1 Checklist

1. For all authors...

    (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] We introduce our dataset in Section 3, analyze online news coverage and how models can leverage it in Section 4, and introduce a baseline for video retrieval on this dataset in Section 5. Evidence can be found in **main paper**.

    (b) Did you describe the limitations of your work? [Yes] We describe performance downsides of our baseline in **Section 5** and detail shortcomings of our dataset in **the datsheet (supplemental material**.

    (c) Did you discuss any potential negative societal impacts of your work? [Yes] We note potential negative consequences of the dataset in the datasheet. Evidence can be found in **the supplementary material**.

    (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] Yes, our paper conforms to the NeurIPS ethics guidelines.

2. If you are including theoretical results...

    (a) Did you state the full set of assumptions of all theoretical results? [N/A] Our paper does not include theoretical results.

    (b) Did you include complete proofs of all theoretical results? [N/A] Our paper does not include theoretical results.

3. If you ran experiments (e.g. for benchmarks)...

    (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We include our baseline and data used in experiments. External code and models can be downloaded and run from their respective open-source repositories. Evidence can be found in **the supplemental material**.

    (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A] Our methods do not include training. We cite the papers where our pretrained models' training information is listed in **Section 5**.

    (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] We run each model on one seed on the same data and report the results.

    (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] The computation time was not significant (<12 hours for all experiments on 1 GPU) and was run on an internal cluster.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

    (a) If your work uses existing assets, did you cite the creators? [N/A] Our work does not use existing assets beyond publicly released videos and articles. We link these videos and articles in our dataset rather than providing the raw files themselves so content authorship is documented.

    (b) Did you mention the license of the assets? [Yes] Videos and articles are subject to their linked sources' respective licenses. Evidence can be found in **the datasheet (supplementary material)**.

    (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include all new assets in the **supplemental material**.

    (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] Video and text content is uploaded to public sharing sites and is removed from the dataset when removed from these sites. Evidence can be found in **the data sheet (supplementary material)**.

1

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] There may be sensitive content due to the nature of news videos. This is elaborated upon in **the data sheet (supplementary material)**

5. If you used crowdsourcing or conducted research with human subjects...

   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] We did not use crowdsourcing for annotations (e.g., Amazon Mechanical Turk). We worked with internal annotators from our research center. We provide a written description of the annotation task in **Section 4**.

   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A] We do not conduct any human subject experiments.

   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A] We did not use crowdsourcing for annotations (e.g., Amazon Mechanical Turk). We worked with internal annotators from our research center.

2