

# Margaret Mitchell

Microsoft Research  
Cognition Group  
Redmond, WA U.S.A.  
margarmitchell@gmail.com · (443) 955 3527

## Research Areas

Artificial/Machine Intelligence, Vision & Language, Assistive/Clinical Technology, Natural Language Generation, Natural Language Processing, Referring Expression Generation/Object Reference, Cognitive Modeling, Named Entity Recognition, Semantic Roles, Sentiment, Syntax

## Work

RESEARCHER 2016

MICROSOFT RESEARCH, COGNITION GROUP

*Manager: Pushmeet Kohli*

Focusing on advancing AI to maximally benefit people. Working on image description, storytelling, and visual descriptions for the visually impaired.

RESEARCHER 2013 - 2015

MICROSOFT RESEARCH, NLP GROUP

*Manager: Bill Dolan*

Worked on natural language processing for Bing & Cortana, focusing on microsummarization. Advanced conversation generation, image description. Led clinical NLP research.

POSTDOCTORAL RESEARCHER 2012 - 2013

JOHNS HOPKINS UNIVERSITY

*Supervisor: Benjamin Van Durme*

Working with graphical models for semantic role labeling, named entity recognition, and sentiment detection. With Matt Gormley, developed the Pacaya<sup>a</sup> graphical modelling toolkit.

JOHNS HOPKINS WORKSHOP GRADUATE Summer

CENTER FOR LANGUAGE AND SPEECH PROCESSING 2011

*Supervisors: Tamara Berg, Alex Berg*

Built generation system that takes vision output as input, producing syntactically/semantically well-formed descriptions of images. Worked on automatic classification of visually descriptive text, and characterizing description.

VISITING SCHOLAR 2009 - 2012

CENTER FOR SPOKEN LANGUAGE UNDERSTANDING

*Supervisors: Brian Roark, Richard Sproat*

Developed system to generate more personal language for teenagers with cerebral palsy. Completed work on system to automatically diagnose Alzheimer's.

RESEARCH ASSOCIATE/ASSISTANT 2005 - 2009

CENTER FOR SPOKEN LANGUAGE UNDERSTANDING

*Supervisors: Brian Roark, John-Paul Hosom, Jan van Santen*

In charge of all transcription work for projects on Alzheimer's, Aphasia, Apraxia, Autism, Dysarthria, Parkinson's. Developed system to automatically diagnose Alzheimer's by extracting syntactic features from narratives.

---

<sup>a</sup><https://github.com/mgormley/pacaya>

<b>Education</b>	UNIVERSITY OF ABERDEEN PHD: COMPUTING SCIENCE Doctoral Thesis: Generating Reference to Visible Objects <i>Advisors: Kees van Deemter, Ehud Reiter</i> Created an algorithm that aims to generate human-like reference to identify real world objects. Examined the factors affecting reference in visual settings, including the production of color and size words, visual perception, and stored object representations. Developed end-to-end vision to language system that provides descriptions of images.	2009 - 2012
	UNIVERSITY OF WASHINGTON MA/MS: COMPUTATIONAL LINGUISTICS Master's Thesis: Towards the Generation of Natural Reference <i>Advisors: Scott Farrar, Emily Bender</i> Introduced an algorithm to generate referring expressions comparable to human-produced expressions, and developed a class-based system for the prenominal ordering of modifiers.	2007 - 2008
	REED COLLEGE BA: LINGUISTICS ALLIED FIELD: PSYCHOLOGY Senior Thesis: On the Generation of Referring Expressions <i>Advisors: John Haviland, Matt Pearson</i> Critiqued the problem of generating human-like reference from a knowledge base.	2001 - 2005

**Publications 2016**

Mostafazadeh, N. and Misra, I. and Devlin, J. and Zitnick, C. L. and Mitchell, M. and He, X. and Vanderwende, L. (2016). Generating Natural Questions About an Image. *Proceedings of ACL 2016*.

Huang, T.-H., and Ferraro, F., and Mostafazadeh, N. and Misra, I. and Agrawal, A. and Devlin, J. and Girshick, R. and He, X. and Kohli, P. and Batra, D. and Zitnick, C. L. and Parikh, Devi and Vanderwende, L. and Galley, M. and Mitchell, M. (2016). Visual Storytelling. *Proceedings of NAACL 2016*.

Misra, I. and Zitnick, C. L. and Mitchell, M. and Girshick, R. (2016). Seeing through the Human Reporting Bias: Visual Classifiers from Noisy Human-Centric Labels. *Proceedings of CVPR 2016*.

Mason, R. and Gaska, B. and Van Durme, B. and Choudhury, P. and Hart, T. and Dolan, B. and Toutanova, K. and Mitchell, M. (2016). Microsummarization of Online Reviews: An Experimental Study. *Proceedings of AAAI 2016*.

**2015**

Antol, S. and Agrawal, A. and Lu, J. and Mitchell, M. and Batra, D. and Zitnick, C. L. and Parikh, D. (2015). VQA: Visual Question Answering. *Proceedings of ICCV 2015*.

Ferraro, F. and Mostafazadeh, N. and Huang, T. and Vanderwende, L. and Devlin, J. and Galley, M. and Mitchell, M. (2015). A Survey of Current Datasets for Vision and Language Research. *Proceedings of EMNLP 2015*.

Devlin, J. and Cheng, H. and Fang, H. and Gupta, S. and Deng, L. and He, X. and Zweig, G. and Mitchell, M. (2015). Language Models for Image Captioning: The Quirks and What Works. *Proceedings of ACL 2015*.

Galley, M. and Brockett, C. and Sordoni, A. and Ji, Y. and Auli, M. and Quirk, C. and Mitchell, M. and Gao, J. and Dolan, B. (2015). deltaBLEU: A Discriminative Metric for Generation Tasks with Intrinsically Diverse Targets. *Proceedings of ACL 2015*.

- Publications cont.** Devlin, J. and Gupta, S. and Girshick, R. and Mitchell, M. and Zitnick, C. L. (2015). Exploring Nearest Neighbor Approaches for Image Captioning. *ArXiv print*.
- Sordoni, A. and Galley, M. and Auli, M. and Brockett, C. and Ji, Y. and Mitchell, M. and Nie, J. and Gao, J. and Dolan, B. (2015). A Neural Network Approach to Context-Sensitive Generation of Conversational Responses. *Proceedings of NAACL 2015*.
- Mitchell, M. and Hollingshead, K. and Coppersmith G. (2015). Quantifying the Language of Schizophrenia in Social Media. *Proceedings of the 2nd CLPsych Workshop, NAACL 2015*.
- Coppersmith, G. and Dredze, M. and Harman, C. and Hollingshead, K. and Mitchell, M. (2015). CLPsych 2015 Shared Task: Depression and PTSD on Twitter. *Proceedings of the 2nd CLPsych Workshop, NAACL 2015*.
- Fang, H. and Gupta, S. and Iandola, F. and Srivastava, Rupesh K. and Deng, L. and Dollar, P. and Gao, J. and He, X. and Mitchell, M. and Platt, J. C. and Zitnick, L. and Zweig, G. (2015). From Captions to Visual Concepts and Back. *Proceedings of CVPR 2015*.
- 2014**
- Mitchell, M. and Bohus, D. and Kamar, E. (2014). Crowdsourcing Language Generation Templates for Dialogue Systems. *Proceedings of INLG and SigDial 2014*.
- Gormley, M., and Mitchell, M., and Van Durme, B., and Dredze, M. (2014). Low Resource Semantic Role Labeling. *Proceedings of ACL 2014*.
- Beller, C., and Knowles, R., and Harman, C., and Bergsma, S., and Mitchell, M., and Van Durme, B. (2014). I'm a Belieber: Social Roles via Self-identification and Conceptual Attributes. *Proceedings of ACL 2014*.
- 2013**
- Mitchell, M., and Aguilar, J., and Wilson, T., and Van Durme, B. (2013). Open Domain Targeted Sentiment. *Proceedings of EMNLP 2013*.
- Mitchell, M., and van Deemter, K., and Reiter, E. (2013). Attributes in Visual Reference. *Proceedings of PRE-CogSci 2013*.
- Viethen, J., and Mitchell, M., and Krahmer, E. (2013). Graphs and Spatial Relations in the Generation of Referring Expressions. *Proceedings of ENLG 2013*.
- Mitchell, M. and Reiter, E., and van Deemter, K. (2013). Typicality and Object Reference. *Proceedings of CogSci 2013*.
- Mitchell, M. and van Deemter, K., and Reiter, E. (2013). Generating Expressions that Refer to Visible Objects. *Proceedings of NAACL 2013*.
- Mitchell, M. Generating Reference to Visible Objects. *Doctoral Thesis*.
- 2012**
- Mitchell, M., Dodge, J., Goyal, A., Han, X., Stratos, K., Mensch, A., Yamaguchi, K., Daumé III, H., Berg, A., and Berg, T. L. (2012). Midge: Generating Descriptions of Images. *Proceedings of INLG 2012*. System demo.
- Dodge, J., Goyal, A., Han, X., Stratos, K., Mensch, A., Mitchell, M., Yamaguchi, K., Choi, Y., Daumé III, H., Berg, A., and Berg, T. L. (2012). Detecting Visual Text. *Proceedings of NAACL 2012*.

**Publications cont.** Stratos, K., Sood, A., Mensch, A., Han, X., Mitchell, M., Yamaguchi, K., Dodge, J., Goyal, A., Daumé III, H., Berg, A., and Berg, T. L. (2012). Understanding and Predicting Importance in Images. *Proceedings of CVPR 2012*.

Mitchell, M., Dodge, J., Goyal, A., Yamaguchi, K., Stratos, K., Han, X., Mensch, A., Berg, A., and Berg, T. L., Daumé III, H. (2012). Midge: Generating Image Descriptions From Computer Vision Detections. *Proceedings of EACL 2012*.

## 2011

Mitchell, M., van Deemter, K., and Reiter, E. (2011). Two Approaches for Generating Size Modifiers. *Proceedings of ENLG 2011*.

Mitchell, M. (2011). From an Image to a Description. *Vision and Language Workshop 2011*.

Mitchell, M., Dunlop, A., and Roark, B. (2011). Semi-Supervised Modeling for Prenominal Modifier Ordering. *Proceedings of ACL 2011*.

Mitchell, M., van Deemter, K., and Reiter, E. (2011). On the Use of Size Modifiers When Referring to Visible Objects. *Proceedings of CogSci 2011*.

Mitchell, M., van Deemter, K., and Reiter, E. (2011). Applying Machine Learning to the Choice of Size Modifiers. *Proceedings of PRE-CogSci 2011*.

Prud'hommeaux, E. T., Mitchell, M., and Roark, B. (2011). Using Patterns of Narrative Recall for Improved Detection of Mild Cognitive Impairment. *Proceedings of FICCDAT 2011*.

Roark, B., Mitchell, M., Hosom, J., Hollingshead, K., and Kaye, J. (2011). Spoken Language Derived Measures for Detecting Mild Cognitive Impairment. *IEEE Transactions on Audio, Speech, and Language Processing*.

## 2010

Mitchell, M., van Deemter, K., and Reiter, E. (2010). Natural Reference to Objects in a Visual Domain. *Proceedings of INLG 2010*.

Mitchell, M. (2010). A Flexible Approach to Class-Based Ordering of Prenominal Modifiers. Krahmer, E. and Theune, M., editors, *Empirical Methods in Natural Language Generation*, volume 5980 of *Lecture Notes in Computer Science*, Springer, Berlin/Heidelberg.

Dunlop, A., Mitchell, M., and Roark, B. (2010). Prenominal Modifier Ordering via Multiple Sequence Alignment. *Proceedings of NAACL-HLT 2010*.

## 2009

Mitchell, M. (2009). Class-Based Ordering of Prenominal Modifiers. *Proceedings of the 12th European Workshop on Natural Language Generation (ENLG 2009)*.

**2008** - Master's Thesis

## 2007

Roark, B., Mitchell, M. and Hollingshead, K. (2007). Syntactic Complexity Measures for Detecting Mild Cognitive Impairment. *Proceedings of the ACL 2007 Workshop on Biomedical Natural Language Processing (BioNLP 2007)*.

Roark, B., Hosom, P., Mitchell, M. and Kaye, J. (2007). Automatically Derived Spoken Language Markers for Detecting Mild Cognitive Impairment. *Proceedings of the 2nd International Conference on Technology and Aging (ICTA 2007)*.

<b>Recent Projects</b>	TECHNICAL & RESEARCH LEAD, Seeing AI	2015 - 2016
	Leveraged image captioning work to produce visual descriptions for the visually impaired. Helped stream technology through cloud services to land on user's mobile device or smartglasses. Connected groups throughout Microsoft to move forward as one: MSR Labs, Cognitive Services, Windows, Narrator, Bing, Garage, Outreach, TExT.	
	MICROSOFT STRATEGIC INTERN PROJECT LEAD, Visual Storytelling	2015 - 2016
	Led group of 8 researchers, 2 visiting professors, and 5 students on summer project to drive forward research on selecting key frames in a photo album and generating a coherent story for them. Project resulted in full working system, with all 4 full-time interns first-authoring a paper at a top-tier conference.	
	SENIOR RESEARCHER, JSALT Summer Workshop 2016	2015 - 2016
	Exploring how we can utilize shared patient data alongside shared patient social media feeds to monitor PTSD and depression.	
	TOP DISHES RESEARCHER, Cortana and Bing Local	2014 - 2015
	Developed end-to-end system to mine social media for positive/negative sentiment expressed towards most common items.	
<b>Chairing &amp; Community Leading</b>	GUEST EDITOR, the International Journal of Computer Vision, Combined Image and Lanugage Understanding Special Issue 2016	2015 - 2016
	AREA CHAIR, Generation ACL 2016	2015 - 2016
	PUBLICATION CHAIR & AREA CHAIR, Generation, NAACL 2016	2015 - 2016
	As Publication Chair, introduced new format inspired by CVPR/ICCV ruler style. Modernized the Word template to remove the typesetting program as a factor in reviewing.	
	AREA CHAIR, GENERATION & SUMMARIZATION, EMNLP 2015	2014 - 2015
	GENERAL, PROGRAM & PUBLICATION CHAIR, Computational Linguistics and Clinical Psychology Workshop (CLPsych) 2015 at NAACL 2015	2014 - 2015
	GENERAL, PROGRAM & PUBLICATION CO-CHAIR, Computational Linguistics and Clinical Psychology Workshop (CLPsych) 2014 at ACL 2014	2013 - 2014
	GENERAL, PROGRAM & PUBLICATION CHAIR, the International Natural Language Generation Conference (INLG) 2014	2013 - 2014
	KNOWLEDGE BASE POPULATION - SENTIMENT SLOT FILLING TRACK OWNER, Text Analysis Conference 2014	2013 - 2014
	Defined and organized task for the National Institute of Standards and Technology (NIST) Text Analysis Conference (TAC) to predict sentiment between knowledge base entities.	
STUDENT REPRESENTATIVE, Special Interest Group on Generation (SIGGEN)	2010 - 2012	
Helped move NLG community forward, served as webmaster.		
<b>Reviewing</b>	<b>Journals</b>	
	IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)	2016
	The Journal of Artificial Intelligence Research (JAIR)	2012, 2015
	<b>Conferences</b>	
	Association for the Advancement of Artificial Intelligence (AAAI)	2016
	Empirical Methods in Natural Language Processing (EMNLP)	2013 - 2015

<b>Reviewing cont.</b>	The Annual Meeting of the Association for Computational Linguistics (ACL)	2013 - 2016
	The North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)	2013 - 2016
	The International Natural Language Generation Conference (INLG)	2012 - 2014
	The European Workshop on Natural Language Generation (ENLG)	2013 - 2015
	iV&L Net Workshop on Language for Vision (VL, L4V)	2013 - 2015
	Workshop on the Production of Referring Expressions (PRE-CogSci)	2013
	Workshop on Vision and Language (WVL)	2013
	The Scottish Informatics and Computer Science Alliance PhD Conference (SICSA)	2010

**Systems & Languages**  
 Program in Python, Bash, C#, Java. Experience with Tcl, C++, HTML, JavaScript, PHP, sed, awk, etc.  
 Comfortable with Windows, MacOS 9 and X, Linux (Ubuntu, Fedora), all MS Office programs, all OpenOffice programs, L<sup>A</sup>T<sub>E</sub>X, Praat, Wavesurfer, version control systems, NLTK, WordNet, etc.

**Recent Press**  
 Seeing AI announcement: <https://www.youtube.com/watch?v=rVF2duPVUTY>  
<http://www.fastcodesign.com/3058905/the-real-reason-microsoft-is-building-so-many-computer-vision-apps>  
<https://www.technologyreview.com/s/601339/will-artificial-intelligence-win-the-caption-contest>  
<http://venturebeat.com/2016/04/14/microsoft-ai-visual-storytelling>  
<http://blogs.microsoft.com/next/2016/04/14/teaching-computers-to-describe-images-as-people-would>  
<http://blogs.microsoft.com/next/2016/03/30/decades-of-computer-vision-research-one-swiss-army-knife>

**Personal Details**  
 Citizenship: U.S.  
 Residence: Seattle, Washington  
 Date of birth: 18th November, 1983  
 Hobbies: Cooking, running, political analysis. Occasionally all at once.  
 Nerdy obsession with collecting vinyl records: Yes