Selection of Optimal Salient Time Steps by Non-negative Tucker Tensor Decomposition

J. Pulido[1] and J. Patchett[1] and M. Bhattarai[1] and B. Alexandrov[1] and J. Ahrens[1]

[1]Los Alamos National Laboratory, USA

Abstract

Choosing salient time steps from spatio-temporal data is useful for summarizing the sequence and developing visualizations for animations prior to committing time and resources to their production on an entire time series. Animations can be developed more quickly with visualization choices that work best for a small set of the important salient timesteps. Here we introduce a new unsupervised learning method for finding such salient timesteps. The volumetric data is represented by a 4-dimensional non-negative tensor, $\mathcal{X}(t,x,y,z)$.The presence of latent (not directly observable) structure in this tensor allows a unique representation and compression of the data. To extract the latent time-features we utilize non-negative Tucker tensor decomposition. We then map these time-features to their maximal values to identify the salient time steps. We demonstrate that this choice of time steps allows a good representation of the time series as a whole.

CCS Concepts

• ***Human-centered computing*** → ***Visualization design and evaluation methods;***

1. Introduction

Supercomputers get ever larger with equitably larger potential for data extracts that allow simulation scientists the ability to analyze and study the evolution of the phenomena they model. Outputs are frequently high spatial resolution timesteps, sampled regularly in simulated time or simulation cycles, leaving hundreds or thousands of timesteps on persistent storage representing terabytes to petabytes of data to read and process for summary results. Depending on the maturity and novelty of the model, those data may be repeatedly analyzed for different purposes. We are concerned with the process of efficiently developing and viewing imagery derived from full spatial resolution timesteps "dumped" from simulations.

In this work, we are interested in identifying a subset of key timesteps, i.e. salient timesteps, that efficiently represent an entire time series. This subset can be used both for developing visualization parameters to process the entire time series, to use as keyframes to summarize important moments in the simulation, and to more efficiently communicate to collaborators key timesteps worth investing further analysis.

We introduce the application of products of unsupervised machine learning based on non-negative tucker factorization (NNTF). This novel NNTF process is adapted to produce a set of features and a set of weights for each of those features. For each feature, we identify the specific timestep that it is most heavily associated with, and tag that time step as a keyframe. For comparison, we propose a set of evaluation criteria to measure the quality of selected key timesteps of our experimental method against other traditional methods.

2. Related work

For workflows, Ma et al. [MLF*12] covers the importance of streamlining the visualization process through scientific storytelling, where the selection of important points in a simulation, i.e. keyframes, can improve the process of production visualization.

A variety of techniques exist for identifying keyframes. Meyers, et al. [MLF*16] look at salient time step selection for in situ use cases, they assume they have an incoming stream of time steps and have to decide whether to keep each time step or not as it is produced. They produce a linear piecewise model and evaluate each new timestep against that model until a new timestep falls out of the precision of that model at which point they store a new timestep and start a new model. Tong, et al. [TLS12] find the best K time steps from a time varying dataset with a user specified K. They define a cost between time steps and using Dynamic Time Warping (DTW) and then use dynamic programming to find the minimum cost for a given k timesteps. Frey and Ertl use a flow based distance metric to quantify the distance to neighboring time steps and minimize the selected time steps difference from the entire time series. Our work is similar to Porter, et al. [PXv*19] who use an unsupervised machine learning autoencoder to develop a feature set based on all time steps, then select keyframes based on those features. We use a Non-negative Tucker Factorization to identify features and their weights, then select keyframes based directly on those weights.

DOI: 10.2312/evs.20211055

In practice, there has been usage of more traditional methods for identifying keyframes. In this paper, we make use of such methods to perform comparative analysis against NNTF.

Wavelets for keyframe selection have traditionally been in used in selecting important frames from a series of videos [CLN99]. By using the Haar basis function, a multiscale wavelet decomposition is performed for every keyframe for all keyframes in a dataset. More recent improvements to this method have been made by the introduction of K-means clustering on the wavelet coefficients [HNMK04]. By performing K-means on the high-pass, fine coefficients for a target K frames, a K amount of clusters are computed that represent keyframe groups. Withing each keyframe group cluster, the keyframe that has the shortest distance to the center cluster is considered the most important keyframe.

For large data simulations, it is important to estimate important keyframes or timesteps within a simulation for post-analysis and in-situ analysis routines. In Meyers et al. a new system of unique-float binning is used to extract important keyframes as time-steps in a simulation [MLF*16].

In information theory, it is believed that an entropy-based approach can provide valuable data points when selecting time steps from a simulation. When evaluating multiple timesteps, the frames of which the highest entropy are selected but may have an unfortunate outcome of selecting subsequent frames in a row that relate to a cluster of keyframes with high activity.

Finally, the most trivial method of selecting time steps from a simulation is to pick every nth cycle from the simulation. This method is most commonly used in AMR codes or Lagrangian codes where the cell sizes can change from timestep to timestep.

3. Non-negative Tucker Tensor Decomposition for Salient Timestep Selection

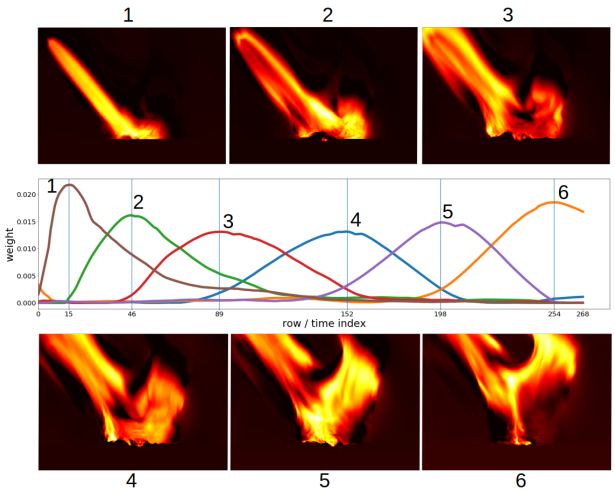

Figure 1: *The weights for each feature plotted over timestep. A rendering of the feature that lines up with each time step is numbered. Vertical lines on the plot show the keyframe (time index) that will be chosen based on that feature.*

3.1. Non-negative Tucker Factorization

In big-data analysis, it is often difficult to link directly the data to the parameters of the generating processes, since the datasets are formed exclusively by directly observable quantities, while the underlying processes/features remain unobserved, hidden, or latent [Eve13]. Extracting these *latent* (not directly observable) features not only reveals valuable information about hidden causality and mechanisms, but it also reduces the dimensionality by revealing the low-dimensional latent structure that represents the whole dataset.

3.2. Matrix Factorization

One of the most powerful tools for extracting latent features is factor analysis. In two dimensions factor analysis can be performed by various versions of Principle Component Analysis (PCA) [Jol86], Independent Component Analysis (ICA) [ACY96], or Non-negative Matrix Factorization (NMF) [PT94]. The presence of the non-negativity constraint in NMF makes the extracted latent features physically interpretable, since they are parts of the data [LS99]. Importantly, many variables, e.g., pixels, density, counts, etc., are naturally non-negative and the extracted features will not have physical meaning if the non-negativity is not at place.

3.3. Tensor Factorization

Most datasets are high-dimensional and are represented by *tensors*, or multidimensional arrays. Such tensors typically describe multiple concurrent latent processes imprinting their signatures in the observable state variables in different dimensions.

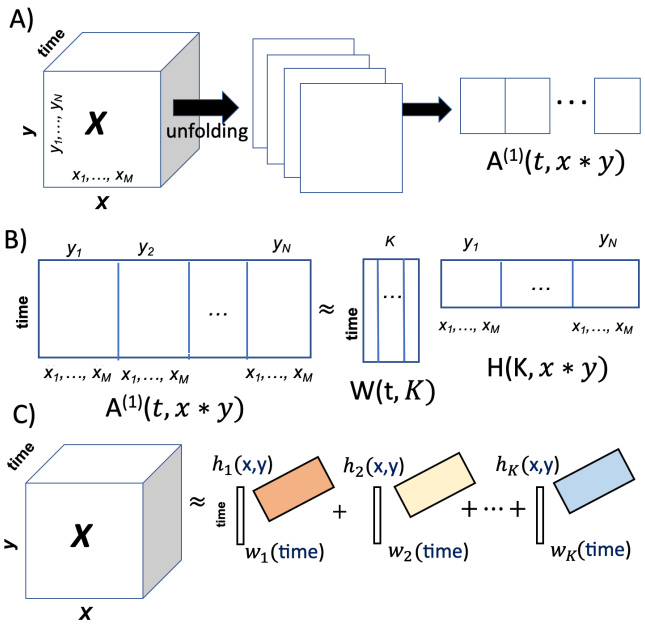

Figure 2: *Illustration of nTD-1 for a 3-dimensional tensor. A) Unfolding of the tensor $\mathcal{X}(t,x,y)$. B) Decomposition of the unfolding, $\boldsymbol{A}^{(1)}(t,x*y)$, via NMF. C) Reshaping the columns of the matrix $\boldsymbol{H}$ to obtain the space-factors.*

Tensor factorization, which is the higher-dimensional analog of

matrix factorization, is an unsupervised learning method that represents a cutting-edge approach for factor analysis. Its main objective is to decompose a high-dimensional tensor into factor matrices and a core-tensor, where the factor matrices carry the latent features in each tensor dimension [KB09]. *Tucker decomposition* is a classical tensor factorization [Tuc66]. Tucker decomposition is not unique but by Tucker decomposition one can extract the minimal subspaces [FH12], which represent the latent features in the corresponding dimensions [ASVR19]. If the data is non-negative, we perform non-negative Tucker decomposition and the minimal subspaces become cones [ADMS19].

3.4. Non-negative Tucker-1 decomposition

In this paper we utilize non-negative Tucker-1 decomposition (nTD-1) [CZPA09], illustrated on Figure (2). nTD-1 unfolds/flattens the original volumetric data-tensor, $\mathcal{X}(t,x,y,z)$, along its time dimension. Thus, nTD-1 converts the 4-dimensional tensor, $\mathcal{X}(t,x,y,z)$, to a $2d$ representation, $\boldsymbol{A}^{(1)}(t,x*y*z)$, see Figure (2), panel A. Further, nTD-1 uses NMF to extract the latent structure of $\boldsymbol{A}^{(1)}(t,x*y*z)$. Specifically, NMF approximates, $\boldsymbol{A}^{(1)} \in \mathbb{R}_{+}^{t\times x*y*z}$, with a product of two non-negative factor matrices, $\boldsymbol{W} \in \mathbb{R}_{+}^{t\times K}$ and $\boldsymbol{H} \in \mathbb{R}_{+}^{K\times x*y*z}$ (Figure (2), panel B), such that the difference,

$$O = \|\boldsymbol{A}^{(1)}(t,x*y*z) - \sum_{s=1}^{K} W_s(t)H_s(x*y*z)\|_{dist} \qquad (1)$$
$$W_s(t) \geqslant 0;\ H_s(x*y*z) \geqslant 0.$$

is minimal under some distance, $\|...\|_{dist}$ and for a given small latent dimension K. Here, for $\|...\|_{dist}$ we use Kullback–Leibler divergence: $D_{KL}(X||Y) = \sum_{i,j} X_{i,j}\log(\frac{X_{i,j}}{Y_{i,j}}) - X_{i,j} + Y_{i,j}$. We utilize the distributed NMF algorithm presented in [BCS*20a] to perform the matrix factorization of the large-scale datasets for the keyframe selection.

3.5. The Salient Time Step Selection

NMF is underpinned by a statistical generative model of superimposed components that can be treated as latent variables of Gaussian, Poisson, or other mixed model [FC09]. In our case, the K-columns, $W_s(t)$ of $\boldsymbol{W}$, represent the latent time-features, while the K-columns, $H_s(x*y*z)$ of the transposed matrix $\boldsymbol{H}$, are the corresponding space-factors. After the extraction of the factor-matrices, $\boldsymbol{W}(t,K)$ and $\boldsymbol{H}(K,x*y*z)$, we reshape the matrix $\boldsymbol{H}(K,x*y*z)$ to construct the 3-dimensional tensors, $\mathcal{H}_s(x,y,z)$. Each $\mathcal{H}_s(x,y,z)$ corresponds to a 3-dimensional space-feature, and we have,

$$\mathcal{X}(t,x,y,z) = \sum_{s=1}^{K} W_s(t)\mathcal{H}_s(x,y,z) + \mathcal{E}(t,x,y,z), \qquad (2)$$

Figure (2), panel C. In (Eq. 2) $\mathcal{E}$ is the tensor error of minimization.

We leverage this observation, (Eq. 2), to select the timesteps that are most strongly associated with each space-feature. This is done simply by finding the index of the largest value in each of the K columns of $\boldsymbol{W}$. This mapping provides us with K easy interpretable features, $\mathcal{H}_s(x,y,z)$, each associated with a specific "influential" time point, the set of which we call the optimal latent salient timesteps. The code for this method is available [BCS*20b].

4. Results

4.1. Evaluation Criteria

To evaluate the selection of K number of keyframes, we use the comparative analysis framework named Foresight [GBP*20]. To evaluate the quality of selected K keyframes, we perform a full temporal dataset reconstruction using only those K keyframes. To reconstruct time-steps in between the selected keyframes, we perform a linear interpolation to the nearest frame. For this initial study, we've selected K=19 because it gives us reasonable results for all of the methods tested in this paper, and provides fair results for traditional methods such as regular and random sampling.

To evaluate the quality of a fully reconstructed temporal dataset from a number of keyframes, we use traditional statistical and image quality algorithm metrics: Total absolute error (TAE), mean-square error (MSE), Peak signal-to-noise ratio (PSNR) [HZ10], signal-to-noise ratio (SNR), the structural similarity index measure (SSIM) [ZBSS04], the multi-scale structural similarity index measure (MS-SSIM) [WSB03], and a Universal-image Quality Index (UQI) [ZB02]. The quality of reconstructed datasets are evaluated by the minimization of TAE and MSE, while maximizing SNR, PSNR, SSIM, MS-SSIM, and UQI metrics.

There are key differences between statistical metrics and image quality algorithm metrics. Traditional statistical methods, i.e. TAE, MSE, SNR, and PSNR, aim in measuring differences between the data points (cells) of the original and approximated datasets. In contrast, image quality metrics, i.e. SSIM, MS-SSIM, and UQI, do the opposite and quantify similarities between data points. Unlike the former, image quality metrics are standardized, i.e. values between 0 to 1, making it easier to evaluate and understand regardless of dataset type. These differences make image quality metrics more suitable for targeting visualization-oriented applications.

4.2. Deep-water impact dataset

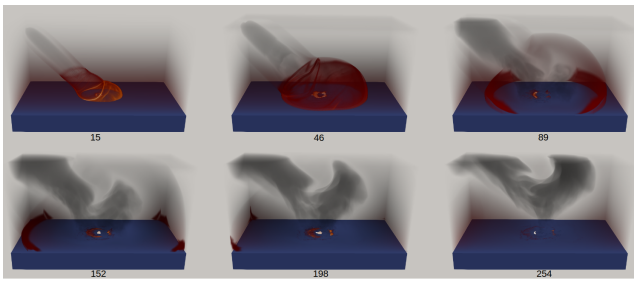

Figure 3: *Volume renderings of six Deep Water Impact time steps chosen by our method.*

The Deep Water Impact Ensemble dataset [PG17], is a set of simulations created to study the propensity of asteroids impacting deep ocean water to create tsunamis. Simulated by xRage [GWC*08], a parallel multi-physics Eulerian hydrodynamics code developed at the Los Alamos National Laboratory, the dataset is produced using an Adaptive Mesh Refinement (AMR) technique that allows higher and lower resolution areas of the simulation grid for computation. We use the yB31 ensemble member- a 250 meter diameter asteroid entering the atmosphere at 45 degrees, and

Method	TAE	PSNR	SNR	MSE*	SSIM	MS-SSIM	UQI
Regular	4,598	48.084	39.742	0.015	0.9291	0.9294	0.7099
Random	9,926	45.093	36.718	0.097	0.9321	0.9330	0.6953
Entropy	44,011	31.936	23.594	0.719	0.9249	0.9293	0.3789
Wavelets(M1)	25,374	35.894	27.552	0.340	0.9266	0.9293	0.5071
Wavelets(M2)	6,090	46.231	37.888	0.023	0.9289	0.9293	0.7202
Unique Floats	11,584	42.345	34.003	0.339	0.9281	0.9293	0.6471
NNTF	8,112	45.202	36.760	0.064	0.9398	0.9405	0.7142

Table 1: *Results for the key-frame selection methods on the Deep-water impact dataset. MSE* is $x10^{-3}$. Wavelets **M1** is coefficient thresholding, while Wavelets **M2** is K-Means clustering.*

Method	TAE*	PSNR	SNR	MSE	SSIM	MS-SSIM	UQI
Regular	318.5	29.09	35.08	735.8	0.8991	0.9029	0.9001
Random	529.1	27.84	32.33	1358	0.8957	0.9011	0.8972
Entropy	1730.0	19.60	24.56	8783	0.8921	0.8966	0.8948
Wavelets(M1)	578.7	23.07	29.09	1832	0.8971	0.9021	0.8989
Wavelets(M2)	220.9	29.99	34.17	517.7	0.9012	0.9044	0.9022
Unique Floats	1068.5	20.79	25.37	3802	0.8924	0.8989	0.8947
NNTF	611.3	25.93	30.56	1984	0.9016	0.9071	0.9035

Table 2: *Results for the key-frame selection methods on Antarctic Ice-sheet ocean dataset. TAE* is ($x10^6$). Wavelets **M1** is coefficient thresholding, while Wavelets **M2** is K-Means clustering.*

air bursting at 5km above the sea. xRage used ParaView Catalyst [ABG*15] to produce the visualization dumps. Our study uses a resampling of the AMR data to a regular grid of $460x280x240$ and has 269 time-steps.

When analyzing the results shown in Table. 1, we found NNTF to excel in image quality metrics compared to the other methods. Performing regular sampling, usually the most looked at choice, presents numerically the best possible case to get an overall view of a simulation but image metrics say otherwise. When using a more explorative and direct approach in selecting keyframes, methods such as wavelets (M2) and NNTF excel in image quality metrics while introducing a small amount of error, with NNTF coming out on top. This data comparison signifies that the approximated data series produced from K=19 keyframes is more visually similar to the original data series, while at the same time introducing higher amount of point-wise error. As explained at the start of Section 4, image quality algorithm metrics focus on capturing visual similarities rather than numerical differences between datasets, therefore making keyframe selection methods desirable.

4.3. Ocean modeling dataset

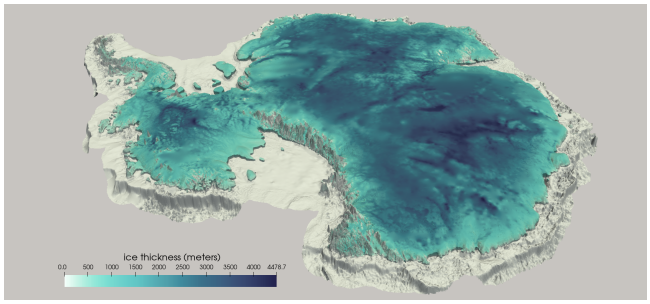

Figure 4: *Early time of the ABUMIP data set showing ice thickness with all floating ice removed.*

The ABUMIP Land ice modeling dataset [SPS*20] represents a simulation from the CMIP6 Ice Sheet Model Intercomparison Project (ISMIP6) [NPL*16] and was run using the MPAS ALbany Land Ice (MALI) simulation code. ABUMIP investigates an extreme scenario where all ice shelves around Antarctica are removed instantaneously and prevented from reforming over a period of 500 years. The ABUMIP experiment is climatologically unrealistic. However, it provides an estimate for the upper-bound response of the Antarctic ice sheet to the loss of its ice shelves. For this model, MALI uses a mesh with 2 km resolution in dynamically important areas near the coast and coarsens to 30 km resolution in the slow moving ice sheet interior, with a total of 1.8 million grid cells in the horizontal over 200 time-steps. The mesh uses ten vertical layers preferentially concentrated near the ice sheet base, where vertical shearing tends to be greatest. The simulation discussed here was run on about 6000 processors at the National Energy Research Scientific Computing Center (NERSC).

The results in Table 2, observe similar quality behavior as the Asteroid Impact dataset. Having processed the selection methods on the raw floating-point surface data, we found that NNTF produces the best image-quality metrics compared to the other keyframe selection methods. Specific to this dataset, we found all methods to generally perform well in image-quality metrics (>0.89) and NNTF to excel above the rest. This can be attributed to slowly evolving numerical data in the original dataset, where hardly significant changes occur compared to the asteroid impact dataset.

5. Conclusions

In this paper, we've shown initial results in keyframe selection methods for explorative visualization. We introduce a set of evaluation criteria that considers statistical and image quality metrics for keyframe quality evaluation. We present a new method, NNTF, that produces the best visual quality metric results compared to other traditional methods in keyframe selection. The analysis also shows that, although traditional methods may have lower numerical error, they may not have the best visual quality metrics relevant visual applications. Overall, by reducing the amount of data needed for analysis by optimally selecting keyframes of interest, this action saves scientists time in exploratory analysis and visualization, and in production workflows.

The results shown in this paper are promising and we intend to do future work. We would like to perform a more thorough analysis on the K selection process, and improve the efficiency of this method to expand it to full 3D simulation datasets and not just subsets. Additionally, we'd like to apply NNTF results to perform data compression and full time-series simulation reconstruction.

6. Acknowledgements

This research was funded by the LANL Laboratory Directed Research and Development (LDRD) grant 20190020DR and Los Alamos National Laboratory Institutional Computing Program, supported by the U.S. Department of Energy National Nuclear Security Administration under Contract No. 89233218CNA000001.

References

[ABG*15] AYACHIT U., BAUER A., GEVECI B., O'LEARY P., MORELAND K., FABIAN N., MAULDIN J.: Paraview catalyst: Enabling in situ data analysis and visualization. In *Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization* (New York, NY, USA, 2015), ISAV2015, Association for Computing Machinery, p. 25–29. 4

[ACY96] AMARI S.-I., CICHOCKI A., YANG H. H.: A new learning algorithm for blind signal separation. In *Advances in neural information processing systems* (1996), pp. 757–763. 2

[ADMS19] ALEXANDROV B., DESANTIS D., MANZINI G., SKAU E.: Nonnegative canonical polyadic decomposition with rank deficient factors. *arXiv preprint arXiv:1909.07570* (2019). 3

[ASVR19] ALEXANDROV B. S., STANEV V. G., VESSELINOV V. V., RASMUSSEN K. Ø.: Nonnegative tensor decomposition with custom clustering for microphase separation of block copolymers. *Statistical Analysis and Data Mining: The ASA Data Science Journal 12*, 4 (2019), 302–310. 3

[BCS*20a] BHATTARAI M., CHENNUPATI G., SKAU E., VANGARA R., DJIDJEV H., ALEXANDROV B. S.: Distributed non-negative tensor train decomposition. In *2020 IEEE High Performance Extreme Computing Conference (HPEC)* (2020), pp. 1–10. 3

[BCS*20b] BHATTARAI M., CHENNUPATI G., SKAU E., VANGARA R., DJIDJEV H., ALEXANDROV B. S.: pydnmfk: Python distributed non negative matrix factorization with determination of hidden features, 2020. [ONLINE]. https://github.com/lanl/pyDNMFk, Last accessed on 2021-04-12. 3

[CLN99] CAMPISI P., LONGARI A., NERI A.: Automatic key frame selection using a wavelet-based approach. In *Wavelet Applications in Signal and Image Processing VII* (Oct 1999), vol. 3813, International Society for Optics and Photonics, p. 861–872. 2

[CZPA09] CICHOCKI A., ZDUNEK R., PHAN A. H., AMARI S.-I.: *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation*. John Wiley & Sons, 2009. 3

[Eve13] EVERETT B.: *An introduction to latent variable models*. Springer Science & Business Media, 2013. 2

[FC09] FÉVOTTE C., CEMGIL A. T.: Nonnegative matrix factorizations as probabilistic inference in composite models. In *2009 17th European Signal Processing Conference* (2009), IEEE, pp. 1913–1917. 3

[FH12] FALCÓ A., HACKBUSCH W.: On minimal subspaces in tensor representations. *Foundations of computational mathematics 12*, 6 (2012), 765–803. 3

[GBP*20] GROSSET P., BIWER C. M., PULIDO J., MOHAN A. T., BISWAS A., PATCHETT J., TURTON T. L., ROGERS D. H., LIVESCU D., AHRENS J.: Foresight: Analysis that matters for data reduction. In *2020 SC20: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)* (Los Alamitos, CA, USA, nov 2020), IEEE Computer Society, pp. 1–15. 3

[GWC*08] GITTINGS M., WEAVER R., CLOVER M., BETLACH T., BYRNE N., COKER R., DENDY E., HUECKSTAEDT R., NEW K., OAKES W. R., ET AL.: The rage radiation-hydrodynamic code. *Computational Science & Discovery 1*, 1 (2008), 015005. 3

[HNMK04] HASEBE S., NAGUMO M., MURAMATSU S., KIKUCHI H.: Video key frame selection by clustering wavelet coefficients. In *2004 12th European Signal Processing Conference* (Sep 2004), p. 2303–2306. 2

[HZ10] HORÉ A., ZIOU D.: Image quality metrics: Psnr vs. ssim. In *2010 20th International Conference on Pattern Recognition* (2010), pp. 2366–2369. 3

[Jol86] JOLLIFFE I. T.: Principal components in regression analysis. In *Principal component analysis*. Springer, 1986, pp. 129–155. 2

[KB09] KOLDA T. G., BADER B. W.: Tensor decompositions and applications. *SIAM review 51*, 3 (2009), 455–500. 3

[LS99] LEE D. D., SEUNG H. S.: Learning the parts of objects by non-negative matrix factorization. *Nature 401*, 6755 (1999), 788–791. 2

[MLF*12] MA K., LIAO I., FRAZIER J., HAUSER H., KOSTIS H.: Scientific storytelling using visualization. *IEEE Computer Graphics and Applications 32*, 1 (2012), 12–19. 1

[MLF*16] MYERS K., LAWRENCE E., FUGATE M., BOWEN C. M., TICKNOR L., WOODRING J., WENDELBERGER J., AHRENS J.: Partitioning a large simulation as it runs. *Technometrics 58*, 3 (Jul 2016), 329–340. Publisher: Taylor Francis. 1, 2

[NPL*16] NOWICKI S. M., PAYNE A., LAROUR E., SEROUSSI H., GOELZER H., LIPSCOMB W., GREGORY J., ABE-OUCHI A., SHEPHERD A.: Ice Sheet Model Intercomparison Project (ISMIP6) contribution to CMIP6. *Geoscientific Model Development 9*, 12 (2016), 4521–4545. 4

[PG17] PATCHETT J., GISLER G.: Deep water impact ensemble data set. *Technical Report LA-UR-17-21595, Los Alamos National Laboratory* (2017). 3

[PT94] PAATERO P., TAPPER U.: Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics 5*, 2 (1994), 111–126. 2

[PXv*19] PORTER W. P., XING Y., VON OHLEN B. R., HAN J., WANG C.: A deep learning approach to selecting representative time steps for time-varying multivariate data. In *2019 IEEE Visualization Conference (VIS)* (2019), pp. 1–5. 1

[SPS*20] SUN S., PATTYN F., SIMON E. G., ALBRECHT T., CORNFORD S., CALOV R., DUMAS C., GILLET-CHAULET F., GOELZER H., GOLLEDGE N. R., GREVE R., HOFFMAN M. J., HUMBERT A., KAZMIERCZAK E., KLEINER T., LEGUY G. R., LIPSCOMB W. H., MARTIN D., MORLIGHEM M., NOWICKI S., POLLARD D., PRICE S., QUIQUET A., SEROUSSI H., SCHLEMM T., SUTTER J., VAN DE WAL R. S. W., WINKELMANN R., ZHANG T.: Antarctic ice sheet response to sudden and sustained ice-shelf collapse (ABUMIP). *Journal of Glaciology* (sep 2020), 1–14. 4

[TLS12] TONG X., LEE T., SHEN H.: Salient time steps selection from large scale time-varying data sets with dynamic time warping. In *IEEE Symposium on Large Data Analysis and Visualization (LDAV)* (2012), pp. 49–56. 1

[Tuc66] TUCKER L. R.: Some mathematical notes on three-mode factor analysis. *Psychometrika 31*, 3 (1966), 279–311. 3

[WSB03] WANG Z., SIMONCELLI E. P., BOVIK A. C.: Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems Computers, 2003* (2003), vol. 2, pp. 1398–1402 Vol.2. 3

[ZB02] ZHOU WANG, BOVIK A. C.: A universal image quality index. *IEEE Signal Processing Letters 9*, 3 (2002), 81–84. 3

[ZBSS04] ZHOU WANG, BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing 13*, 4 (2004), 600–612. 3