Mike Perkowitz

206.915.3138
mike@perkowitz.net

Summary

- Twenty years experience with machine learning, data mining, information extraction, and natural language processing.
- Eight years experience with startup companies.
- Six years experience with product planning and direction.
- Designed and implemented both prototype and production systems for end users, internal analysts, and research projects.
- Experienced with Java, Spring, Maven, Git, etc. Experienced with SQL, Graph DBs and NoSQL databases. Some experience with Titan, Cassandra, Hadoop, MapReduce.
- Excellent communicator – I have written research papers, white papers, and patents, presented at conferences and conventions, and led teams.

Professional Experience

CTO, Senior Data Scientist, ARO Inc. (formerly Kiha Software) *June 2011 – July 2014*

- Led ARO's research and technology team. Defined basic research and experimentation, productized results, published/patented where applicable.
- Defined, designed, and implemented ARO's scalable data management platform. Conducted market research and interviews to plan ARO's data insight product.
- Designed and implemented user modeling and segmentation algorithms.
- Managed ARO's patent portfolio, wrote provisional patent applications and technology descriptions, worked closely with patent attorneys to write and revise applications.

Senior Scientist, Kiha Software, Inc. *August 2008 – June 2011*

- Led the team that designed and implemented the text analysis pipeline, user modeling service, and search service.
- Designed and prototyped algorithms for named entity recognition, personalized entity recognition, user activity and place modeling, social relationship modeling, and search indexing and exploration.
- Managed Kiha's patent portfolio, wrote provisional patent applications and technology descriptions, worked closely with patent attorneys to write and revise applications.

Program Manager, Trailfire, Inc. *November, 2005 – March, 2008*

- Developed product strategy and planned features for Trailfire's web-based social annotation tool.
- Participated in crafting overall company strategy, planning and executing marketing initiatives, and meeting with potential investors and partners.
- Developed clustering and other data mining algorithms for analyzing and organizing web pages.
- Designed and implemented web-based user interface.
- Hired and worked with graphic designers on product look and feel.

Software Design Engineer, Amazon.com *June 2004 – November 2005*

- Designed and implemented a personalization engine that tailors advertising for Amazon affiliate web sites to the sites' content and traffic.
- Explored new tools and algorithms for understanding customer behavior to better recommend relevant products.
- Maintained a detailed customer behavior database, created tools for viewing and analysis of that customer behavior, and assisted others throughout the company in better understanding Amazon customers.

Researcher, Intel Research Seattle *November 2002 – June 2004*

- Helped design a system for identifying activities of daily living among residents of assisted-living facilities.
- Developed algorithms for extracting structured statistical models of activities of daily living from unstructured web sources for use in activity recognition, and created a partial knowledge base of activities.

Senior Engineer, Advanced Development, Appliant, Inc. *August 2000 - July 2002*

- Designed data analysis and reports showing the connection between web site performance and customer retention and sales.
- Prototyped new features and reports and wrote specifications for implementation.
- Created automated diagnostic tools for pinpointing problems in web site infrastructure.
- Performed in-depth analysis of customer data and prepared explanatory presentations for executive team and prospective customers.

Research Assistant, University of Washington Department of Computer Science

September 1993 - August 2000

- Developed new statistical and clustering algorithms for analyzing user behavior data (Ph.D. dissertation work).
- Applied these algorithms in a system to analyze user traffic at a web site and adapt the site structure to facilitate navigation.
- Developed an algorithm for extracting structured information from semi-structured internet information sources, and implemented a prototype system (M.S. work).
- Published numerous papers in artificial intelligence, machine learning, and data mining journals and conferences, and presented results at conferences.

Research Intern, Microsoft Corporation *June 1996 - September 1996*

- Developed working prototype of a system for sharing web-based information via a character-based user interface.

Education

University of Washington, Ph.D., Computer Science, February, 2001.
University of Washington, M.S., Computer Science, June, 1995.
Brown University, B.S., Cognitive Science and Linguistics, May, 1993.

Selected Patents

Kevin Francis Eustice et. al. **Using Mobile Device Data to Create a Storyline, Model User Routine and Personality, and Create Customized Recommendation Agents.** US Patent # WO 2014018687 A1. Filed 2013. Pending.
Michael Perkowitz, Kevin Francis Eustice, Andrew F. Hickl. **Recommendation Agent Using a Personality Model Determined from Mobile Device Data.** US Patent #20140032452A1. Filed 2013. Pending.
Michael Perkowitz et. al. **Dynamic Gazetteers for Entity Recognition and Fact Association.** US Patent #8429099 B1. Filed 2010. Issued 2013.
Michael Perkowitz et. al. **Electronic Message Prioritization.** US Patent # 8095613B1. Filed 2010. Issued 2012.
Jonathan D. Lazarus et. al. **Social Network Model for Semantic Processing.** US Patent #20100030715A1. Filed 2009. Issued 2010.
Raj K. Kizhakkekalathil et. al. **Providing Ad Information Using Plural Content Providers.** US Patent #7882046B1. Filed 2007, Issued 2011.
Michael Perkowitz, John E. Ball. **Sharing and Organizing World Wide Web References Using Distinctive Characters.** US Patent #6138128A. Filed 1997. Issued 2000.

Selected Publications

N. Lawson, K. Eustice, M. Yetisgen-yildiz, M. Perkowitz. **Annotating Large Email Datasets for Named Entity Recognition with Mechanical Turk.** Proceedings of the NAACL HLT Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk , 2010.

M. Philipose, K. Fishkin, M. Perkowitz, D. J. Patterson, D, Fox, H. Kautz, and D. Hahnel. **Inferring Activities from Interactions with Objects.** IEEE Pervasive Computing. Vol. 3, num. 2, 2004. pp. 50-57.

M. Perkowitz, M. Philipose, K. Fishkin, and D. J. Patterson. **Mining Models of Human Activities from the Web.** Proceedings of the Thirteenth International World Wide Web Conference (WWW 2004).

J. B. Chen and M. Perkowitz. **Using End-User Latency to Manage Internet Infrastructure.** Proceedings of the 2nd Workshop on Industrial Experiences with Systems Software, 2002.

M. Perkowitz. **Adaptive Web Sites: Cluster Mining and Conceptual Clustering for Index Page Synthesis.** Ph.D. dissertation, University of Washington, 2001.

M. Perkowitz, and O. Etzioni. **Towards Adaptive Web Sites: Conceptual Framework and Case Study**. Artificial Intelligence Journal, 1999.

M. Perkowitz, R. Doorenbos, O. Etzioni, and D. Weld. **Learning to Understand Information on the Internet: an Example-Based Approach.** Journal of Intelligent Information Systems. Vol. 8, num. 2, 1997. pp. 133-153.

E. Charniak, C. Hendrickson, N. Jacobson, and M. Perkowitz. **Equations for Part-of-Speech Tagging.** Proceedings of AAAI 1993. pp. 784-789.