Datasheet for RenderMe-360

1   Motivation

Q1: For what purpose was the dataset created?

RenderMe-360 is a head-oriented, large-scale, high-resolution, and multi-sensory 4D human head dataset with rich annotations. It is created as a comprehensive digital asset library, to ensure the compatibility of evaluating multiple head avatar tasks in one single dataset. Before the curation of RenderMe-360, the publicly available datasets can only support research on a small set of problems. These datasets are either limited in resolution due to non-head-oriented capture system, or suffer from data volume and diversity. With the release of RenderMe-360 dataset, researchers could train their models with more robust data prior learned from our data, or evaluate their algorithms with more exhaustive observations.

Along with the data, we also provide benchmarks on five vital tasks, *i.e.,* novel view synthesis, novel expression synthesis, hair rendering, head editing, and talking head generation. On the one hand, the benchmarks uncover the strengths and flaws of state-of-the-art methods, which present many new observations and challenges for the research community to catalyze future research on the human head avatar. On the other hand, as no benchmark agreement for human head avatar research is achieved across different institutions yet, we hope our attempt on the five vital tasks could provide unified criteria or paradigms for this domain.

In a nutshell, we create RenderMe-360 to facilitate the development of advanced research on high-fidelity head avatar synthesis, which spurs new opportunities not only from our formalized benchmarks, but also alternative ones that the community might come up with from our dataset.

Q2:Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?

This knowledge base was created by RenderMe-360 Team: Dongwei Pan (Shanghai AI Laboratory and SenseTime), Long Zhuo (Shanghai AI Laboratory), Jingtan Piao (SenseTime and The Chinese University of Hong Kong), Huiwen Luo (Shanghai AI Laboratory), Wei Cheng (Shanghai AI Laboratory and SenseTime), Yuxin Wang (Shanghai AI Laboratory), Siming Fan (SenseTime), Shengqi Liu (SenseTime), Lei Yang (SenseTime), Bo Dai (Shanghai AI Laboratory), Ziwei Liu (S-Lab, NTU), Chen Change Loy (S-Lab, NTU), Chen Qian (Shanghai AI Laboratory), Wayne Wu (Shanghai AI Laboratory), Dahua Lin (Shanghai AI Laboratory and The Chinese University of Hong Kong), and Kwan-Yee Lin (Shanghai AI Laboratory and The Chinese University of Hong Kong).

Q3:What support was needed to make this dataset?

This work was funded by Shanghai AI Laboratory.

Q4:Any other comments?

No.

Submitted to the 37th Conference on Neural Information Processing Systems (NeurIPS 2023) Track on Datasets and Benchmarks. Do not distribute.

2 Composition

Q5: What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)?

(1) Documents: The dataset consists of source video data (multi-view RGB frames, audio files), and rich annotations (camera parameters, matting, 2d landmarks, 3d landmarks, FLAME parameters, UV texture map, scan mesh, and text annotation). We process the data and package the data with annotation together in the form of SMC files.

(2) Countries of captured subjects: the performers come from different countries all over the world. The ethnicity covers White, Asian, Brown (Middle Eastern, Indian, and Latino), and Black.

Q6: How many instances are there in total (of each type, if appropriate)?

We provide human head assets from 500 different identities. The collected subjects vary from different ages, eras, ethnicities, and cultures. Each subject is required to perform various motions with respect to expression, hair movement, and speech activities.

(1) Overall data volume: we provide $243+$ million frames in full-head level and over $800k$ video sequences captured from multi-view cameras. Along with the dataset, we provide rich annotations. For per-frame, we provide annotation of camera parameters, matting, and facial 2D/3D landmarks. The expression data includes additional action unit annotation. For key frames of each identity, we provide FLAME parameters, UV texture map, and scan mesh. The volume of these annotations is over 600k, 600k, 6k respectively. In addition, per-id annotations also include text annotations, which cover four major aspects, each identity contains 97 static facial features, 36 static information of non-facial regions, 25 dynamic facial actions, and $74 - 84$ dynamic video activity descriptions.

(2) Data volume for each type capture: The frame volume consists of more than 6000 expression data, 1300 hairstyle capture data, and more than 20000 sentence collections. Please refer to Figure 3 in the main paper and Section 3 in the PDF file "Supplementary Materials" for more details.

Q7: Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?

We will provide all instances from our RenderMe-360 data repositories.

Q8: What data does each instance consist of?

We provide raw data and their associated annotations.

Q9: Is there a label or target associated with each instance?

Yes. We provide annotations in two levels – per-frame annotations, and per-id annotations. The per-frame annotations are matting, camera parameters, 2d landmarks, and 3d landmarks. The per-id annotations refer to the ones annotated on key frames of each identity, in which scan mesh, FLAME parameters, UV texture map, and text annotations are included. The rich annotation is one of the characteristics of our dataset.

Q10: Is any information missing from individual instances?

No. Note that, for performers who originally wear unique hair accessories or hairstyles that are difficult to take off, the wig recordings are intentionally skipped.

Q11: Are relationships between individual instances made explicit (e.g., users' movie ratings, social network links)?

No.

Q12: Are there recommended data splits (e.g., training, development/validation, testing)?

We classify the dataset into three categories according to the complexity of texture and geometry, which are *Normal Case*, *With Deformable Accessories*, and *With Complex Accessories*. Based on different benchmarks, we further split the dataset into multiple subsets with different dimensions. Detailed information could be found in the PDF file "Supplementary Materials". Still and all, there are many other possible usages of our comprehensive data repository. For example, one could use the

2

entire database for pre-training. We encourage researchers to creatively use our data according to their experimental needs.

Q13: Are there any errors, sources of noise, or redundancies in the dataset?
Due to the hardware limitation, the recording data of a few subjects miss one or two camera views. For more discussion, please refer to Section 2.2 in the PDF file "Supplementary Materials".

Q14: Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)?
Yes. The dataset is self-contained.

Q15: Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals' non-public communications)?
No.

Q16: Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety?
No.

Q17: Does the dataset relate to people?
Yes.

Q18: Does the dataset identify any subpopulations (e.g., by age, gender)?
We provide basic information of each subject, like age, gender, ethnicity, height, and weight. Please refer to Figure S14 in the PDF file "Supplementary Materials" for more details.

Q19: Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset?
Yes. RenderMe-360 is a human head centric dataset for facilitating the development of advanced research on high-fidelity head avatar synthesis. Thus, it is feasible to identify people via face recognition. All performers have assigned files that consenting the data capture and release for non-commercial usage.

Q20: Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; criminal history)?
No.

Q21: Any other comments?
We call for responsible usage of the dataset for research purposes only.

3   Collection

The collection procedure and preprocessing are explained in detail in Section 2.2 and Section 3 of the PDF file "Supplementary Materials".

Q22: How was the data associated with each instance acquired?
The original data are captured directly from the subjects by multi-view cameras and audio equipment. We get the annotation data with the help of our semi-automated annotation tools and manual annotation.

Q23: Over what timeframe was the data collected?
The valid data were collected from August 2022 to May 2023.

Q24: What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)?
We set up a high-end data collection system, named POrtrait Large-scale hIgh-quality Capturing

3

sYstem (POLICY), to capture the raw data. The details of the POLICY and collection procedures refer to the Section 2.1 and 2.2 in the PDF file "Supplementary Materials".

Q25:If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?
No.

Q26: Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?
The performers come from different professions to ensure the diversity of appearance. We employ contractors for manual labeling. The data curation is done by RenderMe-360 co-authors.

Q27: Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., websites)?
We ask each individual to perform under explicit instructions.

Q28: Were the individuals in question notified about the data collection and use of their data?
Each subject has signed the notification of the collection process, and data usage consent before data capture.

Q29: Any other comments?
No.

4 Preprocessing / Cleaning / Labeling

Q30: Was any preprocessing/cleaning/labeling of the data done(e.g.,discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)?
Yes. We preprocess the raw data in the form of images and aligned audio. We clean the data which do not satisfy our requirement by manually checking. We provide rich annotation of the dataset, including camera parameters, human mask, FLAME parameters, 2d landmarks, 3d landmarks, UV texture map, scan mesh, and text annotations.

Q31: Is the software used to preprocess/clean/label the instances available?
Yes. We will release our preprocessing and labeling tools on `https://github.com/RenderMe-360/RenderMe-360`.

Q32:Any other comments?
No.

5 Uses

Q33: Has the dataset been used for any tasks already?
In the experiment part of our paper, we evaluate the state-of-the-art methods in different tasks with our collected dataset.

Q34: Is there a repository that links to any or all papers or systems that use the dataset?
N/A.

Q35: What (other) tasks could the dataset be used for?
Our data library is primarily intended to facilitate the development of advanced research on high-fidelity head avatar synthesis. We set up example benchmarks on five vital tasks: novel view synthesis, novel expression synthesis, hair rendering, hair Editing, and talking head generation. Other directions related to human heads are also possible, such as text-to-3D human head generation, face parsing, matting, *etc* .

Q36: Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?
Nothing we are aware of.

4

Q37: Are there tasks for which the dataset should not be used?

Any research that intentionally generates harmful or toxic content will be not condoned by us.

Q38: Any other comments?

No.

6 Distribution

Q39: Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?

Yes, the dataset will be open-source.

Q40: How will the dataset will be distributed (e.g., tarball on website, API, GitHub)?

The dataset will be distributed by the tarball on the official website `https://renderme-360.github.io/`

Q41: When will the dataset be distributed?

A part of the dataset has been distributed on the official website, and the full dataset will be distributed before the start of the NeurIPS conference.

Q42: Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?

Yes. The license of our dataset is Creative Common CC-BY 4.0 license `https://creativecommons.org/licenses/by/4.0/`

Q43: Have any third parties imposed IP-based or other restrictions on the data associated with the instances?

NO

Q44: Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?

No.

Q45: Any other comments?

No.

7 Maintenance

Q46: Who is supporting/hosting/maintaining the dataset?

The RenderMe-360 team.

Q47: How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

Please contact Dongwei Pan (pandongwei@sensetime.com), Long Zhuo (zhuolong@pjlab.org.cn), Dahua Lin (dhlin@ie.cuhk.edu.hk), and Kwan-Yee Lin (linjunyi9335@gmail.com)

Q48: Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?

Yes. If there is any update, we will make announcements on the official website `https://renderme-360.github.io/`.

Q49: If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?

Yes. We welcome contributors to communicate with us and leave any comments on the official GitHub page `https://github.com/RenderMe-360/RenderMe-360`.

Q50: Any other comments?

No.

5