VIJAY KRISHNAN

Los Altos Hills CA 94022

(650) 796-6388
vijayk@cs.stanford.edu

• Founder & CTO (earlier Co-CEO) of Rover, acquired by Revcontent where I am currently Senior Vice President of Data Science.
• Senior Data Science/Machine Learning/Engineering/Product/Business leader.
• Experienced in building, growing and leading teams applying large scale machine learning methods toward creating massive commercial value.
• 10 years industry experience leading large scale commercial Machine Learning efforts toward problems in text categorization, personalized search, personalized content recommendations, CTR prediction and eCPM maximization for personalized and contextual ad targeting, and optimization of other business objectives as Sr. Scientist at Yahoo, Founder & CTO of Rover and SVP of Data Science at Revcontent.
• Authored highly cited Machine Learning, Web Search, Semantic Search, Entity extraction, Data Mining, Natural Language Processing papers in top conferences and journals.

EMPLOYMENT

Senior Vice President of Data Science   **Revcontent, Sunnyvale CA**   **Feb '17-Present**

Revcontent is the fastest growing native advertising content recommendation networks in the world, reaching 1.5 Billion users monthly. Its publishers include Forbes, Newsweek, Atlantic, Reuters, IBN.

- Joined Revcontent, after Revcontent's acquisition of my startup, Rover. [TechCrunch](#) | [Forbes](#)
- Lead efforts to integrate Rover's Machine Learning based personalized content recommendation technology with Revcontent's systems for:
    - CTR Prediction and eCPM maximization
    - Improving ad quality
    - Offering advertisers audience interest targeting capability to target over 3000 user interests
    - Offering publishers personalization tools to maximize user engagement and sessions lengths.

Founder & CTO   **Rover (acquired by Revcontent), Sunnyvale, CA**   **Feb '08 – Feb '17**

Rover is a Machine Learning based personalized content recommendation company. Rover learns a user's fine grained interest preferences from their past content consumption activity and uses that to recommend highly engaging content on the topics they love.
- Co-founded Rover, raised over $6 Million in Venture Capital.
- Drove metrics driven social and sharing driven growth efforts that grew web app (then called Flipora) to 40 Million registered users.
- Drove revenue efforts via advertising and technology partnerships, generating $5 Million in revenues.
- Led Large Scale Machine Learning efforts toward content categorization, personalization technology, technology to score content for interestingness etc. to deliver highly engaging content recommendations.
- Served initially as Co-CEO, later as CTO, hired a stellar team of engineers and designers, negotiated multiple business partnerships, drove different aspects of the product and ensured a successful exit by way of acquisition by Revcontent.
- Built our personalized search and real-time search engine products (then called Infoaxe), before we decided to shift company focus from web search to content discovery.

Sr. Scientist Associate   **Yahoo! Data Mining & Research, Sunnyvale CA**   **Summer '06, Aug '07-Jan '08**
- Invented a new fast, efficient, sparse multiclass Support Vector Machine.
- Raised categorization accuracy from about 12% to 90% on a taxonomy with many thousands of categories.

- 10x faster training and prediction, 500x more compact models with no accuracy loss.

EDUCATION

- **Stanford University:** M.S. in Computer and Science, AI Specialization. 2005-07. GPA: 3.92/4.00
- **IIT Bombay:** Bachelors & Masters in Computer Science & Engineering. 2000-05. GPA: 9.18/10.00

PUBLICATIONS & PATENTS

- **Vijay Krishnan, R Raj.** Web Spam Detection with Anti-Trust Rank. In **AIRWeb (ACM SIGIR 2006)**
- **Vijay Krishnan, Christopher D Manning.** An Effective Two-Stage Model for Exploiting Non-Local Dependencies in Named Entity Recognition. In **COLING-ACL 2006.**
- **Vijay Krishnan, S Das, Soumen Chakrabarti**. Enhanced Answer Type Inference from Questions using Sequential Models. In **HLT/EMNLP 2005.**
- **Vijay Krishnan.** Shortcomings of Latent Models in Supervised Settings. **In ACM SIGIR 2005.**
- **Shipra Agrawal, Vijay Krishnan, Jayant Haritsa.** On Addressing Efficiency Concerns in Privacy-Preserving Mining. **In DASFAA 2004.**
- **Soumen Chakrabarti, Sujatha Das, Vijay Krishnan, Kriti Puniyani.** Text search enhanced with types and entities. **In Text Mining: Theory, Application, and Visualization, Srivastava and Sahami, eds., 2008.**
- **DJ Ruau, M Mbagwu, Joel T. Dudley, Vijay Krishnan, Atul J. Butte.** Comparison of Automated and Human Assignment of MeSH Terms on Publicly-Available Molecular Datasets. **Journal of Biomedical Informatics 2011. Distinguished Paper award in the AMIA 2011 Summit on Translational Bioinformatics.**
- **Sathiya Keerthi Selvaraj, Dmitry Pavlov, Scott Gaffney, Nicolas Mayoraz, Pavel Berkhin, Vijay Krishnan, S. Sellamanickam.** Method for efficiently building compact models for large multi-class text classification. **US Patent # 20090274376, 2009.**

My patent as well as all peer reviewed publications, cited over 500 times are in the areas of Data Management, Data/Text/Web/Sequence Mining, Graphical Models, Large Scale Machine Learning, Web Search, Information Retrieval, Mathematical Modeling and Natural Language Processing.

Press about me and my recently acquired startup, Rover

- TechCrunch   : Revcontent acquires Rover for better content recommendation tech
- Forbes   : Revcontent To Conquer The Content Discovery Market Through Rover Acquisition
- Huffington Post : Rover & the future of Content Discovery on Mobile
- Forbes   : Flipora Rebrands as Rover: Paves New Path to Content Discovery
- VentureBeat   : Flipora is the fastest-growing social surfing service you've never heard of
- All Things Digital : Flipora Launches a Smarter StumbleUpon Based on Your Web History
- TechCrunch   : Infoaxe's Flipora Passes 8M Registered Users, Adds Discovery Engine
- VentureBeat   : Infoaxe raises $3M to make online bookmarks obsolete
- TechCrunch   : InfoAxe Raises $3 Million For The Search Engine For Your Web History
- CNN: : Infoaxe's Search Engine: More current than Real-Time
- New York Times : Infoaxe — a Real-Time Search Engine That Doesn't Rely on Twitter
- TechCrunch   : Infoaxe Launches Real-Time Search Engine Based On Web History Platform
- Jerusalem Post : Infoaxe, a new service that takes the headache and hassle out of reviewing your browser history
- TechCrunch   : Infoaxe Records Your Browsing Sessions, Lets You Search Your Web History

More press links and other Information available on my Linkedin Page and my Stanford Homepage.