

Speech2Properties2Gestures: Gesture-Property Prediction as a Tool for Generating Representational Gestures from Speech

Taras Kucherenko*
tarask@kth.se

Rajmund Nagy†
rajmundn@kth.se

Patrik Jonell†
pjonell@kth.se

Michael Neff‡
mpneff@ucdavis.edu

Hedvig Kjellström*
hedvig@kth.se

Gustav Eje Henter†
ghe@kth.se

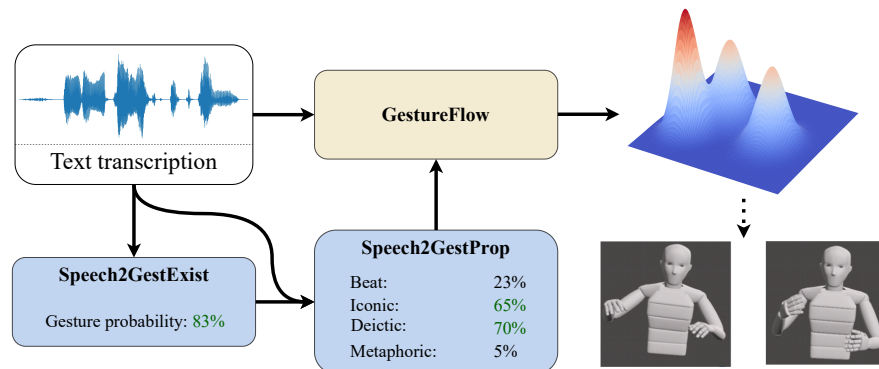


Figure 1: Overview of the proposed framework. We first use speech text and audio to predict whether or not the agent should gesture. After that, we predict several gesture properties, such as gesture type. Finally, gestures are generated by a probabilistic model (e.g., a normalizing flow) conditioned on text, audio, and predicted gesture properties together.

ABSTRACT

We propose a new framework for gesture generation, aiming to allow data-driven approaches to produce more semantically rich gestures. Our approach first predicts whether to gesture, followed by a prediction of the gesture properties. Those properties are then used as conditioning for a modern probabilistic gesture-generation model capable of high-quality output. This empowers the approach to generate gestures that are both diverse and representational. Follow-ups and more information can be found on the project page: <https://svito-zar.github.io/speech2properties2gestures/>

CCS CONCEPTS

• **Human-centered computing** → *Human computer interaction (HCI)*; • **Computing methodologies** → **Animation**.

KEYWORDS

gesture generation, virtual agents, representational gestures

*Robotics, Perception, and Learning, KTH Royal Institute of Technology, Sweden.

†Speech, Music, and Hearing, KTH Royal Institute of Technology, Sweden.

‡University of California, Davis, United States.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

IVA '21, September 14–17, 2021, Virtual Event, Japan

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8619-7/21/09.

<https://doi.org/10.1145/3472306.3478333>

ACM Reference Format:

Taras Kucherenko, Rajmund Nagy, Patrik Jonell, Michael Neff, Hedvig Kjellström, and Gustav Eje Henter. 2021. Speech2Properties2Gestures: Gesture-Property Prediction as a Tool for Generating Representational Gestures from Speech. In *21th ACM International Conference on Intelligent Virtual Agents (IVA '21)*, September 14–17, 2021, Virtual Event, Japan. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3472306.3478333>

1 INTRODUCTION AND BACKGROUND

A large part of human communication is non-verbal [10] and often takes place through co-speech gestures [9, 19]. Co-speech gesture behavior in embodied agents has been shown to help with learning tasks [5] and lead to greater emotional response [24]. Gesture generation is hence an important part of both automated character animation and human-agent interaction.

Early dominance of rule-based approaches [6, 12, 18, 21] has been challenged by data-driven gesture generation systems [1, 4, 7, 14, 20, 26, 27]. These latter systems first only considered a single speech modality (either audio or text) [4, 13, 20, 27], but are now shifting to use both audio and text together [1, 14, 26].

While rule-based systems provide control over the communicative function of output gestures, they lack variability and require much manual effort to design. Data-driven systems, on the other hand, need less manual work and are very flexible, but most existing systems do not provide much control over communicative function and generated gestures have little relation to speech content [15].

This paper continues recent efforts to bridge the gap between the two paradigms [7, 23, 29]. The most similar prior work is Yunus et al. [29] where gesture timing and duration were predicted based

Label	Gesture category [Macro F ₁]				Gesture semantics [Macro F ₁]				Gesture phase [F ₁]				
	deictic	beat	iconic	discourse	amount	shape	direction	size	pre-hold	post-hold	stroke	retr.	prep.
Relative frequency	29.05%	14.47%	72.03%	12.78%	4.7%	13.1%	13.7%	1.9%	0.6%	12.2%	40.9%	14.8%	30.8%
RandomGuess	50% ± 2%	50% ± 2%	50% ± 1.5%	50% ± 2%	49% ± 1%	49% ± 2%	49% ± 2%	50% ± 1%	1.3% ± 4%	12% ± 4%	42% ± 4%	14% ± 5%	30% ± 3%
ProposedModel	60% ± 6%	53% ± 6%	63% ± 5%	59% ± 7%	63% ± 8%	65% ± 6%	62% ± 8%	59% ± 9%	0.5% ± 1.3%	23% ± 12%	47% ± 10%	25% ± 5%	45% ± 6%

Table 1: Gesture-property prediction scores for random guessing and our trained predictors using text and audio modalities together. Bold, coloured numbers indicate that the label in question can be predicted better than chance.

on acoustic features only. The method proposed here differs from their approach in three ways: 1) it considers not only audio but also text as input; 2) it models not only gesture phase, but multiple gesture properties; 3) it also provides a framework for integrating these gesture properties in a data-driven gesture-generation system.

The proposed approach helps decouple different aspects of gesticulation and can leverage database information about gesture timing and content with modern, high-quality data-driven animation.

2 PROPOSED MODEL

Our unified model uses speech text and audio as input to generate gestures as a sequence of 3D poses. As depicted in Figure 1, it is composed of three neural networks:

- (1) **Speech2GestExist**: A temporal CNN which takes speech as input and returns a binary flag indicating if the agent should gesture (similar to [28]);
- (2) **Speech2GestProp**: A temporal CNN which takes speech as input and predicts a set of binary gesture properties, such as gesture type, gesture phase, etc.;
- (3) **GestureFlow**: A normalizing flow [11] that takes both speech and predicted gesture properties as input, and describes a probability distribution over 3D poses, from which motion sequences can be sampled.

In this study, we experiment with the first two neural networks only. We implemented the *Speech2GestProp* and *Speech2GestExist* components using dilated CNNs. Their inputs are sequences of aligned speech text and audio frames, and they return a binary vector of gesture properties (for *Speech2Prop*) or a binary flag of gesture existence (for *Speech2GestExist*) as its output. By sliding a window over the speech and predicting poses, frame-by-frame properties are generated at 5 fps. Text features were extracted using DistilBERT [22]. Audio features were log-scaled mel-spectrograms.

3 PRELIMINARY RESULTS

Dataset. We evaluated our model on the SaGA direction-giving dataset [17] designed to contain many representational gestures. The dataset contains audio/video recordings of 25 participants (all German native speakers) describing the same route to other participants and includes detailed annotations of gesture properties.

We considered the following three gesture properties: 1) *Phase* (preparation, pre-stroke hold, stroke, post-stroke hold, and retraction); 2) *Type* (deictic, beat, iconic [19], and discourse); 3) *Semantic information* (amount, shape, direction, size, as described in [3]).

Experimental Results. For each of our experiments we calculated the mean and standard deviation of the F₁ score across 20-fold cross-validation. The F₁ score is preferable over accuracy here since the data is highly unbalanced and accuracy does not represent overall

performance well. For gesture category and phase we report Macro F₁ score [25], since those properties are not mutually exclusive.

First we validated that gesture presence can be predicted from the speech in our dataset. We achieved a 70% ± 3.7% Macro F₁ score for binary classification, which aligns with previous work [28].

Next, we experimented with predicting gesture properties. Table 1 contains results for predicting the gesture category, gesture semantic information, and gesture phase from speech text and audio. We can see that this is a challenging task, but we are still able to predict most of the values better than chance. This was unexpected given how complex gesture semantics tend to be and could be due to the focused scope of the direction-giving task. For a deeper study with more results and analyses, please see the follow-up work [16].

4 DISCUSSION

In this section we discuss the feasibility of the proposed approach. Our proposal to use probabilistic models (especially normalizing flows) is inspired by a recent application of MoGlow [8] to perform gesture synthesis by Alexanderson et al. [2]. They showed that such models can be seamlessly conditioned on various kinematic gesture properties (such as speed, range, and hand height), suggesting that it is possible to condition gestures on semantic properties as well.

We obtained good results for the gesture-property prediction part of our proposed system, as described in Section 3. Since we can predict several important properties with F₁ scores significantly above chance level, we believe that our predictions are reasonable and will be useful for more appropriate gesture synthesis.

Our two-stage approach lets the machine learning model leverage additional information (such as detailed annotation) about human gestures. It also allows direct control of gesture frequency, by adjusting the threshold on the output of *Speech2GestExist* needed to trigger a gesture. Finally, it helps the model learn from small datasets, since each sub-module has a more straightforward task than learning everything at once and also can be trained separately.

5 CONCLUSION

We presented a novel gesture generation framework aiming to bridge the semantic gap between rule-based and data-driven models. Our method first predicts if a gesture is appropriate for a given point in the speech and what kind of gesture is appropriate. Once this prediction is made, it is used to condition the gesture generation model. Our gesture-property prediction results are promising and indicate that the proposed approach is feasible.

Authors are grateful to Stefan Kopp for providing the SaGA dataset and fruitful discussions about it and to Olga Abramov for advising on its gesture-property processing. This work was partially supported by the Swedish Foundation for Strategic Research Grant No. RIT15-0107 and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

REFERENCES

- [1] Chaitanya Ahuja, Dong Won Lee, Ryo Ishii, and Louis-Philippe Morency. 2020. No gestures left behind: Learning relationships between spoken language and freeform gestures. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*. 1884–1895.
- [2] Simon Alexanderson, Gustav Eje Henter, Taras Kucherenko, and Jonas Beskow. 2020. Style-controllable speech-driven gesture synthesis using normalising flows. *Computer Graphics Forum* 39, 2 (2020), 487–496.
- [3] Kirsten Bergmann and Stefan Kopp. 2006. Verbal or Visual? How Information is Distributed across Speech and Gesture in Spatial Dialog. In *Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*. 90–97.
- [4] Kirsten Bergmann and Stefan Kopp. 2009. GNetlc—Using Bayesian decision networks for iconic gesture generation. In *International Workshop on Intelligent Virtual Agents*. Springer, 76–89.
- [5] Kirsten Bergmann and Manuela Macedonia. 2013. A virtual agent as vocabulary trainer: iconic gestures help to improve learners’ memory performance. In *Proceedings of the International Workshop on Intelligent Virtual Agents*. 139–148.
- [6] Justine Cassell, Catherine Pelachaud, Norman Badler, Mark Steedman, Brett Achorn, Tripp Becket, Brett Douville, Scott Prevost, and Matthew Stone. 1994. Animated conversation: rule-based generation of facial expression, gesture & spoken intonation for multiple conversational agents. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*. 413–420.
- [7] Ylva Ferstl, Michael Neff, and Rachel McDonnell. 2020. Understanding the predictability of gesture parameters from speech and their perceptual importance. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*. Article 19, 8 pages.
- [8] Gustav Eje Henter, Simon Alexanderson, and Jonas Beskow. 2020. MoGlow: Probabilistic and controllable motion synthesis using normalising flows. *ACM Transactions on Graphics* 39 (2020), 236:1–236:14.
- [9] Adam Kendon. 2004. *Gesture: Visible Action as Utterance*. Cambridge University Press.
- [10] Mark L. Knapp, Judith A. Hall, and Terrence G. Horgan. 2013. *Nonverbal Communication in Human Interaction*. Wadsworth, Cengage Learning.
- [11] Ivan Kobayev, Simon Prince, and Marcus Brubaker. 2020. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
- [12] Stefan Kopp and Ipke Wachsmuth. 2004. Synthesizing multimodal utterances for conversational agents. *Computer Animation and Virtual Worlds* 15, 1 (2004), 39–52.
- [13] Taras Kucherenko, Dai Hasegawa, Naoshi Kaneko, Gustav Eje Henter, and Hedvig Kjellström. 2021. Moving fast and slow: Analysis of representations and post-processing in speech-driven automatic gesture generation. *International Journal of Human-Computer Interaction* (2021), 1–17.
- [14] Taras Kucherenko, Patrik Jonell, Sanne van Waveren, Gustav Eje Henter, Simon Alexanderson, Iolanda Leite, and Hedvig Kjellström. 2020. Gesticulator: A framework for semantically-aware speech-driven gesture generation. In *Proceedings of the ACM International Conference on Multimodal Interaction*. 242–250.
- [15] Taras Kucherenko, Patrik Jonell, Youngwoo Yoon, Pieter Wolfert, and Gustav Eje Henter. 2021. A large, crowdsourced evaluation of gesture generation systems on common Data: The GENE Challenge 2020. In *Proceedings of the 26th International Conference on Intelligent User Interfaces*. 11–21.
- [16] Taras Kucherenko, Rajmund Nagy, Michael Neff, Hedvig Kjellström, and Gustav Eje Henter. 2021. Multimodal analysis of the predictability of hand-gesture properties. *arXiv preprint* (2021).
- [17] Andy Lücking, Kirsten Bergman, Florian Hahn, Stefan Kopp, and Hannes Rieser. 2013. Data-based analysis of speech and gesture: The Bielefeld Speech and Gesture Alignment Corpus (SaGA) and its applications. *Journal on Multimodal User Interfaces* 7, 1 (2013), 5–18.
- [18] Stacy Marsella, Yuyu Xu, Margaux Lhommet, Andrew Feng, Stefan Scherer, and Ari Shapiro. 2013. Virtual character performance from speech. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 25–35.
- [19] David McNeill. 1992. *Hand and Mind: What Gestures Reveal about Thought*. University of Chicago Press.
- [20] Michael Neff, Michael Kipp, Irene Albrecht, and Hans-Peter Seidel. 2008. Gesture modeling and animation based on a probabilistic re-creation of speaker style. *ACM Transactions on Graphics* 27, 1, Article 5 (2008), 24 pages.
- [21] Victor Ng-Thow-Hing, Pengcheng Luo, and Sandra Okita. 2010. Synchronized gesture and speech production for humanoid robots. In *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 4617–4624.
- [22] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. In *Proceedings of the 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing*.
- [23] Carolyn Saund, Andrei Birlădeanu, and Stacy Marsella. 2021. CMCF: An architecture for realtime gesture generation by clustering gestures by motion and communicative function. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*. 1136–1144.
- [24] Yanxiang Wu, Sabarish V. Babu, Rowan Armstrong, Jeffrey W. Bertrand, Jun Luo, Tania Roy, Shaundra B. Daily, Lauren Cairco Dukes, Larry F. Hodges, and Tracy Fasolino. 2014. Effects of virtual human animation on emotion contagion in simulated inter-personal experiences. *IEEE Transactions on Visualization and Computer Graphics* 20, 4 (2014), 626–635.
- [25] Yiming Yang and Xin Liu. 1999. A re-examination of text categorization methods. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. 42–49.
- [26] Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. 2020. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Transactions on Graphics* 39 (2020), 222:1–222:16.
- [27] Youngwoo Yoon, Woo-Ri Ko, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. 2019. Robots learn social skills: End-to-end learning of co-speech gesture generation for humanoid robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 4303–4309.
- [28] Fajriani Yunus, Chloé Clavel, and Catherine Pelachaud. 2019. Gesture Class Prediction by Recurrent Neural Network and Attention Mechanism. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*. 233–235.
- [29] Fajriani Yunus, Chloé Clavel, and Catherine Pelachaud. 2021. Sequence-to-Sequence Predictive Model: From Prosody to Communicative Gestures. In *Digital Human Modeling and Applications in Health, Safety, Ergonomics and Risk Management. Human Body, Motion and Behavior*. Springer International Publishing, Cham, 355–374.