

# Countering Multi-modal Representation Collapse through Rank-targeted Fusion

Seulgi Kim, Kiran Kokilepersaud, Mohit Prabhushankar, Ghassan AlRegib  
Georgia Institute of Technology, Atlanta, GA, USA  
{seulgi.kim, kpk6, mohit.p, alregib}@gatech.edu

## Abstract

Multi-modal fusion methods often suffer from two types of representation collapse: feature collapse where individual dimensions lose their discriminative power (as measured by eigenspectra), and modality collapse where one dominant modality overwhelms the other. Applications like human action anticipation that require fusing multifarious sensor data are hindered by both feature and modality collapse. However, existing methods attempt to counter feature collapse and modality collapse separately. This is because there is no unifying framework that efficiently addresses feature and modality collapse in conjunction. In this paper, we posit the utility of effective rank as an informative measure that can be utilized to quantify and counter both the representation collapses. We propose Rank-enhancing Token Fuser, a theoretically grounded fusion framework that selectively blends less informative features from one modality with complementary features from another modality. We show that our method increases the effective rank of the fused representation. To address modality collapse, we evaluate modality combinations that mutually increase each others' effective rank. We show that depth maintains representational balance when fused with RGB, avoiding modality collapse. We validate our method on action anticipation, where we present R3D, a depth-informed fusion framework. Extensive experiments on NTURGBD, UTKinect, and DARai demonstrate that our approach significantly outperforms prior state-of-the-art methods by up to 3.74%. Our code is available at: <https://github.com/olivesgatech/R3D>.

## 1. Introduction

Modern data collection paradigms [33] utilize multifarious sensor modalities to represent an environment, as each modality captures unique aspects of a scene. For example, RGB excels at conveying visual cues in objects such as color or texture, while depth captures geometric structures and directional relationships between objects (Figure 1 (a-b)). Hence, multi-modal learning [70] is indispensable for human-centric tasks such as action anticipation [7, 11, 25,

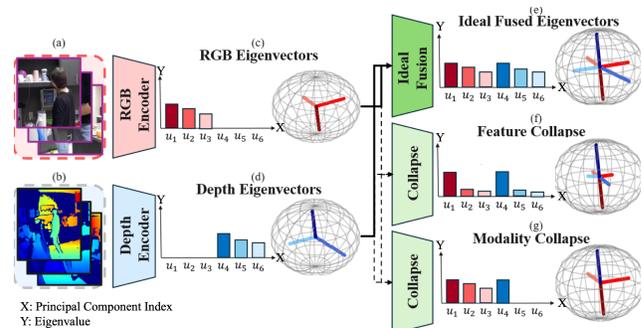


Figure 1. This is a toy figure describing feature and modality collapse using spectral decomposition. (e): The ideal fused representation preserves complementary eigenvectors from both data modalities. (f): In contrast, feature collapse occurs when the fused representation varies along a subset of eigenvectors. (g): Modality collapse occurs when one modality dominates and suppresses the contribution of the eigenvectors of the other modality.

36, 41, 76], video understanding [28, 32], and action recognition [67], where single modalities fall short in expressing subtle semantic variations within or across scenes. Recent multi-modal fusion methods have achieved empirical success through contrastive losses [21, 54], attention blocks [69], or various specialized objectives [18]. However, they often rely on learning representations via indirect pre-text tasks or alignment losses [48], rather than directly targeting the informative content in each modality. This lack of *targeted information fusion* makes multi-modal fusion models susceptible to the common pitfalls: feature collapse [44] and modality collapse [12]. A central technical challenge we tackle in this paper is the lack of a unifying information framework that targets both feature and modality collapse.

We illustrate *targeted information fusion* in Figure 1. In (c) and (d), RGB and Depth are projected into their respective representation spaces using their corresponding encoders. Each space is visualized with a bar plot of eigenvalues associated with eigenvectors  $[u_1, \dots, u_6]$  and a toy 3D spherical projection of the data points. Intuitively, the eigenvectors of each space represent the principal directions of variation in modality-specific representations. For the purposes of discussion, we define each eigenvector as a combination of projections from individual features (chan-

nels). Under this view, a channel is considered *informative* if it contributes significantly to key eigenvector directions, indicating strong alignment with the principal variations in the space. Thus, we interpret eigenvectors in Figure 1 as *information components*, that is, the feature directions encoding useful content from each modality.

An ideal representation space after fusion should reflect Figure 1(e), where each information component is equally represented, resulting in a space that integrates feature directions encoding useful content from both modalities. However, in reality, different types of collapse can occur in the fused representation space. In Figure 1(f), we visualize feature collapse where certain information components become limited in their contribution to the overall space after fusion. This limits the overall feature diversity and the generalizability of the model [10, 40]. In Figure 1(g), we observe modality collapse [12] where information components are expressed for one modality over the other, limiting the sharing of complementary features between the modalities.

To address these challenges, we propose an ideal fusion strategy for action anticipation that addresses both representation collapse. First, to counter feature collapse, we aim to increase the effective rank [57] of the fused representation matrix, as it serves as a proxy for the diversity of information. This is motivated by prior work defining the effective rank as matrix entropy [55, 62, 74]. Based on this, we develop *Rank-enhancing Token fuser*, a theoretically grounded fusion framework that selectively blends *less informative* channels (i.e., channels that contribute little to the principal eigenvector directions) with complementary features from another modality. Second, we posit that compatible modalities mutually increase each other’s effective rank. Hence, to address modality collapse, we perform an analysis of different modalities to pair with RGB for action anticipation. We demonstrate through a relative rank analysis that the depth modality results in the most balanced feature space between both modalities. We use this intuition to develop R3D (Rank-enhancing fusion in 3D), the first depth-informed action anticipation architecture. We demonstrate performance improvements across a wide variety of benchmarks. Our main contributions are summarized as follows:

1. **Rank-enhancing Token Fuser:** We are the first to formulate multi-modal fusion as a problem of *rank-targeted fusion* to simultaneously address feature collapse and modality collapse. We provide theoretical conditions under which selective channel blending provably increases effective rank and correspondingly prevents representation collapse.
2. **Depth-aware 3D Action Anticipation:** We present R3D, the first depth-informed framework for 3D action anticipation and show that depth is the most complemen-

tary modality to pair with RGB for preserving modality-specific features.

3. **State-of-the-art performance:** R3D achieves up to 3.74% performance improvement on NTURGBD, UTKinect, and DARai, setting a new benchmark for multimodal action anticipation.

## 2. Related Works

### 2.1. Action Anticipation

The availability of large-scale video datasets [14, 24, 38, 42, 46] has spurred significant progress in addressing the challenge of action anticipation. These methods can be broadly categorized into two types. First, short-term action anticipation focuses on predicting a single future action that will occur within a few seconds [17, 19, 58]. In contrast, long-term action anticipation seeks to forecast an extended sequence of future actions from a long-range video, aiming to predict events far into the future [1, 2, 23, 28, 36, 58]. To predict long-term action, recent advances in action anticipation leverage multi-modal inputs to capture complementary contextual cues [43]. For instance, [71, 75] explicitly integrated object features alongside RGB data. [76] leveraged audio to capture off-camera events and to disambiguate visually similar actions. [8, 36] integrates text to provide rich and fine-grained semantic context. As humans perform actions in a 3D world, prior studies have explored the use of human pose as a complementary modality [7, 11, 15, 25, 41], often relying on motion capture systems or pose estimation pipelines. In contrast, we are the first to introduce raw depth data as a multi-modal input without requiring additional motion capture hardware, making it practical and deployable in real-world settings where commodity RGB-D cameras are readily available.

### 2.2. Multi-modal Fusion Method

A core challenge in multi-modal learning is the effective fusion of heterogeneous data streams [4, 13, 27, 29, 37]. Fusion strategies are broadly categorized based on their methodology of fusion. Aggregation-based fusion strategies combine features through concatenations, summation, or attention pooling [26, 53, 56, 61, 73]. Alignment-based fusion methods temporally or semantically align modality representations [8, 29, 36, 37, 63, 71, 76]. Finally, hybrid approaches integrate both aggregation and alignment techniques [5]. Despite their architectural advances, these fusion strategies often suffer from feature redundancy where redundant or weakly informative features from one modality can overwhelm the shared representation space, resulting in suboptimal fusion [48]. Our method addresses this limitation by encouraging diverse representations, as evidenced by an increased effective rank in the joint embedding space.

## 2.3. Approximating Information through Rank

The technical novelty in this paper is targeted multi-modal feature fusion. We utilize the effective rank of the feature matrix as a measure of information content. [55, 62] identify that some measure of the uniformity of the singular value spectrum, such as effective rank [57], is representative of the entropy of a matrix [74], and its information content. Intuitively, a more uniform singular value distribution implies that the representation of the data varies along a greater number of feature directions, thus exhibiting greater diversity and countering potential representational collapse [31]. To capture this eigenvalue uniformity as a metric, the effective rank computes the entropy of the normalized eigenspectrum of a matrix  $Z$  with associated eigenvalues  $[\sigma_1, \dots, \sigma_r]$  as follows:

$$\text{ERank}(Z) := \exp \left( - \sum_{j=1}^{\text{rank}(Z)} p_j \log p_j \right), \quad p_j = \frac{\sigma_j(Z)}{\|Z\|_*}. \quad (1)$$

Due to these useful properties, measuring the rank of a representation space has gained traction in various research areas. Within self-supervised learning, [3, 20, 60] use some measure of the eigenvalue spectra as a way to identify better performing models without the need for task specific fine-tuning. [40] show that adapting a loss during training based on rank measurements can improve downstream performance. Other works devise loss functions with the explicit goal of increasing rank [6, 39] and manipulating eigenvalues directly [35]. Despite the potential of integrating rank into training paradigms, there is limited work on applying rank as a mechanism to improve modality fusion. [12] analyze representational collapse in multi-modal architectures through effective rank measurements. However, their method does not involve directly integrating the rank and instead relies on discriminator networks. In contrast, our work is the first to demonstrate that rank information can directly inform targeted feature fusion.

## 3. Analysis

### 3.1. Theory: Rank-enhancing Multi-modal Fusion

We introduce a novel *Rank-enhancing Token Fuser (RTF)* that selectively blends *less informative* channels with features from another modality in a way that provably increases the effective rank of the resulting representation. We show the proof in Section A.

**Notation.** Let  $X \in \mathbb{R}^{T \times D}$  denote the representation matrix of a modality, where  $T$  is the number of timesteps and  $D$  is the channel (feature) dimension. Let the singular value decomposition (SVD) of  $X$  be  $X = U\Sigma V^T$ , where  $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$  contains singular values in decreasing order,  $V = [v_1, \dots, v_r]^T \in \mathbb{R}^{r \times D}$  contains the corresponding right singular vectors. We refer to these right

singular vectors as the principal directions of the representation. Let  $x_c \in \mathbb{R}^T$  denote the  $c$ -th column of  $X$ . Similarly, let  $Y \in \mathbb{R}^{T \times D}$  denote features from another modality with columns  $y_c \in \mathbb{R}^T$ . These serve as candidates for enhancing *less informative* channels in  $X$  through selective blending.

**Definition 1. Channel Informativeness.** We define the *informativeness* of channel  $c$  as its contribution to the top- $k$  singular vectors of the representation:  $I_c := \sum_{i=1}^k \sigma_i^2 v_{i,c}^2$ , where  $v_{i,c}$  is the  $c$ -th component of the  $i$ -th right singular vector  $v_i$ . Channels with low  $I_c$  are considered *less informative* and are potential candidates for fusion with complementary channels from other modalities. Let  $\mathcal{C}_{\text{low}} := \{c \in \{1, \dots, D\} \mid I_c \leq \eta\}$  denote the set of *low informativeness channels*, where  $\eta > 0$  is a predefined threshold.

**Definition 2. Fusion.** We define a fused representation  $X' \in \mathbb{R}^{T \times D}$  by selectively blending channels in  $\mathcal{C}_{\text{low}}$  with signals from modality  $Y$ :

$$x'_c = \begin{cases} \alpha_c x_c + (1 - \alpha_c) y_c & \text{if } c \in \mathcal{C}_{\text{low}}, \\ x_c & \text{otherwise,} \end{cases}$$

where  $\alpha_c \in [0, 1]$  is a learnable channel-wise blending coefficient. The goal is to enhance underutilized directions while preserving already informative ones.

**Theorem 3.1 (Channel Fusion Increases Effective Rank).** Let  $u_1, \dots, u_k \in \mathbb{R}^T$  denote the top- $k$  left singular vectors of  $X$ , and define  $\delta_k := \sigma_k - \sigma_{k+1}$  as the singular value gap, which quantifies the separation between the dominant subspace (top- $k$ ) and the residual space.

Assume:

1. **(Bounding condition)** For all  $c \in \mathcal{C}_{\text{low}}$ ,  $y_c$  is zero-mean and satisfies  $\|y_c\|_2 \leq \beta$  for some constant  $\beta > 0$ .
2. **(Non-trivial modification in the representation)** There exists  $\exists \epsilon > 0$  such that  $\sum_{c \in \mathcal{C}_{\text{low}}} \|x'_c - x_c\|_2^2 \geq \epsilon$
3. **(Dominant subspace preservation)** The updated representation should not distort the dominant subspace of  $X$ . This is ensured by requiring the update to only target low-informativeness channels. Specifically, the threshold  $\eta$  is small enough to satisfy  $\sqrt{\eta} \leq \min \left( \frac{\delta_k}{3\sqrt{|\mathcal{C}_{\text{low}}|}}, \frac{\epsilon}{4|\mathcal{C}_{\text{low}}|\beta} \right)$
4. **(Bounded Dominant Subspace Alignment)** The injected channels  $y_c$  are not perfectly aligned with the dominant subspace of  $X$ . Formally, the projection of each injected channel  $y_c$  to the dominant subspace of  $X$   $\mathcal{U}_k = \text{span}\{u_1, \dots, u_k\}$  is bounded by its projection onto the remaining subspace. i.e.,  $\sum_{i=1}^k |\langle y_c, u_i \rangle|^2 \leq \gamma \sum_{i=k+1}^r |\langle y_c, u_i \rangle|^2$ , where  $0 \leq \gamma < 1$ .

Then  $y_c$  introduces novel directions in the feature space of  $X$ , hence the effective rank satisfies  $\text{ERank}(X') > \text{ERank}(X)$ , where the equation of  $\text{ERank}(X)$  follows (1).

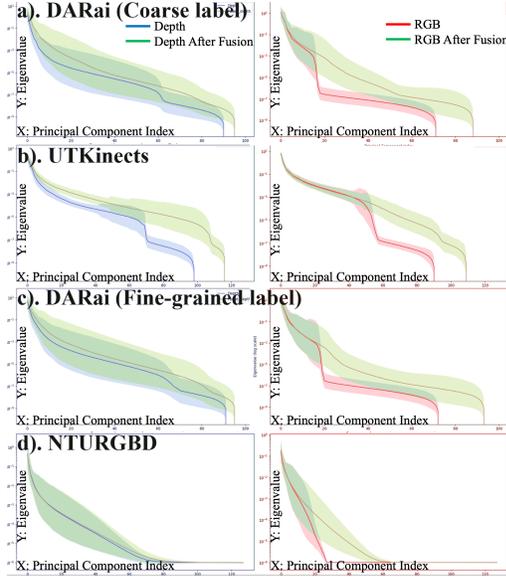


Figure 2. This figure compares the eigenvalue spectra of each modality before (Depth - blue, RGB - red) and after fusion (green) using the formulation in Theorem 3.1. The left column shows the spectrum for the Depth modality, and the right column for RGB. Across all datasets and label granularities, the fused modality consistently exhibits a flatter spectrum in mid-to-lower components as well as the dominant ones.

In short, Theorem 3.1 provides mathematical conditions under which selectively blending *less informative* channels in a modality  $X$  with channels from another modality  $Y$  that have bounded alignment with the dominant subspace of  $X$  leads to increased feature diversity, as quantified by effective rank.

Figure 2 shows that across all datasets, the eigenvalue spectra of both RGB and Depth representations become flatter after fusion. This flattening reflects a more uniform distribution of eigenvalues, which implies that the spectrum exhibits greater entropy. As in Equation 1, since effective rank is defined as the entropy over the spectrum, a flatter spectrum corresponds to higher effective rank, indicating a richer and more balanced representation. We provide detailed analysis of Figure 2 in Section B.1.

Notably, this increase in effective rank occurs in both modalities, as shown in Figure 2 (a-d). In other words, fusion not only enhances the effective rank of RGB through Depth modality, but also enhances Depth through RGB, confirming that the representational benefit is mutual. This mutual increase is desirable because it ensures that the fused representation integrates complementary information from both modalities, rather than favoring one. According to Theorem 3.1, such mutual improvement requires the injected features to be sufficiently complementary to the target modality, formalized by having a bounded projection onto its dominant subspace. This principle motivates the search for optimal modality combinations that maximize

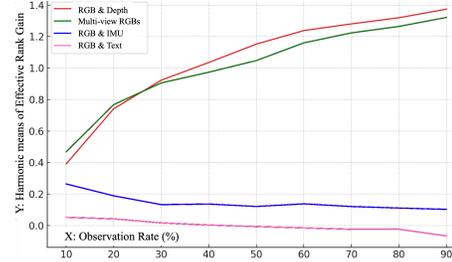


Figure 3. This figure compares the Harmonic Mean of effective rank gain across four modalities: Multi-view RGB, Text, IMU, and Depth. Depth consistently achieves the highest harmonic mean across all observation rates, indicating a more balanced interaction with RGB compared to other modalities.

balanced and mutual rank enhancement.

### 3.2. Modality Selection

In this section, we show that not all modality pairs are equally effective for preserving modality-specific features. To quantify this, we analyze how the effective rank of each modality changes after fusion with RGB across 4 different modalities: Depth, Multi-view RGB, IMU, and Text, each denoted as  $\Delta_{\text{Depth}}$ ,  $\Delta_{\text{MVRGB}}$ ,  $\Delta_{\text{IMU}}$ ,  $\Delta_{\text{Text}}$ , respectively. In this case, all modalities are used within a fusion framework for an action anticipation task. For analysis, we extract the fused representation of each modality at an intermediate point in the network. Then, we introduce the harmonic mean of effective rank gains as our evaluation metric. This is because this score captures both the magnitude and the symmetry of increases in effective rank across both modalities. A high harmonic mean implies that both modalities gain substantially from fusion and do so in a balanced manner. Figure 3 visualizes this harmonic mean score for four modality combinations:  $\frac{2 \cdot \Delta_{\text{RGB}} \cdot \Delta_{\text{Depth}}}{\Delta_{\text{RGB}} + \Delta_{\text{Depth}}}$ ,  $\frac{2 \cdot \Delta_{\text{RGB}} \cdot \Delta_{\text{MVRGB}}}{\Delta_{\text{RGB}} + \Delta_{\text{MVRGB}}}$ ,  $\frac{2 \cdot \Delta_{\text{RGB}} \cdot \Delta_{\text{IMU}}}{\Delta_{\text{RGB}} + \Delta_{\text{IMU}}}$ ,  $\frac{2 \cdot \Delta_{\text{RGB}} \cdot \Delta_{\text{Text}}}{\Delta_{\text{RGB}} + \Delta_{\text{Text}}}$ . As shown, Depth consistently achieves the highest harmonic mean across all observation rates. This suggests that fusion with Depth yields substantial improvements in overall representational capacity (magnitude) and distributes those improvements symmetrically across both RGB and Depth modalities. In contrast, other modalities tend to enhance one modality while degrading or leaving the other unchanged. Based on this analysis, we identify Depth as the most complementary modality to RGB, as it supports mutual enhancement without collapsing the expressive structure of either representation.

## 4. Methodology

This section details the components of R3D. Figure 4 provides a detailed illustration of the overall architecture of R3D. We begin with the RGB Encoder and Depth Encoder, which extracts meaningful features from the input data (RGB video and Depth video, respectively). The Rank-

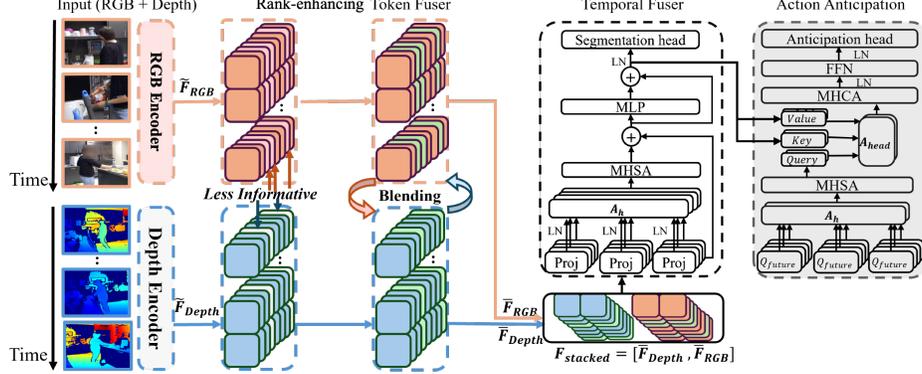


Figure 4. The detailed architecture of R3D. It comprises three components: the Rank-Enhancing Token Fuser (RTF), the Temporal Fuser, and the Action Anticipation Module. The RTF compensates for *less informative* channels in each modality by blending complementary information, while the Temporal Fuser captures continuous temporal dependencies and segments each timestamp. Finally, the Action Anticipation Module predicts future actions based on the integrated multi-modal information.

enhancing Token Fuser (RTF) adaptively blends *less informative* channels of both RGB and Depth modalities. Temporal Fuser integrates this multi-modal information and captures how multi-modal dependencies evolve over time for improved anticipation. Lastly, we describe the Action Anticipation Module.

#### 4.1. RGB and Depth Encoders

We have a dataset consisting of  $N$  paired videos consisting of RGB and Depth modalities. The RGB videos are denoted as  $\{X_n^{\text{RGB}}\}_{n=1}^N$ , and the corresponding Depth videos as  $\{X_n^{\text{Depth}}\}_{n=1}^N$ . Each RGB video  $X_n^{\text{RGB}} \in \mathbb{R}^{B \times T \times H \times W \times 3}$  and Depth video  $X_n^{\text{Depth}} \in \mathbb{R}^{B \times T \times H \times W \times 1}$  is encoded into visual features  $\mathbf{F}^{\text{RGB}} \in \mathbb{R}^{B \times T \times C}$  and  $\mathbf{F}^{\text{Depth}} \in \mathbb{R}^{B \times T \times C}$ , respectively, where  $B$  is the batch size,  $T$  the number of frames,  $H$  height,  $W$  width, and  $C$  the feature embedding dimension. To obtain this feature embeddings, we use a pretrained ResNet50 as the video encoder. To reduce computational cost while maintaining temporal structure, we sample frames at regular intervals using a temporal stride  $\tau$ . This results in a sampled sequences  $F_\tau^{\text{RGB}} \in \mathbb{R}^{B \times T_\tau \times C}$  and  $F_\tau^{\text{Depth}} \in \mathbb{R}^{B \times T_\tau \times C}$ , where  $T_\tau = \lfloor \frac{T}{\tau} \rfloor$  is the number of sampled frames. The features of the sampled frames are then passed through a linear transformation layer  $W^{\text{RGB}} \in \mathbb{R}^{C \times D}$  and  $W^{\text{Depth}} \in \mathbb{R}^{C \times D}$ , followed by a  $\text{ReLU}(\cdot)$  activation function:

$$\tilde{\mathbf{F}}^{\text{RGB}} = \text{ReLU}(F_\tau^{\text{RGB}} W^{\text{RGB}}); \tilde{\mathbf{F}}^{\text{Depth}} = \text{ReLU}(F_\tau^{\text{Depth}} W^{\text{Depth}}), \quad (2)$$

where  $\tilde{\mathbf{F}}^{\text{RGB}} \in \mathbb{R}^{B \times T_\tau \times D}$  and  $\tilde{\mathbf{F}}^{\text{Depth}} \in \mathbb{R}^{B \times T_\tau \times D}$  represent RGB input tokens and Depth input tokens, respectively, and  $D$  is a new feature dimension. The resulting features  $\tilde{\mathbf{F}}^{\text{RGB}}$  and  $\tilde{\mathbf{F}}^{\text{Depth}}$ , as shown in Figure 4, are then fed into RTF for cross-modal integration.

#### 4.2. Rank-enhancing Token Fuser (RTF)

In Section 3, we show that blending channels that are weakly aligned with the singular vectors of one modality with complementary channels from another can increase feature diversity. In this section, we translate this insight into a differentiable fusion module, RTF, as described in Figure 4.

**Channel importance estimation.** As illustrated in *Definition 1* of Theorem 3.1, we quantify the *importance* of each channel based on how much each channel contributes to the singular vector of the representation matrix. Given the SVD of the reshaped RGB feature matrix  $\tilde{\mathbf{F}}^{\text{RGB}} \in \mathbb{R}^{B \times T_\tau \times D}$  as  $\tilde{X}^{\text{RGB}} \in \mathbb{R}^{(B \cdot T_\tau) \times D}$ , we compute  $\tilde{X}^{\text{RGB}} = U^{\text{RGB}} \Sigma^{\text{RGB}} (V^{\text{RGB}})^T$ , where the singular values are  $\sigma_1^{\text{RGB}}, \dots, \sigma_r^{\text{RGB}}$  and right singular vectors are  $v_i^{\text{RGB}} \in \mathbb{R}^D$ . The *channel-wise importance score*  $I_c^{\text{RGB}}$  is then defined as  $I_c^{\text{RGB}} = \sum_{i=1}^r (\sigma_i^{\text{RGB}})^2 (v_{i,c}^{\text{RGB}})^2$ . Similarly, we apply SVD to the reshaped Depth feature matrix  $\tilde{\mathbf{F}}^{\text{Depth}} \in \mathbb{R}^{B \times T_\tau \times D}$ , resulting in the importance score denoted as  $I_c^{\text{Depth}} = \sum_{i=1}^r (\sigma_i^{\text{Depth}})^2 (v_{i,c}^{\text{Depth}})^2$ , where  $I_c^{\text{RGB}}, I_c^{\text{Depth}}$  reflects how much channel  $c$  contributes to the singular vector of  $\tilde{\mathbf{F}}^{\text{RGB}}, \tilde{\mathbf{F}}^{\text{Depth}}$ , respectively. As illustrated in Theorem 3.1, channels with low  $I_c$  are considered *less informative*, thus are blended with complementary features from the other modality. This targeted fusion ensures that *less informative* channels in one modality are enhanced by incorporating complementary cues from the other. Thus, to identify *less informative* channels, we select the bottom  $k'$  channels with the lowest importance scores:

$$\mathcal{I}_{\text{low}}^{\text{RGB}} = \underset{c}{\text{argmin}}_{k'} I_c^{\text{RGB}}, \mathcal{I}_{\text{low}}^{\text{Depth}} = \underset{c}{\text{argmin}}_{k'} I_c^{\text{Depth}}. \quad (3)$$

**Adaptive channel blending.** Once the least important channels are identified, we perform cross-modal feature blending only on the bottom  $k'$  least importance channels as determined by the channel importance estimation. Here, we introduce a learnable scaling parameter  $\alpha_{\mathcal{I}_{\text{low}}^{\text{RGB}}}$  and  $\alpha_{\mathcal{I}_{\text{low}}^{\text{Depth}}}$

that are randomly initialized from a uniform distribution, to enable adaptive channel blending:

$$\bar{\mathbf{F}}_{\dots, \mathcal{I}_{low}^{RGB}}^{RGB} = \alpha_{\mathcal{I}_{low}^{RGB}} \cdot y_{\dots, \mathcal{I}_{low}^{RGB}}^{Depth} + (1 - \alpha_{\mathcal{I}_{low}^{RGB}}) \cdot y_{\dots, \mathcal{I}_{low}^{RGB}}^{RGB}, \quad (4)$$

$$\bar{\mathbf{F}}_{\dots, \mathcal{I}_{low}^{Depth}}^{Depth} = \alpha_{\mathcal{I}_{low}^{Depth}} \cdot y_{\dots, \mathcal{I}_{low}^{Depth}}^{RGB} + (1 - \alpha_{\mathcal{I}_{low}^{Depth}}) \cdot y_{\dots, \mathcal{I}_{low}^{Depth}}^{Depth}, \quad (5)$$

where the updated features,  $\bar{\mathbf{F}}_{\dots, \mathcal{I}_{low}^{RGB}}^{RGB}$  and  $\bar{\mathbf{F}}_{\dots, \mathcal{I}_{low}^{Depth}}^{Depth}$  denote the refined representations after adaptive channel blending for the least informative channels  $\mathcal{I}_{low}^{RGB}, \mathcal{I}_{low}^{Depth}$ . These parameters  $\alpha_{\mathcal{I}_{low}^{RGB}}$  and  $\alpha_{\mathcal{I}_{low}^{Depth}}$  allow the model to adaptively control the degree of feature blending between RGB and Depth channels, rather than enforcing a hard swap.

These partially updated features are then combined with the remaining channels to form the final aggregated features:  $\bar{\mathbf{F}}_{low}^{RGB}$  and  $\bar{\mathbf{F}}_{low}^{Depth} \in \mathbb{R}^{B \times T_r \times D}$ . Finally, the fused features  $\bar{\mathbf{F}}_{low}^{RGB}$  and  $\bar{\mathbf{F}}_{low}^{Depth}$  are stacked together along the channel dimension ( $\mathbf{F}_{stacked}$  in Figure 4) and passed as an input to the Temporal Fuser for sequential modeling.

### 4.3. Temporal Fuser and Action Anticipation

As illustrated in Figure 4, Temporal Fuser takes the output from RTF and captures sequential dependencies over time, to accurately interpret the meaning of each frame. Its primary goal is to integrate temporal information so that the model can determine the meaning of a scene at every time step. Section C.1 details the three components of Temporal Fuser: MHSA, MLP, and Segmentation Head.

Building on these temporal representations, Action Anticipation Module predicts future human actions. Section C.2 describes its main components: MHCA, FFN, Action Anticipation Head.

## 5. Experiments

This section presents experimental analyses and discussion of the proposed method. The detailed dataset configuration is described in Section D.1, the experimental setup is detailed in Section D.2, and the computational cost and scalability are detailed in Section D.3. For all analyses, we set the observation rate  $\alpha \in \{0.2, 0.3\}$  and prediction rate  $\beta \in \{0.1, 0.2, 0.3, 0.5\}$ , and report mean over classes (MoC) accuracy as the evaluation metric. Qualitative analyses are provided in Section F.

### 5.1. Quantitative Analysis

**State-of-the-art comparison.** As we observe in Table 1, R3D consistently outperforms the state-of-the-art across all datasets and evaluation settings, achieving an average accuracy gain of up to 3.74%. We compare the performance of multiple architectures for action anticipation and observe that our approach achieves the best performance among all models, with AFFT [76], m&m-Ant [36], FUTR [22], and

Dataset	Input	Methods	$\beta(\alpha = 0.2)$				$\beta(\alpha = 0.3)$				
			0.1	0.2	0.3	0.5	0.1	0.2	0.3	0.5	
DARai [32] (Coarse)	Uni	RNN [1]	22.40	22.59	20.71	18.38	30.75	25.34	25.99	23.17	
		CNN [1]	13.28	13.70	12.99	13.90	19.85	16.68	18.69	17.48	
		FUTR [22]	25.05	25.11	24.48	23.18	40.71	33.57	33.43	30.79	
		GTAN [72]	27.70	29.13	27.35	26.11	42.47	<b>42.45</b>	42.27	34.37	
		R3D (Uni)	31.52	30.52	28.48	26.66	44.01	37.10	37.58	34.44	
DARai [32] (Fine-grained)	Multi	AFFT [76]	23.14	24.78	23.62	21.02	33.82	29.25	28.33	25.45	
		m&m-Ant [36]	25.75	25.70	24.24	23.12	42.00	34.71	34.49	31.34	
		<b>R3D (Ours)</b>	<b>33.44</b>	<b>32.14</b>	<b>30.56</b>	<b>29.59</b>	<b>46.29</b>	<b>42.05</b>	<b>43.41</b>	<b>40.25</b>	
		Uni	RNN [1]	7.45	5.69	4.66	3.58	9.07	7.35	6.40	5.04
UTKinects [65]	Multi	CNN [1]	7.70	6.28	5.38	3.98	6.69	5.48	5.10	3.86	
		FUTR [22]	17.09	12.34	10.39	8.18	18.59	15.83	13.26	10.81	
		GTAN [72]	25.66	24.06	22.60	<b>22.05</b>	26.30	25.43	22.40	<b>18.88</b>	
		R3D (Uni)	20.12	18.49	16.11	13.28	24.62	22.93	22.78	18.68	
		<b>R3D (Ours)</b>	<b>32.57</b>	<b>25.92</b>	<b>24.02</b>	<b>16.68</b>	<b>26.97</b>	<b>28.43</b>	<b>24.81</b>	<b>18.02</b>	
NTURGBD [59]	Multi	AFFT [76]	14.13	13.38	13.49	9.53	17.01	17.26	15.18	10.91	
		<b>R3D (Ours)</b>	<b>38.96</b>	<b>37.68</b>	<b>37.16</b>	<b>35.95</b>	<b>34.80</b>	<b>32.51</b>	<b>36.03</b>	<b>26.98</b>	
		Uni	RNN [1]	24.42	15.37	15.92	11.10	15.22	15.19	13.12	9.89
		CNN [1]	25.00	16.60	14.79	11.07	15.10	16.67	14.29	8.08	
		FUTR [22]	29.63	20.75	20.53	16.48	16.46	16.90	14.31	9.95	
NTURGBD [59]	Multi	GTAN [72]	34.38	35.71	<b>39.54</b>	34.17	28.27	32.36	27.22	21.16	
		R3D (Uni)	29.63	20.75	20.53	16.48	16.46	16.90	14.31	9.95	
		AFFT [76]	25.00	16.67	16.67	12.50	16.67	16.67	14.29	10.00	
		<b>R3D (Ours)</b>	<b>38.96</b>	<b>37.68</b>	<b>37.16</b>	<b>35.95</b>	<b>34.80</b>	<b>32.51</b>	<b>36.03</b>	<b>26.98</b>	
		Uni	RNN [1]	10.98	10.49	10.53	10.94	8.60	9.55	10.11	11.69
NTURGBD [59]	Multi	CNN [1]	15.11	14.91	14.33	15.12	15.09	15.49	15.52	16.74	
		FUTR [22]	18.59	18.56	18.61	18.61	20.04	20.13	20.07	20.11	
		GTAN [72]	18.93	18.51	18.55	18.70	21.63	<b>21.72</b>	<b>21.95</b>	20.93	
		R3D (Uni)	10.24	10.14	10.18	10.15	10.91	10.97	10.92	10.96	
		<b>R3D (Ours)</b>	<b>20.55</b>	<b>20.76</b>	<b>20.66</b>	<b>20.79</b>	<b>20.72</b>	<b>20.72</b>	<b>20.74</b>	<b>21.27</b>	
NTURGBD [59]	Multi	AFFT [76]	20.55	20.76	20.66	20.79	20.72	20.72	20.74	21.27	
		<b>R3D (Ours)</b>	<b>21.98</b>	20.18	19.72	18.90	<b>23.17</b>	21.27	20.51	19.89	

Table 1. State-of-the-art comparisons of action anticipation performance across three widely used datasets under varying conditions of  $\alpha$  (observation rate) and  $\beta$  (prediction rate). “Uni” denotes Uni-modal input and “Multi” indicates a Multi-modal setting. “R3D (Uni)” refers to the R3D model trained solely with RGB input, without depth modality.

GTAN [72], serving as the state-of-the-art baselines. This gain is particularly notable in scenarios with lower observation rates ( $\alpha = 0.2$ ), where depth information allows the model to leverage subtle cues even when visual input is limited. The performance boost is more pronounced on the DARai dataset than on the UTKinects dataset. This is because DARai contains more dense action sequences and more frequent transitions. In this case, depth information enhances the model’s ability to capture nuanced temporal variations, leading to greater improvements compared to other datasets. In contrast, on NTURGBD, the strength of R3D is less pronounced. This can be attributed to the fact that NTURGBD provides high-quality RGB-Depth pairs with well-trimmed and well-curated action segments. While R3D is designed to be robust in less curated, noisy real-world settings (See 5.4), its relative advantage diminishes when both modalities are already strong, which is less common in real-world settings.

### 5.2. Ablation study

We perform ablation studies to assess the impact of depth modality (Table 1), impact of Rank-enhancing Token Fuser (Table 2), Rank-enhancing Token Fuser modeling (Figure

Dataset	Methods	$\beta(\alpha = 0.2)$				$\beta(\alpha = 0.3)$			
		0.1	0.2	0.3	0.5	0.1	0.2	0.3	0.5
UTKinects	W/O RTF	29.63	20.75	20.53	16.48	16.46	16.90	14.31	9.95
	W/ RTF	<b>38.96</b>	<b>37.68</b>	<b>37.16</b>	<b>35.95</b>	<b>34.80</b>	<b>32.51</b>	<b>36.03</b>	<b>26.98</b>
NTURGBD	W/O RTF	16.86	16.74	16.74	16.74	19.53	19.89	19.74	19.89
	W/ RTF	<b>21.98</b>	<b>20.18</b>	<b>19.72</b>	<b>18.90</b>	<b>23.17</b>	<b>21.27</b>	<b>20.51</b>	<b>19.89</b>
DARai	W/O RTF	31.84	31.40	28.88	27.06	35.32	31.98	32.83	31.91
	W/ RTF	<b>33.44</b>	<b>32.14</b>	<b>30.56</b>	<b>29.59</b>	<b>46.29</b>	<b>42.05</b>	<b>43.41</b>	<b>40.25</b>

Table 2. Ablation Study evaluating the impact of RTF.

Dataset	Methods	$\beta(\alpha = 0.2)$				$\beta(\alpha = 0.3)$			
		0.1	0.2	0.3	0.5	0.1	0.2	0.3	0.5
UTKinects	Static	<b>40.30</b>	25.62	26.42	19.27	13.55	19.10	16.46	12.17
	Adaptive	38.96	<b>37.68</b>	<b>37.16</b>	<b>35.95</b>	<b>34.80</b>	<b>32.51</b>	<b>36.03</b>	<b>26.98</b>
NTURGBD	Static	18.93	17.93	16.90	16.77	18.48	17.83	16.65	16.30
	Adaptive	<b>21.98</b>	<b>20.18</b>	<b>19.72</b>	<b>18.90</b>	<b>23.17</b>	<b>21.27</b>	<b>20.51</b>	<b>19.89</b>
DARai	Static	27.86	26.50	24.62	23.39	41.10	35.35	37.00	35.01
	Adaptive	<b>33.44</b>	<b>32.14</b>	<b>30.56</b>	<b>29.59</b>	<b>46.29</b>	<b>42.05</b>	<b>43.41</b>	<b>40.25</b>

Table 3. Ablation Study on channel exchange method across three datasets. ‘‘Static’’ indicates Static Channel Exchange and ‘‘Adaptive’’ indicates Adaptive Channel Exchange with learnable scaling factor  $\alpha_1$ .

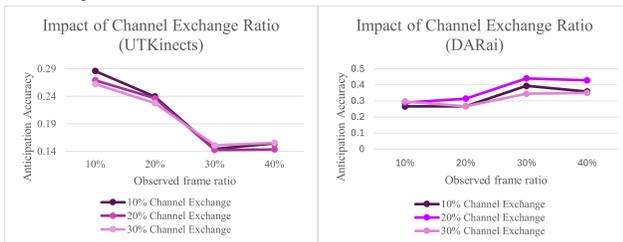


Figure 5. Ablation study on UTKinects and DARai dataset examining the impact of the proportion of exchanged channels (10%, 20%, 30%) in Token Fuser.

5, Table 3), and impact of Temporal Fuser (Table 8). We show the ablation study of Temporal Fuser and additional Rank-enhancing Token Fuser modeling in Section E.1 and Section E.2, respectively.

**RGB only.** First, we evaluate the effect of incorporating depth information by comparing models trained with and without depth features. As shown in Table 1, integrating depth consistently improves performance across all datasets, demonstrating its importance for capturing fine-grained spatial relationships and motion directionality.

**RGB-Depth aggregation without Rank-enhancement.** Next, we analyze the impact of the Token Fuser by removing it from R3D. The results in Table 2 show that models without Rank-enhancing Token Fuser exhibit lower performance, confirming its effectiveness in facilitating cross-modal information exchange and improving representation quality.

**Role of adaptive  $\alpha$  during fusion.** To further investigate Rank-enhancing Token Fuser modeling, we compare static channel exchange with adaptive channel exchange method as shown in Table 3. Across all datasets, adaptive blending yields better performance, as action anticipation task exhibits high variability in object interactions and spatial

Modality pair	Methods	$\beta(\alpha = 0.2)$				$\beta(\alpha = 0.3)$			
		0.1	0.2	0.3	0.5	0.1	0.2	0.3	0.5
RGB - Depth	AFFT [76]	14.13	13.38	13.49	9.53	17.01	17.26	15.18	10.91
	<b>R3D (Ours)</b>	<b>32.57</b>	<b>25.92</b>	<b>24.02</b>	<b>16.68</b>	<b>26.97</b>	<b>28.43</b>	<b>24.81</b>	<b>18.02</b>
Different-view RGBs	AFFT [76]	11.53	8.46	7.47	5.84	10.61	8.51	8.22	6.88
	<b>R3D (Ours)</b>	<b>26.99</b>	<b>23.80</b>	<b>20.73</b>	<b>21.50</b>	<b>35.02</b>	<b>27.37</b>	<b>27.19</b>	<b>22.62</b>
RGB - IMU	AFFT [76]	11.89	<b>10.04</b>	<b>8.81</b>	<b>7.01</b>	<b>11.31</b>	<b>9.72</b>	<b>9.64</b>	<b>7.59</b>
	<b>R3D (Ours)</b>	<b>12.08</b>	9.76	7.87	5.80	10.61	8.46	7.54	5.75
RGB - Text	AFFT [76]	20.21	21.83	21.93	17.07	25.79	25.62	24.90	19.62
	<b>R3D (Ours)</b>	<b>30.24</b>	<b>28.25</b>	<b>22.80</b>	<b>18.01</b>	<b>28.17</b>	<b>27.69</b>	<b>25.69</b>	<b>22.93</b>

Table 4. State-of-the-art comparisons of action anticipation performance on DARai dataset across four different modalities: Depth, Multi-view RGB, IMU, and Text.

configurations, requiring a flexible fusion mechanism.

**Impact of varying channel exchange ratio.** Additionally, we conduct an ablation study on the optimal channel exchange ratio for Token Fuser as shown in Figure 5. Our findings reveal that the optimal exchange ratio varies by dataset: UTKinect achieves the highest performance at 10% channel exchange, while DARai performs best at 20% exchange. This suggests that datasets with more structured, repetitive actions (e.g., UTKinect) benefit from lower exchange rates, whereas datasets with more complex, unstructured activities (e.g., DARai) require a higher degree of adaptive feature replacement.

### 5.3. More SOTA comparisons

**More modality pairs.** To further validate the general effectiveness of R3D with different modalities, as shown in Table 4, we additionally present the State-of-the-art comparisons of action anticipation performance across four different modality pairs: Depth, Multi-view RGB, IMU, and Text at varying observation rates ( $\alpha$ ) and prediction rates ( $\beta$ ). For comparison, we use AFPT [76], the current state-of-the-art method in multi-modal action anticipation. To ensure fairness and reproducibility, we conduct three runs with fixed random seeds (1, 10, 13452) and simply replace the input from the Figure 4. Overall, R3D achieves superior performance to AFPT for modality pairs with higher harmonic means as illustrated in Figure 3. For RGB - IMU pair, R3D underperforms compared to AFPT. We attribute this to the lack of sufficiently balanced interaction between RGB and IMU features, as also reflected in Figure 3. We expect that tailoring the fusion parameters for RGB - IMU would enhance the interaction and further improve performance.

**More downstream tasks.** To further validate the general effectiveness of RTF, we additionally apply our method to Action Segmentation task. Unlike Action Anticipation, Action Segmentation task segments an input video into temporally contiguous segments, assigning a categorical action label to each segment independently [9, 16, 23, 30, 45, 47, 64, 66, 68]. To comprehensively evaluate the task, first, we assess the correctness of predictions at every individual frame

Dataset	Methods	Frame-wise metrics	Segment-wise metrics
		Accuracy	Edit score
DARai [32] (Coarse)	ActFusion [23]	31.38	26.62
	<b>R3D (Ours)</b>	<b>33.32</b>	<b>31.21</b>
DARai [32] (Fine-grained)	ActFusion [23]	17.83	18.16
	<b>R3D (Ours)</b>	<b>20.79</b>	<b>28.98</b>
NTURGBD [59]	ActFusion [23]	13.26	6.00
	<b>R3D (Ours)</b>	<b>13.26</b>	<b>8.93</b>

Table 5. State-of-the-art comparisons of action segmentation performance across multiple datasets.

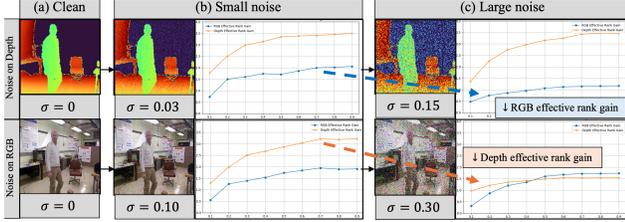


Figure 6. Effective rank gain under various noise perturbations (Depth: (b) small  $\sigma=0.03$ , (c) large  $\sigma=0.15$ , (b) RGB: small  $\sigma=0.10$ , (c) large  $\sigma=0.30$ ). When more noise is added to one modality, the noisy modality’s effective rank gain remains stable, whereas the other (clean) modality’s gain diminishes. For example, increasing noise in Depth keeps its gain stable but reduces RGB’s. This shows that fusion adaptively relies more on the cleaner modality.

by reporting frame-wise accuracy [9, 23, 30, 45, 64, 66]. Additionally, we also measure how well the predicted segments align with the true temporal boundaries and labels of actions by measuring an edit score, a segment-wise metrics [16, 23, 47, 68]. To obtain results from the action segmentation task, we use the output of the Segmentation Head of the Temporal Fuser of R3D, illustrated in Figure 4 and explained in detail in Section C.1. To ensure consistency, we report the average performance over three runs with fixed random seeds (1, 10, 13452).

As shown in Table 5, R3D consistently outperforms the state-of-the-art across different levels of granularity. The performance gap is more pronounced on the fine-grained setting, since the finer temporal boundaries and larger number of action segments make boundary alignment substantially more challenging, amplifying the advantage of R3D in capturing temporal structure and maintaining segment consistency. It is worth noting that R3D is not originally designed for action segmentation task, yet it still demonstrates strong performance on this task.

## 5.4. Robustness

In this section, we simulate the real world scenario where one modality can be unreliable or noisy. Figure 6 illustrates two cases: noise added to the Depth modality (top) and to the RGB (bottom). We assess how much each modality contributes to fusion by measuring the effective rank gain across three conditions: clean (Figure 6 (a)), small noise

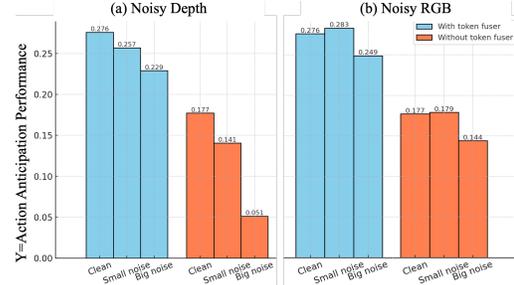


Figure 7. This shows an action anticipation performance under varying noise levels (clean, small noise, large noise) in two settings: with Rank-Enhancing Token Fuser (RTF) and without RTF. Notably, (a) when noise is added to Depth, performance without RTF degrades substantially, while with RTF maintains stability, compared to (b) when noise is added to RGB.

(Figure 6 (b)), and large noise (Figure 6 (c)). To ensure fairness, we calibrate the noise level  $\sigma$  separately for Depth and RGB so that perturbations induce the comparable change in each modality’s representation space.

Figure 7 (a) shows that without RTF, the performance drops sharply when Depth is noisy, whereas with RTF, the performance remains stable under both noisy Depth and noisy RGB. Figure 6 shows the mechanism of how RTF works. As noise increases (Figure 6 (b)  $\rightarrow$  (c)) in one modality, RTF adaptively shifts reliance toward the cleaner modality. For example, when Depth is corrupted, Depth’s effective rank gain remains stable whereas RGB’s gain diminishes (Figure 6 (b)  $\rightarrow$  (c), Top part), indicating that the fusion compensates by down-weighting the unreliable Depth contribution and leveraging the complementary information from RGB, so that the overall representation remains robust despite the corruption.

## 6. Conclusion

We present R3D, a depth-informed framework for action anticipation that addresses two representation collapse in multi-modal fusion. The utility of effective rank as a computational means to counter both feature and modality collapse jointly is emphasized. We introduce RTF, a theoretically grounded mechanism that selectively blends *less informative* channels with complementary signals from another modality, provably increasing representational diversity. The principle of mutual effective rank gain, whereby the effective rank of both modalities increases simultaneously is highlighted. For the application of action anticipation, depth adds background context and directionality to the object-centric visual cues from RGB. R3D is validated extensively across NTURGBD, UTKinect, and DARai. Our method not only outperforms state-of-the-art approaches but also provides new insights into what constitutes effective and balanced multi-modal fusion.

## References

- [1] Yazan Abu Farha, Alexander Richard, and Juergen Gall. When will you do what?-anticipating temporal occurrences of activities. In *CVPR*, pages 5343–5352, 2018. 2, 6, 16
- [2] Yazan Abu Farha, Qihong Ke, Bernt Schiele, and Juergen Gall. Long-term anticipation of activities with cycle consistency. In *PR*, pages 159–173. Springer, 2021. 2, 16
- [3] Kumar K Agrawal, Arnab Kumar Mondal, Arna Ghosh, and Blake Richards. alpha-req: Assessing representation quality in self-supervised learning by measuring eigenspectrum decay. *Advances in Neural Information Processing Systems*, 35:17626–17638, 2022. 3
- [4] Cem Akkus, Luyang Chu, Vladana Djakovic, Steffen Jauch-Walser, Philipp Koch, Giacomo Loss, Christopher Marquardt, Marco Moldovan, Nadja Sauter, Maximilian Schneider, et al. Multimodal deep learning. *arXiv preprint arXiv:2301.04856*, 2023. 2
- [5] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2018. 2
- [6] Adrien Bardes, Jean Ponce, and Yann LeCun. Vi-creg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021. 3
- [7] Emad Barsoum, John Kender, and Zicheng Liu. 3d human motion anticipation and classification. *arXiv preprint arXiv:2012.15378*, 2020. 1, 2
- [8] Apoorva Beedu, Harish Haresamudram, Karan Samel, and Irfan Essa. On the efficacy of text-based input modalities for action anticipation. *arXiv preprint arXiv:2401.12972*, 2024. 2
- [9] Nadine Behrmann, S Alireza Golestaneh, Zico Kolter, Juergen Gall, and Mehdi Noroozi. Unified fully and timestamp supervised temporal action segmentation via sequence to sequence translation. In *European conference on computer vision*, pages 52–68. Springer, 2022. 7, 8
- [10] Ryan Benkert, Mohit Prabhushankar, and Ghassan AlRegib. Transitional uncertainty with layered intermediate predictions. *arXiv preprint arXiv:2405.17494*, 2024. 2
- [11] Wenming Cao, Shuangshuang Li, and Jianqi Zhong. Qmednet: A quaternion-based multi-order differential encoder-decoder model for 3d human motion prediction. *Neural Networks*, 154:141–151, 2022. 1, 2
- [12] Abhra Chaudhuri, Anjan Dutta, Tu Bui, and Serban Georgescu. A closer look at multimodal representation collapse. *arXiv preprint arXiv:2505.22483*, 2025. 1, 2, 3
- [13] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198, 2024. 2
- [14] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *IJCV*, pages 1–23, 2022. 2
- [15] Haodong Duan, Yue Zhao, Kai Chen, Dahua Lin, and Bo Dai. Revisiting skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2969–2978, 2022. 2
- [16] Yazan Abu Farha and Juergen Gall. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In *CVPR*, pages 3575–3584, 2019. 7, 8
- [17] Basura Fernando and Samitha Herath. Anticipating human actions by correlating past with the future with jaccard similarity measures. In *CVPR*, pages 13224–13233, 2021. 2
- [18] Yanping Fu, Qingan Yan, Jie Liao, Huajian Zhou, Jin Tang, and Chunxia Xiao. Seamless texture optimization for rgb-d reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 29(3):1845–1859, 2021. 1
- [19] Antonino Furnari and Giovanni Maria Farinella. What would you expect? anticipating egocentric actions with rolling-unrolling lstms and modality attention. In *ICCV*, pages 6252–6261, 2019. 2
- [20] Quentin Garrido, Randall Balestriero, Laurent Najman, and Yann Lecun. Rankme: Assessing the downstream performance of pretrained self-supervised representations by their rank. In *International conference on machine learning*, pages 10929–10974. PMLR, 2023. 3
- [21] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15180–15190, 2023. 1
- [22] Dayoung Gong, Joonseok Lee, Manjin Kim, Seong Jong Ha, and Minsu Cho. Future transformer for long-term action anticipation. In *CVPR*, pages 3052–3061, 2022. 6, 15
- [23] Dayoung Gong, Suha Kwak, and Minsu Cho. Actfusion: a unified diffusion model for action segmentation and anticipation. *arXiv:2412.04353*, 2024. 2, 7, 8
- [24] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *CVPR*, pages 18995–19012, 2022. 2
- [25] Shannan Guan, Haiyan Lu, Linchao Zhu, and Gengfa Fang. Afe-cnn: 3d skeleton-based action recognition with action feature enhancement. *Neurocomputing*, 514:256–267, 2022. 1, 2
- [26] Caner Hazirbas, Lingni Ma, Csaba Domokos, and Daniel Cremers. Fusetnet: Incorporating depth into semantic segmentation via fusion-based cnn architecture. In *Asian conference on computer vision*, pages 213–228. Springer, 2016. 2
- [27] Ronghang Hu and Amanpreet Singh. Unit: Multimodal multitask learning with a unified transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1439–1449, 2021. 2
- [28] Jia-Hong Huang. Multi-modal video summarization. In *Proceedings of the 2024 International Conference on Multimedia Retrieval*, pages 1214–1218, 2024. 1, 2

- [29] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020. 2
- [30] Yuchi Ishikawa, Seito Kasai, Yoshimitsu Aoki, and Hirokatsu Kataoka. Alleviating over-segmentation errors by detecting action boundaries. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2322–2331, 2021. 7, 8
- [31] Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. Understanding dimensional collapse in contrastive self-supervised learning. *arXiv preprint arXiv:2110.09348*, 2021. 3
- [32] Ghazal Kaviani, Yavuz Yarici, Seulgi Kim, Mohit Prabhushankar, Ghassan AlRegib, Mashhour Solh, and Ameya Patil. Hierarchical and multimodal data for daily activity understanding. *arXiv preprint arXiv:2504.17696*, 2025. 1, 6, 8, 15
- [33] Ghazal Kaviani, Yavuz Yarici, Seulgi Kim, Mohit Prabhushankar, and Ghassan AlRegib. Exploring human daily life through a hierarchical multimodal lens, 2025. 1
- [34] Qihong Ke, Mario Fritz, and Bernt Schiele. Time-conditioned action anticipation in one shot. In *CVPR*, pages 9925–9934, 2019. 16
- [35] Jaeil Kim, Suhyun Kang, Duhun Hwang, Jungwook Shin, and Wonjong Rhee. Vne: An effective method for improving deep representation by manipulating eigenvalue distribution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3799–3810, 2023. 3
- [36] Seulgi Kim, Ghazal Kaviani, Mohit Prabhushankar, and Ghassan AlRegib. Multi-level and multi-modal action anticipation. *arXiv preprint arXiv:2506.02382*, 2025. 1, 2, 6
- [37] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International conference on machine learning*, pages 5583–5594. PMLR, 2021. 2
- [38] Kiran Kokilepersaud, Yash-Yee Logan, Ryan Benkert, Chen Zhou, Mohit Prabhushankar, Ghassan AlRegib, Enrique Corona, Kunjan Singh, and Mostafa Parchami. Focal: A cost-aware video dataset for active learning. In *2023 IEEE International Conference on Big Data (BigData)*, pages 1269–1278. IEEE, 2023. 2
- [39] Kiran Kokilepersaud, Seulgi Kim, Mohit Prabhushankar, and Ghassan AlRegib. Hex: Hierarchical emergence exploitation in self-supervised algorithms. *arXiv preprint arXiv:2410.23200*, 2024. 3
- [40] Kiran Kokilepersaud, Mohit Prabhushankar, and Ghassan AlRegib. Adadim: Dimensionality adaptation for ssl representational dynamics. *arXiv preprint arXiv:2505.12576*, 2025. 2, 3
- [41] Hema Koppula and Ashutosh Saxena. Learning spatio-temporal structure from rgb-d videos for human activity detection and anticipation. In *Proceedings of the 30th International Conference on Machine Learning*, pages 792–800, Atlanta, Georgia, USA, 2013. PMLR. 1, 2
- [42] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, 2014. 2
- [43] Bolin Lai, Sam Toyer, Tushar Nagarajan, Rohit Girdhar, Shengxin Zha, James M Rehg, Kris Kitani, Kristen Grauman, Ruta Desai, and Miao Liu. Human action anticipation: A survey. *arXiv preprint arXiv:2410.14045*, 2024. 2
- [44] Thomas Laurent, James H von Brecht, and Xavier Bresson. Feature collapse. *arXiv preprint arXiv:2305.16162*, 2023. 1
- [45] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 156–165, 2017. 7, 8
- [46] Yin Li, Miao Liu, and James M Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *ECCV*, pages 619–635, 2018. 2
- [47] Daochang Liu, Qiyue Li, Anh-Dung Dinh, Tingting Jiang, Mubarak Shah, and Chang Xu. Diffusion action segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10139–10149, 2023. 7, 8
- [48] Feng Liu, Ziwang Fu, Yunlong Wang, and Qijian Zheng. Tacfn: transformer-based adaptive cross-modal fusion network for multimodal emotion recognition. *arXiv preprint arXiv:2505.06536*, 2025. 1, 2
- [49] Ilya Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 16
- [50] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 16
- [51] Albert W Marshall, Ingram Olkin, and Barry C Arnold. Inequalities: theory of majorization and its applications. *Spring Series in Statistics*, 1979. 14
- [52] Sean O’Rourke, Van Vu, and Ke Wang. Matrices with gaussian noise: Optimal estimates for singular subspace perturbation. *IEEE Transactions on Information Theory*, 70(3): 1978–2002, 2023. 12
- [53] Mohit Prabhushankar, Kiran Kokilepersaud, Yash-yee Logan, Stephanie Trejo Corona, Ghassan AlRegib, and Charles Wykoff. Olives dataset: Ophthalmic labels for investigating visual eye semantics. *Advances in Neural Information Processing Systems*, 35:9201–9216, 2022. 2
- [54] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1
- [55] Alfréd Rényi. On measures of entropy and information. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability, volume 1: contributions to the theory of statistics*, pages 547–562. University of California Press, 1961. 2, 3
- [56] Tiziana Rotondo, Giovanni Maria Farinella, Valeria Tomaselli, and Sebastiano Battiato. Action anticipation from multimodal data. In *VISIGRAPP (4: VISAPP)*, pages 154–161, 2019. 2
- [57] Olivier Roy and Martin Vetterli. The effective rank: A measure of effective dimensionality. In *2007 15th European signal processing conference*, pages 606–610. IEEE, 2007. 2, 3

- [58] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *ECCV*. Springer, 2020. 2, 16
- [59] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019, 2016. 6, 8, 15
- [60] Vimal Thilak, Chen Huang, Omid Saremi, Laurent Dinh, Hanlin Goh, Preetum Nakkiran, Joshua M Susskind, and Etai Littwin. Lidar: Sensing linear probing performance in joint embedding ssl architectures. *arXiv preprint arXiv:2312.04000*, 2023. 3
- [61] Abhinav Valada, Rohit Mohan, and Wolfram Burgard. Self-supervised model adaptation for multimodal semantic segmentation. *International Journal of Computer Vision*, 128(5):1239–1285, 2020. 2
- [62] John Von Neumann. *Mathematical foundations of quantum mechanics: New edition*. Princeton university press, 2018. 2, 3
- [63] Jinghua Wang, Zhenhua Wang, Dacheng Tao, Simon See, and Gang Wang. Learning common and specific features for rgb-d semantic segmentation with deconvolutional networks. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 664–679. Springer, 2016. 2
- [64] Zhenzhi Wang, Ziteng Gao, Limin Wang, Zhifeng Li, and Gangshan Wu. Boundary-aware cascade networks for temporal action segmentation. In *European Conference on Computer Vision*, pages 34–51. Springer, 2020. 7, 8
- [65] L. Xia, C.C. Chen, and JK Aggarwal. View invariant human action recognition using histograms of 3d joints. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2012 IEEE Computer Society Conference on*, pages 20–27. IEEE, 2012. 6, 15
- [66] Ziwei Xu, Yogesh Rawat, Yongkang Wong, Mohan S Kankanhalli, and Mubarak Shah. Don’t pour cereal into coffee: Differentiable temporal logic for temporal action segmentation. *Advances in Neural Information Processing Systems*, 35:14890–14903, 2022. 7, 8
- [67] Yavuz Yarici, Kiran Kokilepersaud, Mohit Prabhushankar, and Ghassan AlRegib. Subject invariant contrastive learning for human activity recognition. In *2025 IEEE 35th International Workshop on Machine Learning for Signal Processing (MLSP)*, Istanbul, Turkey, 2025. IEEE. 1
- [68] Fangqiu Yi, Hongyu Wen, and Tingting Jiang. As-former: Transformer for action segmentation. *arXiv preprint arXiv:2110.08568*, 2021. 7, 8
- [69] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6281–6290, 2019. 1
- [70] Yuan Yuan, Zhaojian Li, and Bin Zhao. A survey of multimodal learning: Methods, applications, and future. *ACM Computing Surveys*, 57(7):1–34, 2025. 1
- [71] Olga Zatsarynna, Yazan Abu Farha, and Juergen Gall. Multi-modal temporal convolutional network for anticipating actions in egocentric videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2249–2258, 2021. 2
- [72] Olga Zatsarynna, Emad Bahrami, Yazan Abu Farha, Gianpiero Francesca, and Juergen Gall. Gated temporal diffusion for stochastic long-term dense anticipation. In *European Conference on Computer Vision*, pages 454–472. Springer, 2024. 6, 16
- [73] Jin Zeng, Yanfeng Tong, Yunmu Huang, Qiong Yan, Wenxiu Sun, Jing Chen, and Yongtian Wang. Deep surface normal estimation with hierarchical rgb-d fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6153–6162, 2019. 2
- [74] Yifan Zhang, Zhiquan Tan, Jingqin Yang, Weiran Huang, and Yang Yuan. Matrix information theory for self-supervised learning. *arXiv preprint arXiv:2305.17326*, 2023. 2, 3
- [75] Yuqi Zhang, Xiucheng Li, Hao Xie, Weijun Zhuang, Shihui Guo, and Zhijun Li. Multi-label action anticipation for real-world videos with scene understanding. *IEEE Transactions on Image Processing*, 2024. 2
- [76] Zeyun Zhong, David Schneider, Michael Voit, Rainer Stiefelhagen, and Jürgen Beyerer. Anticipative feature fusion transformer for multi-modal action anticipation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6068–6077, 2023. 1, 2, 6, 7