

# Enabling Robots to Draw and Tell: Towards Visually Grounded Multimodal Description Generation

Ting Han  
ting.han@aist.go.jp  
Artificial Intelligence Research Center  
AIST, Tokyo, Japan

Sina Zarriëß  
sina.zarriess@uni-jena.de  
Friedrich Schiller University Jena  
Jena, Germany

## ABSTRACT


Socially competent robots should be equipped with the ability to perceive the world that surrounds them and communicate about it in a human-like manner. Representative skills that exhibit such ability include generating image descriptions and visually grounded referring expressions. In the NLG community, these generation tasks are largely investigated in non-interactive and language-only settings. However, in face-to-face interaction, humans often deploy multiple modalities to communicate, forming seamless integration of natural language, hand gestures and other modalities like sketches. To enable robots to describe what they perceive with speech and sketches/gestures, we propose to model the task of generating natural language together with free-hand sketches/hand gestures to describe visual scenes and real-life objects, namely, visually-grounded multimodal description generation. In this paper, we discuss the challenges and evaluation metrics of the task, and how the task can benefit from progress recently made in the natural language processing and computer vision realms, where related topics such as visually grounded NLG, distributional semantics, and photo-based sketch generation have been extensively studied.

## 1 INTRODUCTION

A key way for robots to exhibit social intelligence is to show the ability of communicating what they perceive in the surrounding world as humans do, with both verbal and non-verbal means such as drawing a trajectory in shared space or on a piece of paper (i.e., gesturing and sketching) to illustrate the contour of an object. Research in human-robot interaction has a long-standing interest in embodied multimodal interaction. A lot of work has been done on studying the complex interplay between speech and gestures, as well as understanding and generating speech and symbolic gestures in interaction [5, 6, 10–12, 15, 20]. In comparison, much less work has focused on generating iconic gestures/sketches together with natural language to describe real-life objects, albeit such descriptions are multimodal in nature.

Generating visually grounded multimodal descriptions is a challenging task: 1) It requires the generation of both natural language and gesture/sketch, each of which is a standalone challenging task. Given a photo, the NLG component generates a sequence of words to convey the content symbolically; The gesture/sketch generation component translates the content into a sequence of sketch strokes to communicate iconically. It decides what, where and in which order to draw. 2) The two modalities are not independent, but bear close semantic and temporal relations [16]. Iconic gestures complement or supplement the semantic content of the accompanied speech through their formal relevance to referents in the speech, thus they do not exhibit meanings on their own. Instead, their

meanings are largely designated by the accompanied speech, making multimodal alignment reasoning a critical component of the generation task. 3) Moreover, in situated interaction, the generated descriptions unfold as timing sequences. The durations of a gesture/sketch and the accompanied utterance must fit with each other, in order to be temporally aligned. Therefore, generating aligned descriptions requires reasoning of both the semantics and the timing of speech and gestures/sketches, making the generation task extremely challenging.



**Figure 1: A pseudo multimodal description of an elephant, composed of an utterance and a sketch stroke selected from a full sketch. The red arrows indicate the sketch direction.**

For instance, in Figure 1, the utterance “*elephant, facing right, trunk coiled towards mouth*” and the sketch compose a multimodal description of the elephant in the photo. While the utterance describes the object symbolically, the sketch illustrates the shape of the trunk by visualizing its contour [4, 8], which is difficult to convey symbolically with the word “coiled”. Yet, the meaning of the sketch is ambiguous on its own. It only becomes evident when jointly interpreting the utterance and the sketch, grounding the iconic concepts encoded in both modalities to the same region of the photo (i.e., the trunk).

With the above observation in mind, we propose to decompose the multimodal generation task into three sub-tasks: visually-grounded NLG, visually-grounded gesture/sketch generation, and multimodal alignment reasoning. The first two sub-tasks generate utterances and full sketches respectively. The multimodal alignment sub-task selects representative sketch strokes from full sketches, infers optimal semantic and temporal relations between the utterances and the sketches, then outputs orchestrated descriptions. Finally, we convert the utterances to speech and pair the utterances with sketches displayed on a screen or projected to a gesture space together and duly to obtain aligned multimodal descriptions.

Such a formulation will also benefit the proposed task from available pre-trained large scale language models, sketch generation models, and large scale datasets, e.g., language and vision datasets such as MSCOCO and Visual Genome [9, 13], and free-hand sketching dataset such as QuickDraw [18]. To our knowledge, no large scale multimodal description datasets have been publicly available yet. Utilizing existing datasets is a good starting point for the proposed task.

## 2 THE CHALLENGES

### 2.1 Visually Grounded NLG

The visually grounded NLG sub-task not only generates verbal object descriptions, it also provides important information for multimodal alignment. As aforementioned, sketch/gestures convey meaning through their formal relevance to referents in accompanied speech. Thus, we assume words and sketches that describe the same photo region are likely to align on semantic level. Grounding generated words to visual inputs enables the identification of the region they describe.

Attention-based image caption models are a promising solution to our generation task. With an attention mechanism, image caption models focus on the relevant part of a given photo when generating each word. The “focus” is represented as a weight matrix: The region a word describes is with higher attention weights ([14, 22]). With the attention matrix of each word, we identify the region an utterance describes, then reason out sketch strokes that can be aligned with the utterance (see discussion in Section 2.3).

### 2.2 Photo-based Sketch Generation

The sketch generation sub-task generates a sequence of sketch strokes that exhibit contours of objects in a given photo (as shown in Figure 1), mimicking human sketching process by resembling various levels of sophistication and abstractness.

In computer vision, photo-based sketch generation is posited as a weakly supervised or unsupervised photo-to-sequence translation task [3, 19, 23, 24]. Photo-sketch pairs are not strictly required for training sketch generation models. In recent works, neural network sketching models with encoding-decoding architecture typically learn an encoder network and a decoder network: the former encodes photos into feature vectors; the latter decodes photo vectors to sketch sequences of target objects.

Note that existing neural sketching models generate detailed sketches of objects, as they are trained on sketching datasets collected without timing constraints which might be too sophisticated to draw in situated interaction. In our task, we aim to accompany utterances with sketch strokes in human-robot interaction, where communication is often under timing pressure. Hence, a challenge to be addressed here is to constrain the sophistication of sketch strokes either during model training or via post-editing, so that they suit for situated communication.

### 2.3 Composing Multimodal Descriptions

To produce orchestrated descriptions, the multimodal alignment sub-task first selects sketch strokes that represent salient iconic concepts of the target object, then determines which words are semantically and temporally compatible with the selected strokes. Several candidate descriptions might be initially proposed, from which the one that best suits the communicating purpose (e.g., emphasize a particular concept) is selected as the optimal description.

**Selecting representative sketch strokes.** Given a full sketch of a photo, we select one or two sketch strokes to form an optimal multimodal description. The selection strategy here depends on the communication intention. For instance, to supplement a particular utterance, sketch strokes in the region the utterance describes

should be selected; To be most informative, sketch strokes of the most salient part of the target object should be selected.

**Multimodal alignment reasoning.** Regarding multimodal alignment, although word attention weights indicate which words correlate with selected strokes, note that not every word in natural language correlates with visual features in theory, hence, the attention weights only provide a very rough estimation of the alignment. To achieve more accurate alignment, we propose to reason the alignment according to semantics of words and sketches, as accompanied speech and gestures bear close semantic relations [7]. Works in *Distributional semantics* have shown that content from different modalities can be mapped to a joint embedding space, where vectors with within shorted distances share similar semantic meaning [2]. Clearly, iconic gestures and sketches are different, although they both convey iconicity. This approach could also be applied to measure semantic similarity of utterances and sketches.

In a pilot study, we generated several multimodal descriptions with a simplified approach based on the above framework, and evaluated them under two settings: 1) display speech (converted from text with a TTS software) and sketches on a monitor, 2) enable a Nao robot to speak and gesture by projecting generated sketch strokes into Nao’s gesture space. We found that, with right alignment, sketches/gestures are helpful in conveying iconicity. Iconic gestures transformed from shorter and less sophisticated sketches are easier to interpret and more helpful. Generating iconic gestures based on sketches would be an interesting topic for future research.

### 2.4 Evaluation

An informative evaluation should reflect the quality of individual modalities and how well the two modalities are aligned. We propose to combine human evaluation with automatic evaluation to assess the overall generation quality [1].

**Automatic evaluation** suits for fast evaluations during model development. We suggest to evaluate NLG quality with metrics in NLG such as BLEU[17, 21], while evaluating the informativeness of sketches and multimodal descriptions with sketch-based and multimodal image retrieving tasks: A higher retrieving accuracy indicates better description quality. Note that these automatic evaluations can not measure temporal alignment between modalities, therefore, we resort to human evaluation to assess multimodal alignment quality. **Human evaluation** offers insights into multimodal alignment quality as miss-alignment impairs overall interpretation of the descriptions, especially the interpretation of sketches/gestures. Interactive referring games between humans and robots that reveal how humans understand the generated descriptions, are good testbeds for such evaluation [9].

## 3 CONCLUSION AND FUTURE WORK

We proposed the task of visually grounded multimodal description generation. Along with discussions on the challenges regarding NLG, free-hand sketch generation, multimodal alignment, and evaluation methods, we pointed out how existing works of visually grounded NLG, photo-based sketch generation, distributional semantics, and their respective evaluation metrics benefit the proposed task. We believe, solving the task will lead to more natural multimodal human-robot interaction in scenarios where robots can communicate what they perceive in the world that surrounds them.

## ACKNOWLEDGMENTS

This work is supported by the New Energy and Industrial Technology Development Organization (NEDO) of Japan.

## REFERENCES

- [1] Anja Belz and Ehud Reiter. 2006. Comparing automatic and human evaluation of NLG systems. In *11th Conference of the European Chapter of the Association for Computational Linguistics*.
- [2] Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research* 49 (2014), 1–47.
- [3] Mathias Eitz, James Hays, and Marc Alexa. 2012. How do humans sketch objects? *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.
- [4] Ting Han, Julian Hough, and David Schlangen. 2017. Natural Language Informs the Interpretation of Iconic Gestures. A Computational Approach. In *The 8th International Joint Conference on Natural Language Processing. Proceedings of the Conference. Vol. 2: Short Papers*.
- [5] Ting Han, Casey Kennington, and David Schlangen. 2015. Building and Applying Perceptually-Grounded Representations of Multimodal Scene Descriptions. In *Proceedings of the 19th SemDial Workshop on the Semantics and Pragmatics of Dialogue (goDIAL)*.
- [6] Ting Han, Casey Kennington, and David Schlangen. 2018. Placing Objects in Gesture Space: Toward Real-Time Understanding of Spatial Descriptions. In *Proceedings of the thirty-second AAAI conference on artificial intelligence (AAAI18)*.
- [7] Ting Han and David Schlangen. 2017. Draw and tell. multimodal descriptions outperform verbal-or sketch-only descriptions in an image retrieval task. In *The 8th International Joint Conference on Natural Language Processing. Proceedings of the Conference. Vol. 2: Short Papers*.
- [8] Ting Han, Sina Zarrieß, Kazunori Komatani, and David Schlangen. 2018. Learning to describe multimodally from parallel unimodal data? A pilot study on verbal and sketched object descriptions. In *Proceedings of the 22nd Workshop on the Semantics and Pragmatics of Dialogue (AixDial)*.
- [9] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 787–798.
- [10] Stefan Kopp, Kirsten Bergmann, and Ipke Wachsmuth. 2008. Multimodal communication from multimodal thinking—towards an integrated model of speech and gesture production. *International Journal of Semantic Computing* 2, 01 (2008), 115–136.
- [11] Stefan Kopp, Brigitte Krenn, Stacy Marsella, Andrew N Marshall, Catherine Pelachaud, Hannes Pirker, Kristinn R Thorisson, and Hannes Vilhjalmsson. 2006. Towards a common framework for multimodal generation: The behavior markup language. *Proceedings Intelligent Virtual Agents 2006* 4133 (2006).
- [12] Alfred Kranstedt and Ipke Wachsmuth. 2005. Incremental generation of multimodal deixis referring to objects. In *Proceedings of the Tenth European Workshop on Natural Language Generation (ENLG-05)*.
- [13] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [14] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. 2017. Knowing When to Look: Adaptive Attention via A Visual Sentinel for Image Captioning. *CVPR*.
- [15] Andy Lücking, Kirsten Bergmann, Florian Hahn, Stefan Kopp, and Hannes Rieser. 2010. The Bielefeld speech and gesture alignment corpus (SaGA). In *LREC 2010 workshop: Multimodal corpora—advances in capturing, coding and analyzing multimodality*.
- [16] David McNeill. 1992. *Hand and mind: What gestures reveal about thought*. University of Chicago press.
- [17] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [18] Patsorn Sangkloy, Nathan Burnell, Cusuh Ham, and James Hays. 2016. The sketchy database: learning to retrieve badly drawn bunnies. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–12.
- [19] Jifei Song, Kaiyue Pang, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. 2018. Learning to sketch with shortcut cycle consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 801–810.
- [20] Timo Sowa and Ipke Wachsmuth. 2009. A computational model for the representation and processing of shape in coverbal iconic gestures. In *Spatial Language and Dialogue*.
- [21] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4566–4575.
- [22] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*. 2048–2057.
- [23] Peng Xu. 2020. Deep learning for free-hand sketch: A survey. *arXiv preprint arXiv:2001.02600* (2020).
- [24] Liliang Zhang, Liang Lin, Xian Wu, Shengyong Ding, and Lei Zhang. 2015. End-to-end photo-sketch generation via fully convolutional representation learning. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval*. 627–634.