

Doug Bryan

Doug.Bryan456@gmail.com, Doug-Bryan.com

I'm an AI/machine learning executive, evangelist, and business transformation consultant focused on rapid ROI with sustainable change. I have experience as a SVP of products at a global ad agency, personalization product manager and tech lead at Amazon.com, ecommerce R&D manager at Accenture Labs, startup co-founder, and lecturer in computer science at Stanford University. My 25 years of data science experience covers hundreds of projects across many industries.

Experience

Digital^Shift, remote

Fractional Account Executive and AI Consultant

July 2025 –

Business development and thought leadership for a multinational consulting firm focused on rapid AI ROI.

Touchpoint Strategies, remote

Partner

April 2024 – April 2025

Self-funded startup. I developed marketing strategy and sales enablement products using Bombora B2B intent data, BigQuery and four AI LLMs (Claude, Gemini, ChatGPT and Perplexity).

Why of AI, remote

Strategy Consultant

May 2024 – April 2025

Interdisciplinary workshops and consulting to help organizations develop effective AI strategies.

Dataiku, remote

Field CDO

April 2021 – February 2024

Strategic advisor to large customers. Common topics addressed include machine learning / AI use case ideation and prioritization, generating significant positive data science ROI, and change management programs for migrating to a data mesh operating model. Wrote articles and white papers on topics such as [data mesh](#), [upskilling](#), and [AI trust](#). Conducted seminars and brainstorming workshops. My customers included Ralph Lauren, NBCUniversal, Takeda, Love's Travel Stops and DigiKey.

Nitrogen.ai, Chicago

Co-founder and CTO

April 2019 – April 2021

Pre-funding, pre-revenue startup developing a two-sided online marketplace of external data for data scientists. Product owner responsible for product management, architecture and development. Managed an offshore development team in India using Kanban in JIRA. Net promoter score from 20 beta customers was good but then COVID hit and shut us down.

Qinship, Chicago

Product Strategy Consultant

June 2020 – November 2020

Developed a product strategy for a startup that pivoted from CRM to direct-to-consumer marketing and sales automation.

Merkle, remote

SVP, Data Science Products

April 2018 – April 2019

VP, Data Products

October 2016 – April 2018

Led the data science product development team for M1, a platform for multi-channel, person-level marketing with thousands of attributes on every U.S. adult. Grew my team from 1 to 6 data scientists.

Acted as product owner and had dotted line responsibility to multiple scrum teams. Data science products included self-service automatic look-alike models, customer segmentation, and panel-to-census data fusion.

- 100% automatic look-alike models using 8,000 variables. Generated the best accuracy of three Merkle modeling techniques tested at a Dentsu International hackathon at the Googleplex.
- Used customer segmentation to optimize ad campaigns, increasing conversion rate 46% on average during in-market tests. Moved Merkle from strong performer to a leader in the Forrester Wave for search marketing agencies.
- Automatic fusion of 38k survey respondents to a base population of 280m people
- Migrated modeling processes from on-premise SAS to AWS and Spark, [reducing costs 88%](#) and time-to-market 95%

RichRelevance, remote

Sales Engineering Manager for North America

January 2016 – October 2016

Principal Solutions Architect

January 2013 – October 2016

RichRelevance is San Francisco-based, venture-backed, software-as-a-service (SaaS) company that provides real-time omnichannel personalization and product recommendations to more than 250 multinational retailers. They served product recommendations to 600 million shoppers per month, 20 billion recommendations on Black Friday 2015 and managed 3 petabytes of shopper behavioral data. I had a variety of pre- and post-sales client-facing roles, making extensive use of Hadoop, Hive SQL and R. Projects included:

- Predictive analytics for a holiday [gift finder app at Marks & Spencer](#) that achieved 12% conversion rate on Cyber Monday
- Omnichannel strategic customer segmentation for Barneys New York, a luxury department store
- Sold and consulted on personalization for Barneys New York mobile app and sales associate app ([press article](#) and [client interview](#))
- Sold the first project where we partnered with Accenture. It delivered real-time product and experience recommendations on [Carnival Cruise Line's mobile app](#).
- President's Club 2015
- Helped sell a \$600,000 per year SaaS deal to L.L.Bean and [presented](#) with the client at Strata + Hadoop World ([press article](#))

Core Audience, remote

VP Data Sciences and Analytics

June 2012 – January 2013

A DMP, owned by Hearst, for brands to manage audience-based marketing across bought, earned and owned media. I managed a text mining team in St. Petersburg, Russia, developed and managed a data science team in Chicago, produced DMP reports for advertisers and publishers, developed audience segmentation algorithms, and led pilot project design and execution.

iCrossing, remote

VP Analytics

June 2011 – January 2013

iCrossing is a digital ad agency owned by Hearst. My responsibilities included managing a team of 15 people in four cities, product owner during the implementation of a Netezza SQL data warehouse, performance marketing reporting, measuring ad effectiveness in new media such as social media, attribution modeling algorithms, and sales. Clients included PNC Bank, DirecTV, Mazda and Sears.

OptiMine Software, remote

Co-founder

July 2009 – May 2011

OptiMine forecasted the performance of each paid search ad placement each day and automatically set optimal bids. It was used by leading organizations such as Overstock.com and Wells Fargo, and could be implemented in-house, hosted, or as software-as-a-service (SaaS).

- Co-developed data science algorithms, logical data model and SQL queries
- Optimized a large retailer's paid search advertising spend for 18 months, increasing profit-based return on ad spend (ROAS) by 40%
- Helped customers select the marketing performance metric to use, such as leads, qualified leads, revenue or profit

RichRelevance, Founding Board of Advisors

2007 – 2010

KXEN (acquired by SAP), remote

Technical Director

January 2005 – June 2009

KXEN was the first autoML software vendor. Their solution is based on Vapnik's Statistical Learning Theory. Responsible for pre-sales, pilot studies, proof-of-concept workshops, post-sales and collaboration with customers. My customers included Experian, Discover Financial Services, Washington Mutual, Rogers, and Meredith.

- Regularly achieved a 100% increase in marketing campaign response rate
- Developed and validated an analytics process for Motorola that **reduced root-cause analysis time by 90%**. It was actually 99% but marketing thought that would be unbelievable.
- Designed an automated, real-time analytics method for an online call center optimization service (LivePerson) that increased online sales 200 – 700%
- Designed an analytics process for Overstock.com that tested a dozen email marketing campaign variations simultaneously
- Developed and validated an analytics process for Discover Financial Services that used 10 times more variables (10,000 total) in their customer behavior models
- Developed and validated a predictive model for a hedge fund that predicted the success of commodity trading advisors
- Developed new techniques for **"long tail"** marketing using techniques that later became known as hyper-parameter optimization and weighted ensembles. Presented them to key customers in Germany, France, U.K and the U.S.

Northwestern University Institute on Complex Systems,**Complexity in Action Network**, Evanston, IL

Industry Board of Advisors, founding member

2005 – 2007

Amazon.com, Seattle

September 2003 – December 2004

Product Manager, Personalized Product Recommendations

- Managed the user experience and data science teams that produced product similarity data and personalized product recommendations. Generated over \$2 billion in incremental revenue per year.
- Managed the migration of 40 production data pipelines from a Seattle data center to a Virginia data center (what's now AWS East) with only 20 minutes of downtime
- Departmental representative on a company-wide program to develop an ecommerce platform for other online retailers, including Target. It was the largest new program in the company at the time.

Independent Consultant

2002 – 2003

- Assessed the software needs of a new Texas Instruments product line, working with team members in three countries. Developed and presented a 5-day workshop on real-time software in Nice, France.
- Technology assessment and vendor evaluation for ontology development tools for the College of American Pathologists. Developed custom tools to support a very large medical ontology, SNOMED, using XML, Java, DAML+OIL, OWL and Jena.

SPSS, Chicago

April 2002 – September 2002

Senior sales engineer for text mining. Created demos, case studies, white papers, business cases and presentations for a new text mining product line. Did pre-sales work with Bristol-Myers Squibb, Johnson & Johnson and HSBC.

Accenture, Center for Strategic Technology Research, Northbrook, IL

June 1995 – August 2001

Senior manager responsible for identifying important emerging technologies and communicating them through business cases, presentations, and elaborate prototypes. Regularly presented to the press (Wall Street Journal, New York Times, Internet Week, ZDNet, CNet, Chicago Tribune, Associated Press, PBS' Nightly Business News), Fortune 500 customers (AT&T, Mellon Bank, Best Buy, Texas Instrument, Motorola, Dow, AstraZeneca), startups (Blue Martini, Receipt City, Web Wombat) and government (U.S. Dept. of Education, U.S. Dept. of Defense). Worked directly with clients in Australia, U.S., and U.K.

- Technical manager for a \$4 million, federally funded software R&D project
- Developed a novel user interface and real-time predictive analytics algorithm for dynamic merchandising on the Web using information retrieval and information visualization techniques. It included an associative network model of 12,000 consumer products and a real-time, 1,500-variable model of user behavior based on click-stream data. Presented to internal and external clients worldwide.
- Worked with Sun Microsystems, IBM, HP and Microsoft to develop UDDI, an international Web services standard

Stanford University, Palo Alto, CA

June 1984 – June 1995

Technical lead of a \$1.5 million per year, 20-person research team, managing all aspects of development. Designed and implemented experimental software development languages and environments. Also was a lecturer in computer science.

- Developed and taught the first object-oriented design class at Stanford
- Gave the first object-oriented design lecture at HP Labs
- Selected, procured and managed the first network of personal workstations on campus (15 Sun Microsystems 3/110's and 2 Sun 3/280's)
- Published three textbooks and many papers
- One summer taught FBI agents how to surveil the internet

Link Flight Simulation, Sunnyvale, CA

1988 – 1990

- Management consultant on a real-time radar simulator built on a pipeline of five DEC VAX's
- Adopted major new design technologies while the project was underway
- Significantly changed team behavior in mid project using chalk talks, coaching, ad hoc design reviews, and peer code reviews

Lockheed Missiles & Space Company, Sunnyvale, CA

June 1982 – June 1984

- Developed simulator training for NASA Space Shuttle astronauts using early IBM PC's
- Advanced development of battlefield command and control systems ([screen shot](#))
- Advanced development of a new computer graphics standard for a new programming language

Education: B.S., Computer Science and Mathematics with honors, Carroll College, Waukesha, WI

A few recommendations from [LinkedIn](#)

“One of the smartest guys I know. He brings a level of experience that is hard to find within the world of big data and marketing solutions.” — RichRelevance VP Strategy & Corporate Development

“Rare combination of client facing skills, deep marketing experience, intellectual rigor, self-starting behavior, imaginative problem solving, personal responsibility and value focus that make him a great data product and team lead. Always one of the smartest people in the room... gift is letting you know that gently enough to not be daunted by it.” — Merkle CTO

“A rare combination of sound technical knowledge and practical business judgment... ability to communicate complicated material simply. But his vast experience and sharp mind are his greatest strength because he is the one that people seek out to bounce ideas against.” — U.S. Federal Reserve Bank risk analyst

“Very talented and at ease with easy tasks and very complex ones as well. He excels at taking out complexity and makes things simpler.” — KXEN CEO and co-founder

History of thought leadership

- Co-founded a venture-backed company that automatically generates tens of thousands of predictive models per week to optimize ad placements
- Founding advisors to Northwestern University’s Institute on Complex Systems industry consortium and wrote their first case studies
- [Presented](#) at the first Society of Industrial and Applied Mathematics (SIAM) data mining conference
- [Presented](#) at the first Association for Computing Machinery (ACM) ecommerce conference
- Managed the Amazon.com team that launched the first asymmetric product similarities (people who viewed this item bought these other items...)
- Co-developed the first industry standard for Web services registries
- Developed and taught the first object-oriented design class at Stanford University

Patents

- U.S. Patent 6,664,980, "Visual navigation utilizing Web technology," Douglas L. Bryan and Anatole V. Gershman, December 16, 2003
- U.S. Patent Application 2012/0130,798, "Model sequencing for managing advertising pricing," Robert Cooley and Douglas Bryan, filed November 23, 2010
- U.S. Patent Application 2012/0130,828, "Source of decision considerations for managing advertising pricing," Robert Cooley and Douglas Bryan, filed December 31, 2010

Books published

1. *Software Engineering with Ada*. Grady Booch and Doug Bryan, 3rd ed., Benjamin-Cummings (1994)
2. *Exploring Ada, Volume 2*. G.O. Mendal and D.L. Bryan, Prentice-Hall (1992)
3. *Exploring Ada, Volume 1*. D.L. Bryan and G.O. Mendal, Prentice-Hall (1990)

Select publications and presentations

1. "Upskilling & Hiring for AI: Build Unicorn Teams, Not Unicorn People," Dataiku, 2023. [pdf](#), [web](#), [operating model summary](#)
2. "Build Trustworthy AI," Dataiku, 2023 [pdf](#), [web](#)
3. "Meet a cloud native," Merkle Inc. blog, June 2017 [↗](#)
4. "XGBoost performance for customer look-alike models," Tom Tom Festival Machine Learning Conference, Charlottesville, VA, keynote address, April 13, 2017 [↗](#)
5. "Five types of user matching challenges and how to solve them: cross-device, cross-channel and cross-brand," RichRelevance blog, August 2016 [↗](#)
6. Global sales engineer training, RichRelevance, 2015.
7. "Omnichannel personalization means mobile first, and this time we mean it!", RichRelevance blog, May 2016
8. "Inside King of The Hill: The Science Behind RichRelevance's Decisioning Engine," RichRelevance, July 2015 [↗](#)
9. Marks & Spencer gift finder case study, January 2015 [↗](#)
10. "Transitioning from Original Big Data to the New Big Data: L.L.Bean's Journey," Strata + Hadoop World, New York City, October 2014. [↗](#), [press coverage](#)
11. Wine.com case study. Coached them to develop their own recommendation algorithms in Hadoop and Hive, September 2014 [↗](#)
12. "Ad Men + Data Specialists: When Two Worlds Collide," Direct Marketing Association Conference panel, October 2013. Press coverage at [Ad Exchanger](#) and [Ad Age](#).
13. "Hadoop success requires avoidance of past data mistakes," *TechRepublic*, October 2013
14. "Acxiom's inaccurate data and why it's so useful," *Direct Marketing News*, September 19, 2013 [↗](#)
15. "How CMOs can measure return on ad spend with better modeling and conversion attribution," iCrossing blog, January 2012 [↗](#)
16. "A new framework for measuring social media," iCrossing blog, May 2012. Also presented at a DMA workshop and WOMMA certification class. [↗](#)
17. "Always be reactivating," RichRelevance blog, Sept. 2008
18. "Pennies from eBay: the determinants of price in online auctions," David Lucking-Reiley, Doug Bryan, Naghi Prasad, and Daniel Reeves. *The Journal of Industrial Economics*, June 2007 (First draft Nov. 1999.) [↗](#)
19. "Data mining for quality improvement," Françoise Fogelman Soulié and Doug Bryan, ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) 2007, workshop paper [↗](#)
20. "Extreme Customer Lifecycle Management for Long Tail Retailers," KXEN Users Group meeting, San Francisco, March 2007. Best presentation award. [↗](#)
21. "KXEN Analytics reduce analysis time by 90%," KXEN telecommunications case study, 2006 [↗](#)
22. "Industry case studies of applying complexity science," Northwestern University Institute on Complex Systems, Complexity in Action Network, 2006 [↗](#)
23. "Lifting predictive analytics productivity at Sears," Paul Bibler and Doug Bryan, *Target Marketing* magazine, October 2005 [↗](#)
24. "Experimental politics: ways of virtual worldmaking," in *Cognitive Technology: Instruments of Mind* (4th Intl Conf. on Cognitive Technology), Springer-Verlag, Lecture Notes in Artificial Intelligence (LNAI) 2117, August 2001, 432–441 (Probably the most bizarre paper I've written.) [Article](#) and [paper](#)
25. "Twenty use scenarios for Web services," Doug Bryan (ed.), UDDI Advisors Group, October 2001
26. "Mapping vendor spaces using high-level relations," Doug Bryan, Society of Industrial and Applied Mathematics (SIAM) *1st Intl. Conf. on Data Mining, workshop on Web Mining*, April 2001 [↗](#)

27. "The Aquarium: a novel user interface metaphor for large online stores," Doug Bryan and Anatole Gershman, in 2nd Intl Workshop on Web-based Information Visualization (WebVis 2000), in *Proc. of the 11th Intl Conf. on Database and Expert Systems Applications (DEXA)*, IEEE Computer Society Press (2000) ↗
28. "Courtney misses the logic," Accenture internal essay about Napster, July 2000 ↗
29. "Opportunistic exploration of large consumer product spaces," Doug Bryan and Anatole Gershman, in *Proc. of the 1st ACM Conf. on Electronic Commerce*, ACM Press (1999), 41–47 ↗
30. Accenture internal essay on how long it took the term "Monica Lewinsky" to spread through traditional and digital media. This was pre-Google. Feb. 1998 ↗
31. "Toward software plug-and-play," F. Bronsard, et al., in *Proc. of the Symposium on Software Reusability*, ACM Press (1997)
32. "On the need for 'required interfaces' of components," Á. Ólafsson and D. Bryan, in M. Muehlhaeuser (ed.), *Special Issues in Object-Oriented Programming*, dpunkt Verlag, Heidelberg (1997)
33. "Exactness and clarity in a component-based specification language," D. Bryan, in H. Kilov and B. Harvey (eds.), *Specification of Behavioral Semantics in Object-oriented Information Modeling*, Kluwer Academic Publishers (1996), 1–16
34. "Specification and analysis of system architecture using Rapide," D.C. Luckham, J.J. Kenney, L.M. Augustin, J. Vera, D. Bryan, and W. Mann. *IEEE Transactions on Software Engineering*, 21(4):336-355 (April 1995)
35. Demonstration of the Rapide toolset, 16th International Conference on Software Engineering, IEEE, ACM, Sorrento, Italy, May 1994
36. "Partial orderings of event sets and their application to prototyping concurrent, timed systems," D.C. Luckham, J. Vera, D. Bryan, L. Augustin and F. Belz. *Journal of Systems and Software*, 21(3):253–265, North-Holland Press (1993). Also as Computer Systems Lab technical report CSL–TR–92–515, Stanford University, April 1992.
37. "Rapide-0.2 language and tool-set overview." D. Bryan, Computer Systems Lab technical note CSL–TN–92–387, Stanford University, February 1992
38. "Design of run-time monitors for concurrent programs," D.P. Helmbold and D.L. Bryan, Computer Systems Lab technical report CSL–TR–89–395, Stanford University, October 1989
39. "An introduction to Task Sequencing Language, TSL 1.5, (preliminary version)," D.C. Luckham, S. Meldal, D.P. Helmbold, D.L. Bryan, and W. Mann. Report No. 38, Department of Informatics, University of Bergen, Bergen, Norway, July 1989
40. "An algebraic specification of the partial orders generated by concurrent Ada computations," D. Bryan, in *Tri-Ada '89 Conference Proceedings*, ACM Press (1989), 225–241
41. Demonstration of the TSL tools, Tenth International Conference on Software Engineering, IEEE, ACM, Singapore, April 1988
42. "Task Sequencing Language for specifying distributed Ada systems: TSL-1," D.C. Luckham, D.P. Helmbold, S. Meldal, D.L. Bryan, and M.A. Haberler, in *Proc. of the Conf. on Parallel Architectures and Languages*, Lecture Notes in Computer Science No. 259, Springer-Verlag (1987), 444–463
43. "Dear Ada" column in the Association of Computing Machinery (ACM) journal *Ada Letters*, 1986 – 1992
44. "Object-oriented design with an example using Ada," D. Bryan. *MILSTAR Ada Study in Software Engineering Issues, Phase II*, LMSC–D976358, Lockheed Missiles & Space Co., 1984 ↗
45. "An Ada implementation of the Graphics Kernel System (GKS)," D. Bryan, AdaJUG/AdaTEC (ACM SIGPLAN) meeting, Dallas, October 1983 ([screen shot](#))