

# ViDA-MAN: Visual Dialog with Digital Humans

Tong Shen<sup>1</sup>, Jiawei Zuo<sup>1</sup>, Fan Shi<sup>1</sup>, Jin Zhang<sup>2</sup>, Liqin Jiang<sup>2</sup>,  
Meng Chen<sup>1</sup>, Zhengchen Zhang<sup>1</sup>, Wei Zhang<sup>1</sup>, Xiaodong He<sup>1</sup>, Tao Mei<sup>1</sup>  
<sup>1</sup> JD AI Research, Beijing, China, <sup>2</sup> Migu Culture Technology, Beijing, China

## ABSTRACT

We demonstrate ViDA-MAN, a digital-human agent for multi-modal interaction, which offers realtime audio-visual responses to instant speech inquiries. Compared to traditional text or voice-based system, ViDA-MAN offers human-like interactions (e.g, vivid voice, natural facial expression and body gestures). Given a speech request, the demonstration is able to respond with high quality videos in sub-second latency. To deliver immersive user experience, ViDA-MAN seamlessly integrates multi-modal techniques including Acoustic Speech Recognition (ASR), multi-turn dialog, Text To Speech (TTS), talking heads video generation. Backed with large knowledge base, ViDA-MAN is able to chat with users on a number of topics including chit-chat, weather, device control, News recommendations, booking hotels, as well as answering questions via structured knowledge.

## CCS CONCEPTS

• **Computing methodologies** → **Computer vision**; **Natural language processing**; • **Human-centered computing** → **Human computer interaction (HCI)**.

## KEYWORDS

Multimodal Interaction, Digital Human, Dialog System, Speech Recognition, Text to Speech, Talking-head Generation

### ACM Reference Format:

Tong Shen<sup>1</sup>, Jiawei Zuo<sup>1</sup>, Fan Shi<sup>1</sup>, Jin Zhang<sup>2</sup>, Liqin Jiang<sup>2</sup>, Meng Chen<sup>1</sup>, Zhengchen Zhang<sup>1</sup>, Wei Zhang<sup>1</sup>, Xiaodong He<sup>1</sup>, Tao Mei<sup>1</sup>. 2021. ViDA-MAN: Visual Dialog with Digital Humans. In *Proceedings of the 29th ACM Int'l Conference on Multimedia (MM '21)*, Oct. 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3474085.3478560>

## 1 INTRODUCTION

Digital humans are virtual avatars backed by Artificial Intelligence (AI), which are designed to behave like a real human. Agents powered by such systems can be applied in a wide range of scenarios such as personal assistant, customer service and News broadcasting. In this paper, we present ViDA-MAN, a multi-modal interaction system for digital humans. The system is complex by nature, integrating multimodal techniques such as ASR, TTS, dialog system, visual synthesis.

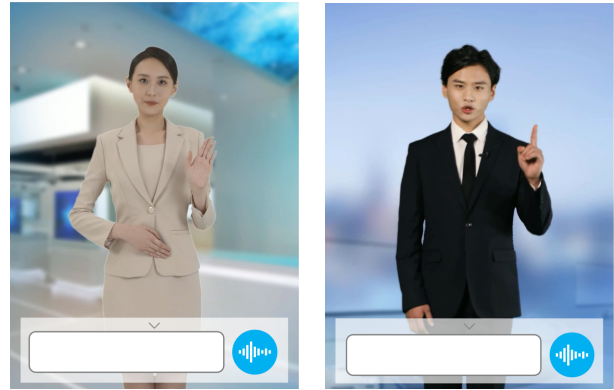


Figure 1: Illustration of our interactive demo. The digital characters have realistic appearance, voice, natural facial expression and body motions, offering lifelike interaction experiences.

One of the core parts of a digital human system is the ability to receive signals from the user and output the corresponding feedback, which can be viewed as a chatbot. We develop a multi-type spoken dialogue system that can handle user requests by multiple dialog skills, such as chit-chat, task-oriented dialog, and question answering based on knowledge graph.

What makes our ViDA-MAN different from a pure chatbot is its concrete visual appearance and voice, which expresses far more information than a pure text-based system, e.g. body language or facial expressions. The voice is empowered by a high quality TTS system, consisting a novel Duration Informed Auto-regressive Network (DIAN)[13] and a speaker-specific neural vocoder LPCNet. The appearance is powered by neural rendering techniques [3, 5, 8–10, 15]. Different from a graphics rendering engines [1, 2], neural renderers do not require a specific high-quality 3D model and are able to produce far more realistic visual results. Figure 1 demonstrates some examples. In this paper, we present our digital human system, ViDA-MAN, to draw more attention on multi-modal interaction systems.

## 2 SYSTEM ARCHITECTURE

The whole system is designed to pursue low latency and high visual fidelity, seeking intelligent and real-time interactions with a lifelike digital character. As shown in Figure 2, the system mainly consists of six modules. The system accepts human voice by an ASR module and feeds it to a dialog system. The response is further translated to realistic voice using TTS. A driving system and a rendering system are responsible for updating the appearance. A streaming service is adopted to integrate everything and encode it into media stream back to the user.

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).

MM '21, October 20–24, 2021, Virtual Event, China.

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

ACM ISBN 978-1-4503-8651-7/21/10

<https://doi.org/10.1145/3474085.3478560>

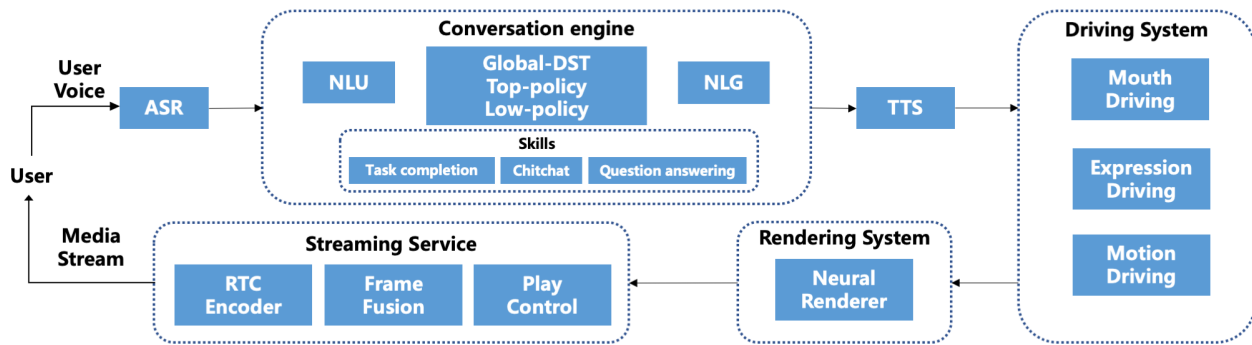


Figure 2: System architecture of ViDA-MAN. The system interacts with the user through human-like voice and appearances. Major sub-modules include ASR, Dialog, TTS, Driving & Visual Renderer, and Streaming service.

## 2.1 Speech Module: ASR and TTS

The speech module (ASR and TTS) acts as interpreters of the whole system. ASR has been widely studied for decades. We choose a widely-used HMM-DNN framework and use TDNN+LSTM [12] to implement the system. The TTS module is essential for human-like voice output. We adopt DIAN [13] and a speaker-specific neural vocoder LPCNet [14]. By adopting the multi-speaker duration model and acoustic model with the speaker-dependent data, the TTS system is able to clone and generate the voice of every new speaker or speaking style, which greatly enriches diversity of choices under different scenarios and can better cooperate with other modules to present digital humans.

## 2.2 Dialog Module: Conversation Engine

Our conversation engine is equipped with multiple dialog skills, which covers various domains including news, weather, movie, etc. As shown in Figure 2, after the ASR, the transcriptions of the user’s speech are firstly passed to a Natural Language Understanding (NLU) unit to identify the user’s intention and related entity information. The information is then formatted as the input to DST (Dialog State Tracker), which maintains the current state of the dialog. In order to coordinate different skills naturally, a hierarchical dialog policy controller [6] is adopted to select next dialogue actions. The main skills are controlled by top-level policy, whereas primitive actions are controlled by low-level policy. Finally, a model-based Natural Language Generator (NLG) [16] is involved to convert agent actions to natural language responses.

## 2.3 Visual Module: Driving and Rendering

The appearance of our digital human actually originates from our real human model. We first record videos of the human model speaking or making different poses. The footage is then used to reconstruct the person with intermediate representations. The role of the driving system is to produce desired representations given speech signals. Our driving system is mainly focused on the face region, i.e. making facial expressions, moving lips etc., but it can be extended to the whole body. The face is reconstructed and represented by a 3D Morphable Model (3DMM) [4, 7, 11].

The rendering system is mainly composed of a neural renderer, which converts the intermediate representation (3DMM mesh) into realistic images. We adopt a modified version of Vid2Vid [15], which is a video-to-video synthesis framework considering both spatial

and temporal consistency. The model is trained by sequences of paired images, real face images and the reconstructed ones.

## 2.4 Realtime Streaming

Our whole system is streaming-based, therefore the streaming service unit plays an essential role in controlling the pace, encoding the frames and broadcasting the media. There are three modules involved. A play control unit is used to control the general status of the digital human. Since our driving system is mainly related to facial movement, the body frames (actions, transitions, etc.) are pre-loaded and controlled by the unit. The render system only starts to render when there are facial movements, which saves computational cost and increases scalability. The face images are then processed and fused with the body frames through a frame fusion unit and further encoded by RTC (Real-Time Communication) encoder. The frames with the synced audio content are compressed into H.246 format and broadcasted by our own RTC platform.

## 3 SYSTEM AND USER INTERFACE

Our dialog system is trained on multiple types of data including a knowledge graph with 180k triples, 22k entities, 20k QA pairs and million-scale NLU data, which covers topics such as celebrities, music, movies. ViDA-MAN supports tasks like enquiring weather, reserving hotels, searching news, etc. Moreover, it is also extendable to more domains in the future.

Voice and appearance are two crucial components for providing human-like interactions. Our TTS benefits a lot from the novel framework and can generate high-quality voice. On a scale from 0 to 5, our TTS system reaches 4.2 MOS (mean opinion score) in our user study, which shows superior quality. Our rendering system also benefits from our joint 2D and 3D based generation networks that produce realistic talking faces in realtime. Figure 1 shows examples of our user interface.

ViDA-MAN supports streaming of 1080p video at 25 FPS (frames per second). The overall latency is around 400 milliseconds when properly deployed. One single GPU (Tesla P40) is able to support four instances simultaneously.

## ACKNOWLEDGMENTS

This work was partially supported by the National Key R&D Program of China under Grant No. 2020AAA0108600, Migu-JD Joint AI-Lab on Multimodal Interaction.

## REFERENCES

- [1] Unity. [www.unity.com](http://www.unity.com).
- [2] Unreal Engine. [www.unrealengine.com](http://www.unrealengine.com).
- [3] Kara-Ali Aliev, Dmitry Ulyanov, and Victor S. Lempitsky. 2019. Neural Point-Based Graphics. *CoRR* abs/1906.08240 (2019). arXiv:1906.08240 <http://arxiv.org/abs/1906.08240>
- [4] Volker Blanz and Thomas Vetter. 1999. A Morphable Model for the Synthesis of 3D Faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., USA, 187–194. <https://doi.org/10.1145/311535.311556>
- [5] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. 2019. Everybody Dance Now. In *IEEE International Conference on Computer Vision (ICCV)*.
- [6] Tejas D Kulkarni, Karthik R Narasimhan, Ardavan Saeedi, and Joshua B Tenenbaum. 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. *arXiv preprint arXiv:1604.06057* (2016).
- [7] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.* 36, 6 (2017), 194:1–194:17. <https://doi.org/10.1145/3130800.3130813>
- [8] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. 2018. Deep Appearance Models for Face Rendering. *ACM Trans. Graph.* 37, 4, Article 68 (July 2018), 13 pages. <https://doi.org/10.1145/3197517.3201401>
- [9] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. 2019. Neural Volumes: Learning Dynamic Renderable Volumes from Images. *ACM Trans. Graph.* 38, 4, Article 65 (July 2019), 14 pages. <https://doi.org/10.1145/3306346.3323020>
- [10] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. 2019. Semantic Image Synthesis With Spatially-Adaptive Normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [11] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. 2009. A 3D Face Model for Pose and Illumination Invariant Face Recognition. In *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*.
- [12] Vijayaditya Peddinti, Daniel Povey, and S. Khudanpur. 2015. A time delay neural network architecture for efficient modeling of long temporal contexts. In *INTERSPEECH*.
- [13] Wei Song, Xin Yuan, Zhengchen Zhang, Chao Zhang, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2021. Dian: duration informed auto-regressive network for voice cloning. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 8598–8602.
- [14] J. Valin and J. Skoglund. 2019. LPCNET: Improving Neural Speech Synthesis through Linear Prediction. *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2019), 5891–5895.
- [15] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc.
- [16] Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *arXiv preprint arXiv:1508.01745* (2015).