowledge of hacker behavior into a codified solution that forms a core part of its Active Cyber Defense System product.

Manufacturing and Supply Chain



Applications that generate predictive analytics are also valuable in manufacturing. Products from

companies like Uptake and Sight Machine ingest data from sensors on machines and other sources, and use these to forecast potential operational outages. As well as minimizing unplanned downtime, the algorithms can also help optimize process efficiency and drive improvements in quality control.

Teams developing new materials to be used in manufacturing can also benefit from an AI-powered Data-First approach. Citrine Informatics has created a vast proprietary database of patents, research papers and technical reports relating to materials. Its machine learning algorithms leverage this database to help customers predict the performance of new materials they are thinking of using, accelerating product development and manufacturing processes.

Behind Citrine and the other 49 business applications on our list are an impressive bunch of entrepreneurs who are charting new territory in their categories. The next section of the report provides a brief overview of their backgrounds.

The Wing Data-First 50 Founders

When we looked at the bios of Wing Data-First 50 founders, we discovered that a significant number had worked in consumer internet companies rather than business tech ones before creating their startups. This contrasts with the founders of leading SaaS applications, many of whom came from an enterprise tech background. Examples of the latter include Marc Benioff, a former Oracle executive who spearheaded the creation of Salesforce.com, and Dave Duffield and Aneel Bhusri of PeopleSoft, who went on to found Workday.

As we noted at the start of the report, Data-First approaches were pioneered in consumer web businesses, which helps explain the presence of a significant number of their alumni. Google was the company whose former employees cropped up most frequently in founding teams of the startups on our list. Amazon, Groupon, and Yahoo! were among the other prominent consumer web companies that featured on founder CVs.

A few categories like IT Security and IT Operations were nevertheless largely dominated by entrepreneurs who had previously worked for business tech companies. The sensitive and highly technical nature of these areas mean that customers here may be more inclined to place their faith in young companies led by founders with deep domain expertise.

Academia was another source of Data-First leaders, with a number of founders coming directly

from universities such as Carnegie Mellon, MIT, and UC Berkeley that are among the pioneers in AI research. We also found a few entrepreneurs who had previously worked in government agencies like NASA and the NSA. The founders from this sub-group are often the leaders of the data science teams that are a critical source of competitive advantage for their startups.

Data-First AI Applications And Digital Transformation

In an interview with the Wall Street Journal earlier this year, Philip Fasano, the then CIO of AIG, was quoted as saying: “We’re at a point where artificial intelligence has finally come of age. Any CIO...has to be considering what AI and knowledge-based systems mean to their business.”

That is wise advice—and should apply to all of the positions in the C-suite. It’s true that there have been a number of false dawns for AI. But as the applications on the Wing 50 List show, AI-powered Data-First applications can augment the capabilities of employees significantly, and in some cases perform tasks far more efficiently than human operators. The potential for dramatic productivity gains are clear: Moogsoft, for instance, has helped one of its financial services clients increase the number of servers that an IT Operations employee can monitor efficiently from just over 1,000 to over 10,000. Next-generation applications can also help increase revenue, reduce operational risk, and improve customer experiences.

Many more companies are now thinking seriously about how to use them to reinvent elements of their operations as part of their digital transformation strategies. Those that fail to move swiftly enough will find themselves at a disadvantage to early adopters. But the entrepreneurial teams behind the Wing Data-First 50 and other next-generation AI-powered business applications will still have to work hard to educate potential enterprise customers who may be wary of entrusting core processes to “black boxes” of algorithms.

To ease adoption, most are positioning themselves as complementary to existing systems. This also helps them get access to legacy data. Their added value is in the new synthetic data sets they generate and the advanced analytics they provide. However, this is likely to be just the first step in what ultimately becomes an exciting wholesale transformation of the business application landscape.

Future Opportunities For Entrepreneurs

As this transformation unfolds, we foresee many more exciting young companies emerging. Here are just a few of the opportunities we think ambitious founders will take advantage of:

- **Expansion across application categories**

Expect more AI-powered Data-First applications to appear in segments where we found only a handful of potential candidates for the Wing Data-First 50. Financial Planning & Analysis is one promising category. Other categories such as Legal, Risk Management, and Knowledge Management are also ripe for Data-First disruption. We're also likely to see more activity in some busier categories such as Manufacturing & Supply Chain, where AI technologies can be applied to plentiful external and internal data related to products and components.

- **Increased automation**

Most of the marketing messages behind the applications on our list position them as ways of making workers more efficient and productive. "Machine-assisted" employees can be freed from the drudgery of certain basic tasks in order to focus on more important ones. But as the quality of datasets and algorithms improves over time—and as more companies learn to trust machine-driven decision-making—next-generation applications will increasingly automate more higher-order tasks too.

- **Growth of the Internet Of Things (IOT)**

Only a tiny fraction of products and machines currently have sensors embedded in them. As that changes over time, the volume of data available about them will grow exponentially. This will make it possible to offer intelligent, automated customer support to many more things, including a wide range of consumer products. It will also create incredibly rich datasets that can be mined by next-generation AI-powered applications for other purposes, from sales intelligence to security.

We would very much like to hear from startups that are thinking about these and other opportunities being created by what Satya Nadella, the CEO of Microsoft, has referred to as "the third run time". At a conference earlier this year, Nadella called operating systems the first run time and browsers the second. Smart agents, he explained, were the third run time "because in some sense, the agent knows you, your work context, and knows the work". The business applications that harness this third run time to the needs of the enterprise most effectively will be the first that

leap to mind when we compile future Wing Data-First 50 lists.

The Wing Data-First 50 Company List

Sales

Aviso

www.aviso.com (<http://www.aviso.com>)

Chorus

www.chorus.ai (<http://www.chorus.ai>)

Clari

www.clari.com (<http://www.clari.com>)

Conversica

www.conversica.com (<http://www.conversica.com>)

Gong

www.gong.io (<http://www.gong.io>)

Marketing

6sense

www.6sense.com (<http://www.6sense.com>)

AgilOne

www.agilone.com (<http://www.agilone.com>)

Drawbridge

www.drawbridge.com (<http://www.drawbridge.com>)

Mintigo

www.mintigo.com (<http://www.mintigo.com>)

Persado

www.persado.com (<http://www.persado.com>)

Customer Relationship Management

Appuri

www.appuri.com (<http://www.appuri.com>)

Blueshift

www.getblueshift.com (<http://www.getblueshift.com>)

Juvo

www.juvo.com (<http://www.juvo.com>)

Retention Science

www.retentionscience.com (<http://www.retentionscience.com>)

Sailthru

www.sailthru.com (<http://www.sailthru.com>)

Tact

www.tactile.com (<http://www.tactile.com>)

Customer Experience Management

DigitalGenius

www.digitalgenius.com (<http://www.digitalgenius.com>)

Solvvy

www.solvvy.com (<http://www.solvvy.com>)

Takt

www.takt.com (<http://www.takt.com>)

E-commerce

BloomReach

www.bloomreach.com (<http://www.bloomreach.com>)

Eversight

www.eversightlabs.com (<http://www.eversightlabs.com>)

Reflektion

www.reflektion.com (<http://www.reflektion.com>)

IT Security

Area 1 Security

www.area1security.com (<http://www.area1security.com>)

Cybereason

www.cybereason.com (<http://www.cybereason.com>)

Exabeam

www.exabeam.com (<http://www.exabeam.com>)

PatternEx

www.patternex.com (<http://www.patternex.com>)

SlashNext

www.slashnext.com (<http://www.slashnext.com>)

Fraud Prevention/Detection

DataVisor

www.datavisor.com (<http://www.datavisor.com>)

Feedzai

www.feedzai.com (<http://www.feedzai.com>)

Forter

www.forter.com (<http://www.forter.com>)

Sift Science

www.siftscience.com (<http://www.siftscience.com>)

Signifyd

www.signifyd.com (<http://www.signifyd.com>)

Financial Planning & Analysis

Prevedere

www.prevedere.com (<http://www.prevedere.com>)

TravelBank

www.travelbank.com (<http://www.travelbank.com>)

IT Operations

Moogsoft

www.moogsoft.com (<http://www.moogsoft.com>)

Rocana

www.rocana.com (<http://www.rocana.com>)

SignalFX

www.signalfx.com (<http://www.signalfx.com>)

Human Resources

Checkr

www.checkr.com (<http://www.checkr.com>)

Entelo

www.entelo.com (<http://www.entelo.com>)

Glint

www.glintinc.com (<http://www.glintinc.com>)

HiQ Labs

www.hiqlabs.com (<http://www.hiqlabs.com>)

Textio

www.textio.com (<http://www.textio.com>)

Legal

Casetext

www.casetext.com (<http://www.casetext.com>)

Ravel Law

www.ravellaw.com (<http://www.ravellaw.com>)

Manufacturing & Supply Chain

Citrine Informatics

www.citrine.io (<http://www.citrine.io>)

Clear Labs

www.clearlabs.com (<http://www.clearlabs.com>)

ClearMetal

www.clearmetal.com (<http://www.sightmachine.com>)

Sight Machine

www.sightmachine.com (<http://www.sightmachine.com>)

Uptake

www.uptake.com (<http://www.uptake.com>)

Zymergen

www.zymergen.com (<http://www.zymergen.com>)

<http://wing.vc/blog/announcing-the-wing-data-first-50-ai-powered-business-applications>

© 1990 — 2019 Wing, Menlo Park, CA