

Spatiotemporal Representation Learning for Short and Long Medical Image Time Series

Chengzhi Shen^{1(✉)}, Martin J. Menten^{1,7}, Hrvoje Bogunović², Ursula Schmidt-Erfurth², Hendrik P.N. Scholl^{3,4,5}, Sobha Sivaprasad⁶, Andrew Lotery⁷, Daniel Rueckert^{1,8}, Paul Hager^{*1}, Robbie Holland^{*8}

¹ Technical University of Munich, Germany

chengzhi.shen@tum.de

² Laboratory for Ophthalmic Image Analysis, Medical University of Vienna, Austria

³ Institute of Molecular and Clinical Ophthalmology Basel, Switzerland

⁴ Department of Ophthalmology, University of Basel, Switzerland

⁵ Department of Clinical Pharmacology, Medical University of Vienna, Austria

⁶ Moorfields Eye Hospital NHS Foundation Trust, London, United Kingdom

⁷ Clinical and Experimental Sciences, University of Southampton, United Kingdom

⁸ BioMedIA, Imperial College London, United Kingdom

Abstract. Analyzing temporal developments is crucial for the accurate prognosis of many medical conditions. Temporal changes that occur over short time scales are key to assessing the health of physiological functions, such as the cardiac cycle. Moreover, tracking longer term developments that occur over months or years in evolving processes, such as age-related macular degeneration (AMD), is essential for accurate prognosis. Despite the importance of both short and long term analysis to clinical decision making, they remain understudied in medical deep learning. State of the art methods for spatiotemporal representation learning, developed for short natural videos, prioritize the detection of temporal constants rather than temporal developments. Moreover, they do not account for varying time intervals between acquisitions, which are essential for contextualizing observed changes. To address these issues, we propose two approaches. First, we combine clip-level contrastive learning with a novel temporal embedding to adapt to irregular time series. Second, we propose masking and predicting latent frame representations of the temporal sequence. Our two approaches outperform all prior methods on temporally-dependent tasks including cardiac output estimation and three prognostic AMD tasks. Overall, this enables the automated analysis of temporal patterns which are typically overlooked in applications of deep learning to medicine. Code is available at <https://github.com/Leooo-Shen/tvrl>.

1 Introduction

Analyzing temporal developments in medical images is crucial to the practice of medicine. Clinicians analyze temporal change over *short* time scales, such as arrhythmia in the cardiac cycle in MR videos [3,22,23] or prenatal movement in ultrasound [27], in order to analyze dynamic physiological processes and function.

Clinicians also track the development of processes that evolve gradually, such as disease progression [21,19,18] or a patient’s response to treatment [5,14], over *long* time frames. In this longitudinal data, images are acquired at sparse and typically irregular intervals that may be months or even years apart. Despite the importance of both temporal dynamics and historical observations to the present diagnosis and prognosis of patients, the development of representation learning methods capable of learning these trends in medical data has thus far been overlooked.

Current state of the art approaches for spatiotemporal learning [13,28] were primarily developed on the Kinetics dataset [20], a human action recognition dataset which depicts only a single action or event per video. To predict this singular event, these methods prioritize the detection of video-level features that persist over time, and in doing so explicitly ignore changes between frames. As such, these methods are unsuitable for learning temporally variant features over the timeline. Furthermore, these methods were designed for videos with a fixed framerate and are thus unable to handle varying time intervals in irregularly sampled longitudinal sequences.

To address these issues, we contribute the following:

- We demonstrate that, while established strategies for representation learning in natural videos perform adequately on short cardiac videos, they are especially poor at modeling long-range developments in longitudinal retinal images due to the neglecting of temporally variant features.
- We then propose a simple clip-level contrastive learning strategy that uses a single clip with time embedding to encode varying time intervals, outperforming all natural video baselines.
- Finally, we propose a novel method that combines the efficacy of clip-level contrastive learning with a frame-level latent feature prediction task. Our method outperforms all natural video baselines in cardiac output estimation and three prognostic tasks for late stage AMD.

2 Related Work

2.1 Self-Supervised Video Representation Learning

In the image domain, contrastive learning has emerged as a strong and robust method for learning global imaging features without labels [7,15,16,8,4]. Similarly, contrastive approaches have achieved the state-of-the-art on the standard benchmark dataset for natural videos, Kinetics [20]. By minimizing the feature distance between differing segments, or clips, sampled from different time points in the same video, these approaches learn to extract the single human action recorded over the span of each video [28,13,35]. Specifically, Qian et al. [28] adopt a probabilistic function to vary the time between the two sampled clips, Feichtenhofer et al. [13] sample up to four clips, and Wang et al. sample clips of varying lengths [35]. However, by learning features that persist across the entire




Fig. 1: Clinicians use spatiotemporal data to observe temporal variations over both *short* and *long* time scales. To observe dynamic physiological processes such as a beating heart, short cardiac MR videos can be captured (top). Tracking long term developments, such as disease progression in retinal OCT scans, require longitudinal acquisitions that typically occur at irregular intervals of years (bottom). Modeling and extrapolating the trajectory of historical change is crucial for the prognosis of late stage disease.

video, these methods ignore temporal variations which are associated with many medical conditions, including those involved in longitudinal disease progression.

To address this limitation, some recent works modify the contrastive framework to instead maximize the distance between the sampled clips [10,40], or by parameterizing the video with an evolving stochastic process [26,38]. Yang et al. in Latent time navigation (LTN) [37] introduce an extra orthogonal basis that explicitly models the temporal shift between sampled clips of the same video. Generative approaches, such as VideoMAE [34] and v-JEPA [2], implicitly model temporal variations by masking and reconstructing a subset of frame-level features. Still, it is unclear whether the ability of these methods to model temporal developments on natural videos translates to spatiotemporal medical data.

2.2 Representation Learning for Spatiotemporal Medical Data

In the medical field, spatiotemporal representation learning has many applications, such as organ segmentation [31], physiological development analysis [39], and disease progression prediction [9]. Some works leverage the assumption that certain disease progressions share similar trajectories in the latent space [25,24].

In addition, [29] and [36] leverage redundancies between successive brain scans to improve segmentation accuracy. However, there is no unified approach that learns self-supervised representations of spatiotemporal data in the medical domain. In this work, we propose a method that combines the efficacy of contrastive approaches with a masked latent reconstruction loss to model trajectories of frame-level temporal changes that are crucial for downstream medical analysis.

3 Materials and Methods

3.1 Short and Long Time Scale Datasets

Cardiac MR video To study temporal variation over short time frames we chose to use short-axis cardiac MR videos sourced from the UK Biobank (UKBB) [32] (see Figure 1). The dataset contains 46,137 videos where each captures at least one full contraction, from end diastole to end systole, in 50 frames taken over 1.6 seconds. After resampling the images to 1.8 mm in the sagittal and coronal dimensions, we used segmentation masks to crop a region of 128×128 about the center of mass of the hearts.

Longitudinal Retinal OCT To study temporal change over long time frames using sparse, irregular data, we use an in-house retinal OCT dataset tracking patients with age-related macular degeneration (see Figure 1) collected by the Southampton Eye Unit in the scope of the PINNACLE study [33]. The dataset includes 48,825 OCT scans, monitoring 6,368 eyes in 3,498 patients. On average, each eye was scanned 7.7 times over 1.9 years, where eyes with AMD were observed for an average of 2.3 years.

3.2 Spatiotemporal Encoder

As the basis of all baselines and our new approaches, we use a spatiotemporal encoder that extracts representation of the sequence (see Figure 2 A). For each image x_i in the sequence, we employ a spatial Vision Transformer (ViT) [11] E_S to extract a feature vector per image $z_i^s = E_S(x_i)$ using a learnable CLS token. The temporal Transformer E_T then models temporal change of features and generates a global representation of the sequence $z^t = E_T(z_1^s \dots z_i^s)$ using a distinct CLS token. Our two-encoder design efficiently operates on spatiotemporal data by reducing attention token length instead of processing all video patches at once.

3.3 Clip-level Contrastive Learning

In Figure 2 B., we illustrates that established contrastive learning methods prioritize temporally-invariant features by minimizing the feature distance between *multiple* temporally distant clips from the same sequence [28,13]. To model temporal variations rather than temporal constants, we construct a clip-level contrastive approach that augments only a *single* randomly sampled clip c . We then




Fig. 2: **A.** Our spatiotemporal encoder extracts representation of the sequence using the CLS token of the temporal Transformer E_T . **B.** Standard contrastive approaches prioritize learning features that persist across different segments of the same sequence (left). Our approaches use a single clip to construct contrastive pairs combined with time embedding (middle) and a frame-level feature predictive approach (right) to model temporal variation over the sequence.

generate a pair of augmented views of our clip, c_i and c_j , where the augmentations in each are applied consistently across timepoints. We pass these views through the spatiotemporal encoder to extract two clip-level contrastive tokens z_i^t and z_j^t , which are learnable CLS tokens in the temporal encoder E_T . Finally, we optimize the NT-Xent [7] loss over the N samples in the batch using equation 1, where sim is the cosine similarity, τ is the temperature, p is a non-linear projection head following the design of [7].

$$L^C = -\log \frac{\exp(sim(p(z_i^t), p(z_j^t))/\tau)}{\sum_{k=1, k \neq i}^{2N} \exp(sim(p(z_i^t), p(z_k^t))/\tau)} \quad (1)$$

3.4 Time Embedding

To model the irregular time intervals between scans, we propose time embedding (TE) for the temporal Transformer E_T (see Figure 2 B). For each clip, we calculate the relative time between each scan and the first scan. We then apply a learnable non-linear mapping based on Ho et al. [17] to map the time points to the feature dimension of the model. Similar to a position embedding, we add TE to the input frame-level features in E_T . This enables the temporal encoder to estimate the rate of change in longitudinal sequences with irregular sampling.

3.5 Temporally-Variant Feature Prediction

Finally, to explicitly model frame-level temporal variations, we propose a temporal feature prediction task inspired by Assran et al. [1] (see Figure 2 B). Specifically, given a sequence of frame-level tokens output by our spatial encoder, we

replace a random subset Y with a learnable mask token using a masking ratio of 0.15. We then use a projection head q on top of the temporal encoder output tokens to reconstruct the masked tokens. We then compute the cosine similarity loss using equation 2 between the reconstructed tokens Y' and the original spatial tokens Y , where m denotes the number of masked tokens.

$$L^M = \frac{1}{m} \sum_{i=1}^m \text{sim}(Y'_i, Y_i) \quad (2)$$

The final loss, termed temporally-variant representation loss (TVRL), combines clip-level contrastive learning with frame-level feature prediction and is defined in equation 3, where λ is a weight term empirically set to 0.5.

$$L = (1 - \lambda)L^C + \lambda L^M \quad (3)$$

4 Experiments and Results

4.1 Experimental Setup

Pretraining Protocol We evaluate our two pretraining approaches: the clip-level contrastive strategy described in section 3.3 which is denoted as *cSimCLR*, and our TVRL extension introduced in section 3.5. For *cSimCLR*, we also evaluate the effect of adding time embedding (TE) from section 3.4 and denote it as *cSimCLR-TE*. Note that our TVRL strategy does not use TE. We leverage standard SimCLR augmentations [7] and discard those unsuitable for grayscale images. We compare our approaches against the leading spatiotemporal representation learning strategies, which includes CVRL [28], LTN [37], and VideoMAE [34]. All pretraining strategies use a spatial encoder E_S parameterized by a ViT-S [11] with 384 hidden dimensions, six heads, and 12 attention layers, while the temporal encoder E_T shares the same architecture but with three layers of attention. Both encoders are randomly initialized before pretraining. In total the model has 27.7M parameters.

Before pretraining, we first partition each dataset into training, validation, and test sets, comprising 70%, 15%, and 15% of the data, respectively, ensuring that each patient appears in only one subset. Each model accepts eight images as input. We sample every other frame from the cardiac videos, and every successive scan in the retinal longitudes. We employ separate projection heads p and q for our contrastive and masked reconstruction losses, following the two-layer configuration in [7]. All models are pretrained for 200 epochs with a batch size of 256 using the AdamW optimizer, and a cosine annealing learning rate of $2 \cdot 10^{-4}$ with 20 epochs of linear warm-up.

Finetuning Protocol We follow the standard linear probing protocol [6] and finetune each pretrained model by freezing the network and appending a new trainable CLS token to the temporal Transformer which serves as input to a single linear layer. We finetune each model with 1% and 100% of the data, and report the area under the receiver operating characteristic curve (AUC) for classification tasks and mean absolute error (MAE) for regression tasks, with mean and standard deviation over five random seeds. We finetune with AdamW for 100 epochs with a learning rate of 10^{-3} and a batch size of 256 and report test performance of the best validation epoch. As baselines to the aforementioned strategies, we also evaluate a fully supervised spatiotemporal encoder trained end-to-end, in addition to a 3D ResNet50 [12] (31.6M parameters) pre-trained on the Kinetics-400 dataset [20]. During inference, in sequences with more than eight images, we adapt the common practice in evaluating video models [13,28] and apply a sliding window with 50% overlap to contiguous clips before averaging these predictions.

4.2 Downstream Evaluation

We first evaluate different pretraining strategies on a set of control tasks that are marginally associated with temporal change. In cardiac MR video, we predict coronary artery disease (CAD) and fibrillation from ICD codes recorded during hospital admissions of all subjects. From retinal OCT longitudinal data we predict the presence of late vs. early AMD given the present and historical set of scans of each patient.

We then evaluate the models on tasks that are strongly associated with temporal change. In the short time scale, we estimate left ventricular ejection fraction (LVEF) and total cardiac output, which measures the volume of blood pumped in one minute. Finally, we focus on three OCT-based prognostic tasks that require extrapolating long term trends from the patient’s history of disease progression. These involve predicting the development of late stage wet AMD (defined by choroidal neovascularisation), late stage dry AMD (defined by cRORA $> 1000\mu m$ [30]) and retinal scarring and fibrosis post treatment, all within the next 3 years.

4.3 Simple Clip-level Contrastive Learning Outperforms Prior Approaches

In Table 1 and Table 2, we find in all tasks that the time-invariant representations learned by CVRL [28] degrade performance compared to our simple time-dependent cSimCLR [7] approach using a single clip. This was most pronounced on the temporally dependent tasks such as prognosis of late dry AMD (74.3% vs. 67.5% AUC). While LTN [37] and VideoMAE [34] improve the modeling of temporal variations, both methods perform comparably to or worse than clip-level cSimCLR across all tasks in both the short and long time scales.

Table 1: Performance on diagnosis and prognosis on the long term AMD prediction tasks in longitudinal retinal OCT. Modeling variable intervals between scans is especially beneficial for prognosis tasks that requires the extrapolation of trajectories of disease progression captured in historical scans.

Model	Non-temporal control task		Temporally dependent prognostic tasks					
	Late vs. Early AUC (%) $\uparrow$		Late Dry AMD AUC (%) $\uparrow$		Late Wet AMD AUC (%) $\uparrow$		Scarring and Fibrosis AUC (%) $\uparrow$	
	1%	100%	1%	100%	1%	100%	1%	100%
Supervised	60.2 $\pm$ 1.1	85.9 $\pm$ 0.3	50.2 $\pm$ 1.1	64.7 $\pm$ 0.5	54.5 $\pm$ 0.2	59.2 $\pm$ 0.9	49.1 $\pm$ 0.8	52.9 $\pm$ 0.6
R3D [12] (Kinetics)	54.3 $\pm$ 1.1	73.8 $\pm$ 0.3	51.8 $\pm$ 5.1	54.4 $\pm$ 0.6	52.2 $\pm$ 2.1	64.9 $\pm$ 0.8	50.6 $\pm$ 5.1	64.4 $\pm$ 1.3
CVRL [28]	75.3 $\pm$ 1.7	86.2 $\pm$ 0.3	52.4 $\pm$ 2.3	67.5 $\pm$ 5.3	56.8 $\pm$ 6.5	64.7 $\pm$ 2.6	59.9 $\pm$ 3.9	71.2 $\pm$ 1.7
LTN [37]	76.9$\pm$0.6	85.7 $\pm$ 0.1	61.8$\pm$3.4	63.1 $\pm$ 1.5	57.6 $\pm$ 6.1	64.8 $\pm$ 4.1	71.6$\pm$0.9	72.9 $\pm$ 0.8
VideoMAE [34]	72.0 $\pm$ 1.1	81.8 $\pm$ 0.1	58.7 $\pm$ 4.1	56.8 $\pm$ 0.6	54.3 $\pm$ 9.2	56.2 $\pm$ 7.8	61.4 $\pm$ 3.4	67.3 $\pm$ 1.1
cSimCLR	76.3 $\pm$ 2.7	88.6$\pm$0.4	53.6 $\pm$ 2.8	74.3 $\pm$ 1.3	58.3 $\pm$ 4.2	66.8 $\pm$ 1.3	68.9 $\pm$ 4.5	72.2 $\pm$ 0.6
cSimCLR-TE	73.0 $\pm$ 0.9	88.3 $\pm$ 0.1	50.4 $\pm$ 5.6	76.8 $\pm$ 1.5	58.7$\pm$5.5	66.5 $\pm$ 1.1	64.4 $\pm$ 3.3	77.4$\pm$1.5
TVRL	76.1 $\pm$ 1.6	87.7 $\pm$ 0.1	55.7 $\pm$ 4.0	80.2$\pm$0.8	53.5 $\pm$ 4.8	67.8$\pm$2.3	65.9 $\pm$ 3.0	73.2 $\pm$ 1.4

Table 2: In short term cardiac video, standard natural video approaches perform comparably on diagnosis but underperform in estimating cardiac output, which has the strongest temporal component.

Model	Diagnostic control tasks				Temporally dependent biometric tasks			
	CAD AUC (%) $\uparrow$		Fibrillation AUC (%) $\uparrow$		LVEF MAE $\downarrow$		Cardiac Output MAE $\downarrow$	
	1%	100%	1%	100%	1%	100%	1%	100%
Supervised	56.3 $\pm$ 0.8	64.9 $\pm$ 3.1	53.2 $\pm$ 7.2	62.7 $\pm$ 0.6	4.75 $\pm$ 0.01	3.58 $\pm$ 0.21	4.33 $\pm$ 3.38	1.80 $\pm$ 0.04
R3D [12] (Kinetics)	52.9 $\pm$ 1.8	64.1 $\pm$ 4.2	54.6 $\pm$ 0.2	66.3 $\pm$ 2.1	8.27 $\pm$ 1.24	7.13 $\pm$ 0.58	2.77 $\pm$ 0.09	1.86 $\pm$ 0.05
CVRL [28]	59.4 $\pm$ 5.8	69.4 $\pm$ 0.2	54.4 $\pm$ 0.9	70.3 $\pm$ 0.3	5.14 $\pm$ 0.24	3.95 $\pm$ 0.03	1.85 $\pm$ 0.02	1.76 $\pm$ 0.08
LTN [37]	56.6 $\pm$ 8.7	69.4 $\pm$ 0.2	54.0 $\pm$ 1.3	70.5$\pm$0.3	4.83 $\pm$ 0.04	3.88$\pm$0.08	1.85 $\pm$ 0.01	1.68 $\pm$ 0.03
VideoMAE [34]	57.1 $\pm$ 4.3	66.9 $\pm$ 0.6	55.1 $\pm$ 3.1	63.1 $\pm$ 0.9	6.89 $\pm$ 0.05	5.28 $\pm$ 0.07	1.98 $\pm$ 0.02	1.86 $\pm$ 0.06
cSimCLR	60.9 $\pm$ 0.4	69.4 $\pm$ 0.1	54.5 $\pm$ 0.4	70.4 $\pm$ 0.6	4.88 $\pm$ 0.35	3.89 $\pm$ 0.04	1.88 $\pm$ 0.01	1.69 $\pm$ 0.12
cSimCLR-TE	62.6$\pm$0.9	67.9 $\pm$ 0.5	57.3$\pm$0.2	66.7 $\pm$ 0.1	4.62$\pm$0.01	4.05 $\pm$ 0.01	1.84$\pm$0.01	1.63$\pm$0.01
TVRL	61.7 $\pm$ 0.1	69.3 $\pm$ 0.1	55.5 $\pm$ 0.7	69.8 $\pm$ 0.1	4.86 $\pm$ 0.02	4.03 $\pm$ 0.01	1.92 $\pm$ 0.02	1.65 $\pm$ 0.02

4.4 Explicit Temporal Modeling Improves Long-term Prognostic Tasks

The clip-level contrastive approach with time embedding (cSimCLR-TE) outperforms all prior methods on the irregularly sampled retinal OCT longitudes. Add TE notably improves the prognosis of scarring from 72.2% to 77.4% AUC as shown in Table 1. It also performed best in estimating the cardiac output, achieving 1.63 MAE as shown in Table 2. Finally, TVRL further boosts the performance on prognosis for late dry AMD to 80.2% AUC and late wet AMD to 67.8% AUC. Additionally, in Figure A.1 and A.2 in the supplementary, we visualize the feature trajectories of sequences from both datasets learned by TVRL.

Regarding Table 2, on cardiac videos with regular frame intervals, TE is expected to provide no further benefit over positional embedding, leading to comparable performance between cSimCLR and cSimCLR-TE. Similarly, TVRL

performs comparably to contrastive baselines, as cardiac videos contain limited temporal variation within a single cardiac cycle.

5 Discussion and Conclusion

In this paper, we identified two common variants of spatiotemporal data used in clinical practice. We found that while state-of-the-art spatiotemporal learning methods developed for natural videos perform comparably in short term analysis on cardiac MR video, they failed to track long term developments and disease trajectories in longitudinal series, leading to poor prognostic utility for AMD. To address this, we propose a simple clip-level contrastive learning strategy that leverages time embeddings in irregular and variable length time series, and a new temporally-variant approach that explicitly models frame-level variation. Models pretrained with our strategies had improved assessment of cardiac output, and substantially improved prognosis for three variants of late stage AMD.

One limitation of this study is that we have not fully determined in which scenario to use each of our two approaches. Future work aims to expand the number and diversity of tasks to provide better guidance.

In conclusion, we envision that our approaches will create representations that leverage the full range of temporal dynamics and patient history that are typically neglected in current diagnostic and prognostic solutions.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Assran, M., et al.: Self-supervised learning from images with a joint-embedding predictive architecture. In: CVPR. pp. 15619–15629 (2023)
2. Bardes, A., et al.: V-jepa: Latent video prediction for visual representation learning (2023)
3. Bernard, O., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? IEEE TMI **37**(11), 2514–2525 (2018)
4. Caron, M., et al.: Emerging properties in self-supervised vision transformers. In: CVPR. pp. 9650–9660 (2021)
5. Chainani-Wu, N., et al.: Oral lichen planus: patient profile, disease progression and treatment responses. The Journal of the American Dental Association **132**(7), 901–909 (2001)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. CoRR **abs/2002.05709** (2020), <https://arxiv.org/abs/2002.05709>
7. Chen, T., et al.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607. PMLR (2020)
8. Chen, X., et al.: Exploring simple siamese representation learning. In: CVPR. pp. 15750–15758 (2021)
9. Couronné, R., et al.: Longitudinal self-supervision to disentangle inter-patient variability from disease progression. In: MICCAI. pp. 231–241. Springer (2021)
10. Dave, I., et al.: Tclr: Temporal contrastive learning for video representation. CVIU **219**, 103406 (2022)
11. Dosovitskiy, A., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
12. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: ICCV. pp. 6202–6211 (2019)
13. Feichtenhofer, C., et al.: A large-scale study on unsupervised spatiotemporal representation learning. In: CVPR. pp. 3299–3309 (2021)
14. Graham, L.J., et al.: Current approaches and challenges in monitoring treatment responses in breast cancer. Journal of Cancer **5**(1), 58 (2014)
15. Grill, J.B., et al.: Bootstrap your own latent—a new approach to self-supervised learning. NeurIPS **33**, 21271–21284 (2020)
16. He, K., et al.: Momentum contrast for unsupervised visual representation learning. In: CVPR. pp. 9729–9738 (2020)
17. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
18. Holland, R., et al.: Clustering disease trajectories in contrastive feature space for biomarker proposal in age-related macular degeneration. In: MICCAI. pp. 724–734. Springer (2023)
19. Jensen, A.B., et al.: Temporal disease trajectories condensed from population-wide registry data covering 6.2 million patients. Nature communications **5**(1), 4022 (2014)
20. Kay, W., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
21. Koch, L., et al.: Fear of recurrence and disease progression in long-term (≥ 5 years) cancer survivors—a systematic review of quantitative studies. Psycho-oncology **22**(1), 1–11 (2013)

22. Komatsu, M., et al.: Detection of cardiac structural abnormalities in fetal ultrasound videos using deep learning. *Applied Sciences* **11**(1), 371 (2021)
23. Mondéjar-Guerra, V., et al.: Heartbeat classification fusing temporal and morphological information of ecgs via ensemble of classifiers. *Biomedical Signal Processing and Control* **47**, 41–48 (2019)
24. Ouyang, J., et al.: Self-supervised longitudinal neighbourhood embedding. In: *MICCAI*. pp. 80–89. Springer (2021)
25. Ouyang, J., et al.: Self-supervised learning of neighborhood embedding for longitudinal mri. *Medical image analysis* **82**, 102571 (2022)
26. Park, J., et al.: Probabilistic representations for video contrastive learning. In: *CVPR*. pp. 14711–14721 (2022)
27. Pugash, D., et al.: Prenatal ultrasound and fetal mri: the comparative value of each modality in prenatal diagnosis. *European journal of radiology* **68**(2), 214–226 (2008)
28. Qian, R., et al.: Spatiotemporal contrastive video representation learning. In: *CVPR*. pp. 6964–6974 (2021)
29. Ren, M., et al.: Local spatiotemporal representation learning for longitudinally-consistent neuroimage analysis. *NeurIPS* **35**, 13541–13556 (2022)
30. Sadda, S.R., et al.: Consensus definition for atrophy associated with age-related macular degeneration on oct: classification of atrophy report 3. *Ophthalmology* **125**(4), 537–548 (2018)
31. Shin, H.C., et al.: Stacked autoencoders for unsupervised feature learning and multiple organ detection in a pilot study using 4d patient data. *IEEE TPAMI* **35**(8), 1930–1943 (2012)
32. Sudlow, C., et al.: Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine* **12**(3), e1001779 (2015)
33. Sutton, J., et al.: Developing and validating a multivariable prediction model which predicts progression of intermediate to late age-related macular degeneration—the pinnacle trial protocol. *Eye* pp. 1–9 (2022)
34. Tong, Z., et al.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *NeurIPS* **35**, 10078–10093 (2022)
35. Wang, J., et al.: Long-short temporal contrastive learning of video transformers. In: *CVPR*. pp. 14010–14020 (2022)
36. Wei, J., et al.: Consistent segmentation of longitudinal brain mr images with spatiotemporal constrained networks. In: *MICCAI*. pp. 89–98. Springer (2021)
37. Yang, D., et al.: Self-supervised video representation learning via latent time navigation. *arXiv preprint arXiv:2305.06437* (2023)
38. Zhang, H., et al.: Modeling video as stochastic processes for fine-grained video representation learning. In: *CVPR*. pp. 2225–2234 (2023)
39. Zhao, Q., et al.: Longitudinal correlation analysis for decoding multi-modal brain development. In: *MICCAI*. pp. 400–409. Springer (2021)
40. Zhuang, C., et al.: Unsupervised learning from video with deep neural embeddings. In: *CVPR*. pp. 9563–9572 (2020)