Evolving Space-Time Neural Architectures for Videos

AJ Piergiovanni, Anelia Angelova, Alexander Toshev, Michael S. Ryoo
Google Brain
{ajpiergi,anelia,toshev,mryoo}@google.com

Abstract

In this paper, we present a new method for evolving video CNN models to find architectures that more optimally captures rich spatio-temporal information in videos. Previous work, taking advantage of 3D convolutional layers, obtained promising results by manually designing CNN architectures for videos. We here develop an evolutionary algorithm that automatically explores models with different types and combinations of space-time convolutional layers to jointly capture various spatial and temporal aspects of video representations. We further propose a new key component in video model evolution, the iTGM layer, which more efficiently utilizes its parameters to allow learning of space-time interactions over longer time horizons. The experiments confirm the advantages of our video CNN architecture evolution, with results outperforming previous state-of-the-art models. Our algorithm discovers new and interesting video architecture structures.

1. Introduction

Video understanding tasks, such as video object detection and activity recognition, are important for many societal applications of computer vision including robot perception, smart cities, medical analysis, and more. Convolutional neural networks (CNNs) have been popular for video understanding. More specifically, representing a video input as a 3D XYT tensor by concatenating frames along time axis T, video CNNs learn intermediate representations as 4D XYTC tensors (where C is the feature dimension) using 3D conv. layers. Previous approaches including C3D [29], I3D [1], R(2+1)D [32], S3D [37], and others [3, 9] have been very successful particularly in video classification. These approaches focused on manually designing CNN architectures specialized for videos, for example by extending known 2D architectures such as Inception [27] and ResNet [5] to 3D [1, 32]. However, designing new, larger or more advanced architectures is a challenging problem, especially when the complexity of the tasks (e.g., continuous activity recognition) requires deeper and

wider architectures with more complex layers, as the exact configuration of layers is not clear. Furthermore, the available known network structures may not necessarily capture joint space-time information in complex videos sufficiently, which is needed to properly represent the fine-grained details that video understanding entails.

We here propose a new method for *evolving* video CNN architectures to harness the rich *spatio-temporal* information present in videos. We develop an evolutionary algorithm that automatically explores models with different types and combinations of space-time convolutional layers. The mutations in our evolutionary algorithm are particularly designed to search for CNN architectures that jointly capture various spatial and temporal aspects of video representations. This requires consideration of more diverse space-time layer components, and searching for architectures with heterogeneous modules with different roles. The proposed algorithm results in novel architectures which present interesting sub-structures, not known before (Figure 1). Furthermore, our approach encourages exploration of more *diverse* architectures by applying nontrivial mutations at the earlier stages of evolution while constraining the mutations at the later stages. This enables discovering multiple very different (but similarly good) architectures, allowing us to form a better ensemble model by combining them.

Neural architecture search and evolution have been applied successfully previously, but were limited to text and 2D image classifications [28, 40]. A naive extension of architecture search to video understanding will encounter multiple challenges: 1) the additional time dimension in video representations makes the architecture search space prohibitively large, even with contemporary compute capabilities – smarter units and combinations are needed; 2) training of each video model itself is computationally expensive, and architecture search requires training thousands more efficiently – layers which utilize more efficiently parameters to allow faster training are needed. Our approach addresses both these issues. To our knowledge this is the first automated neural architecture search algorithm for video understanding.

To enrich the search space for video inputs, we introduce a new key element which is specifically designed to capture
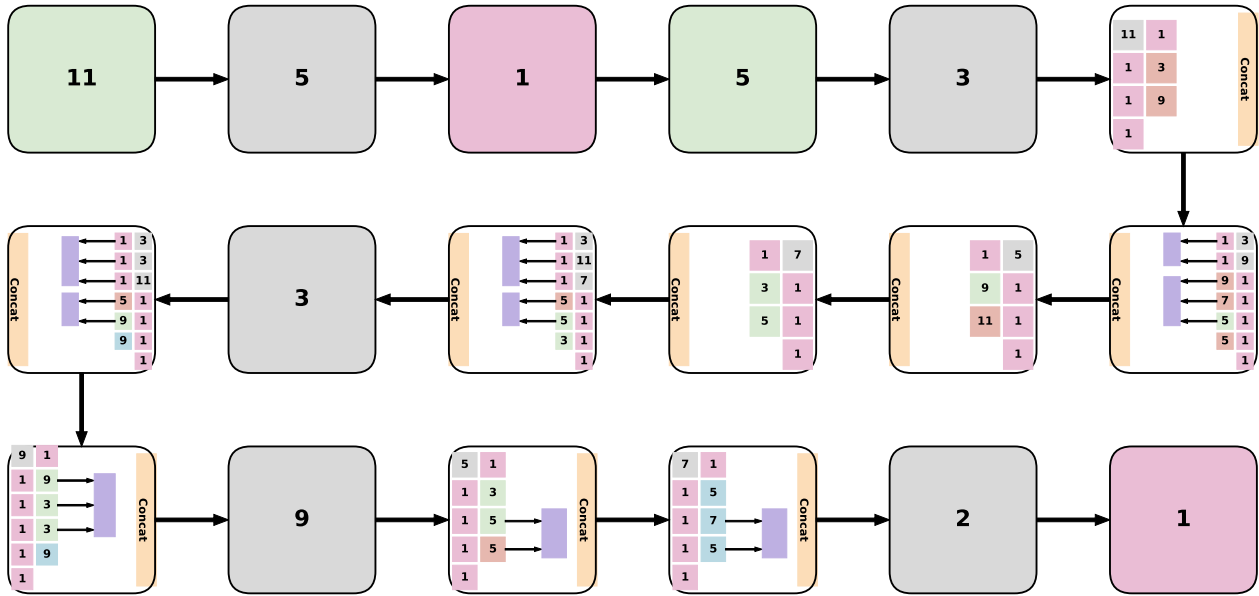
Figure 1. An example new video CNN architecture obtained with evolution (using Charades dataset). Several interesting substructures were discovered: modules combining multiple space-time pooling layers with different temporal intervals (middle row) or modules combining multiple new Inflated TGM layers with (2+1)D or 3D layers (bottom row). Colored squares are different types of convolutional layers, while grey squares are space-time pooling layers. The numbers in the squares indicate the temporal length of the filters. More details can be found in the "architecture findings" paragraph of Section 5.4.

space-time features interactions. We introduce an *Inflated TGM* (iTGM) layer as a possible component our video architecture evolution considers and evolves. The iTGM layer is a layer motivated by the original 1D TGM (Temporal Gaussian Mixture) layer [20]. In our inflated TGM layer, we learn 2D XY convolutional filters in addition to the temporal Gaussian mixture values, and inflate the 2D filter temporally to allow learning of joint features in 3D XYT (Figure 2). The 2D filter is inflated non-uniformly, by following the weights according to the learned 1-D temporal Gaussian mixture pattern. This allows to explore interaction between space and time more effectively and with much fewer number of parameters, and at the same time capture longer temporal information in videos.

Our approach outperforms previous state-of-the-art numbers on all three public action recognition datasets we tested (i.e., HMDB, Charades and Kinetcics), establishing a new state-of-the-art. We also discover complex substructures from the evolved CNNs including modules with multiple parallel space-time conv/pooling layers focusing on different temporal resolutions of video representations. Finally, unlike architecture search, independent evolution attempts generate a diverse set and ensembling them (which is for free as they are already available) increases recognition accuracy beyond what other homogeneous-architecture ensembles can do. The main technical contributions of this paper are:

- We propose an evolutionary approach for developing convolutional modules, specifically designed for targeting motion understanding in videos. To our knowledge this is the first paper to consider neural architecture search or evolution for video understanding.

- We introduce a new space-time convolutional layer, the Inflated TGM layer, designed to capture longer-term temporal information, and design the search space to specifically explore space-time convolutional layers and their combinations.

- New state-of-the-art on several video datasets and new architectures and components which can be reused for future work. Ensembling the discovered diverse architectures further increases the recognition accuracy.

2. Related work

CNNs for video understanding. Approaches considering a video as a space-time volume have been particularly successful [1, 4, 29, 30], with a direct application of 3D CNNs to videos. C3D [29] learned 3x3x3 XYT filters, which was not only applied for action recognition but also for video object recognition. I3D [1] extended the Inception architecture to 3D, obtaining successful results on multiple activity recognition video datasets including Kinetics. S3D [37] investigated the usage of 1D and 2D convolutional layers in addition to the 3D layers. R(2+1)D [32] used the

2D conv. layers followed by 1D conv. layers while following the ResNet structure. Two-stream CNN design is also widely adopted in action recognition, which takes optical flow inputs in addition to raw RGBs [3, 26]. There also are works focusing on capturing longer temporal information in continuous videos using pooling [18], attention [19], and convolution [9]. Recurrent neural networks (e.g., LSTMs) were also used to sequentially represent videos [18, 38].

Neural architecture search. Neural network architectures have advanced significantly since the early convolutional neural network concepts of LeCun et al. [13] and Krizhevsky et al. [11]: from developing wider modules, e.g., Inception [27], or introducing duplicated modules [14], residual connections [5, 36], densely connected networks [6, 7], or multi-task architectures: e.g., Faster-RCNN and RetinaNet for detection, and many others [15, 16, 23]. Recently several ground-breaking approaches have been proposed for automated learning/searching of neural network architectures, rather than manually designing them [22, 28, 40, 41]. Successful architecture search has been demonstrated for images and text [40, 41], including object classification. Tran et al. [31] analyze action recognition experiments with different settings, e.g., input resolution, frame rate, number of frames, network depth, all within the 3D ResNet architecture. In the context of video understanding, we are not aware of any prior work that has attempted developing an automated algorithm for data-driven architecture search/evolution.

3. Convolutional layers for action recognition

In this section, we review standard space-time convolutional layers and introduce a new efficient 3D space-time convolutional layer named the *Inflated Temporal Gaussian Mixture* (iTGM) layer. These convolutional layers take 4D XYTC space-time tensors both as an input and output, and they serve as basic components of the CNN architectures.

Standard 3D convolutional layer. The 3D convolutional layer [8] learns a standard 3D XYT convolutional kernel. It takes a $T \times Y \times X \times C_{in}$ input and learns C_{out} number of filters to be convolved in 3D (XYT) directions. Each filter size is $L \times H \times W \times C_{in}$, where C_{in} is the number of input channels, W are H are spatial width/height of the kernel and L is the temporal length of the kernel. Further, expanding 2D kernels to 3D has been explored [17]. I3D expanded kernels by stacking the 2D kernels L times, results in state-of-the-art performance [1]. The 3D convolutional layer has many parameters (i.e., $L \times H \times W \times C_{in} \times C_{out}$), which makes it challenging to learn, especially for small datasets.

(2+1)D convolutional layer. Recently (2+1)D layers have been proposed to decompose a 3D kernel into a 2D spatial kernel followed by a 1D temporal kernel [32, 37]. It is implemented as a 2D spatial convolution, followed by a 1D
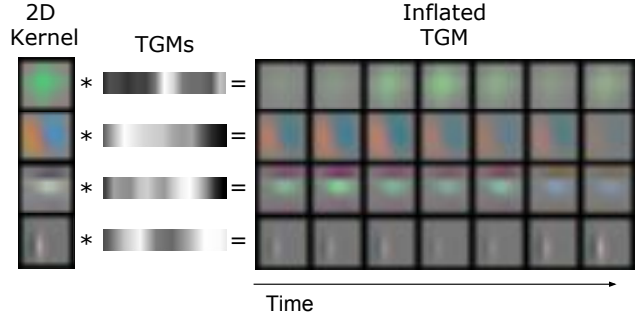

Figure 2. Example of inflated TGM kernels.

convolution in time, thus learning fewer parameters than the standard 3D convolution. It has $H \times W \times C_{in} \times C_{out} + L \times C_{out} \times C_{out}$ parameters.

3D Inflated TGM layer. Following a recent idea on Temporal Gaussian Mixture (TGM) layers [20], we propose to 'inflate' 2D image kernels into 3D kernels based on a mixture of Gaussians. Rather than directly learning the $L \times H \times W \times C_{in} \times C_{out}$ parameters used in the standard 3D convolution, we learn a 2D kernel with a size of $H \times W \times C_{in} \times C_{out}$ together with 1-D temporal Gaussian mixture parameters whose size is $2 \times M \times C_{in} + M \times C_{out}$. We then temporally inflate this 2D kernel according to the Gaussian mixture. M Gaussians are parameterized by a center μ and width σ, and a set of soft-attention mixing-weights, $a_{i,m}$, where $i \in [0, C_{out}]$, are learned (M is a hyper-parameter, typically smaller than L). In practice, μ is constrained to be between 0 and L and σ is positive:

$$\mu = (L-1) \cdot \frac{\tanh(\hat{\mu}) + 1}{2}, \;\; \sigma^2 = \exp(\hat{\sigma}). \quad (1)$$

The temporal Gaussians are constructed as:

$$\hat{K}_{m,l} = \frac{1}{Z} \exp -\frac{(l - \mu_m)^2}{2\sigma_m^2} \quad (2)$$

where Z is a normalization constant such that $\sum_{l=0}^{L} \hat{K}_{m,l} = 1$. We then apply the mixing weights,

$$K_{i,l} = \frac{\exp a_{i,m}}{\sum_j \exp a_{i,j}} \hat{K}_{m,l}. \quad (3)$$

This results in K being a $C_{out} \times L$ kernel. Given S, a 2D spatial kernel, we inflate S to a 3D as:

$$k = S * K \quad (4)$$

by convolving S with K. This results in a $L \times H \times W \times C_{in} \times C_{out}$ kernel. The parameters of the inflated TGM (i.e., S, μ_m, σ_m, and $a_{i,m}$) are differentiable; they are learned from the training data and become optimized for the task. Once learned, it behaves exactly like the standard 3D XYT convolution. Note that this layer learns fewer parameters than both 3D and (2+1)D convolution, and can learn longer kernels as the number of parameters is independent of the length, L. Examples of inflated TGMs are shown in Fig. 2.

4. Neural architecture evolution for videos

The objective of neural architecture evolution is to find better-performing CNN architectures by iteratively modifying a pool of architectures. It starts from a set of random, diverse architectures and continues to mutate them over multiple rounds, while only letting the better performing models survive. Recent studies [21] show that evolutionary algorithms can find good object classification architectures from a smaller number of samples, as opposed to model search algorithms using reinforcement learning [40]. This makes evolution more suitable for video CNN architecture search, as training video CNNs is expensive. Further, it allows for mutating architectures by selecting and combining various space-time layers which more effectively process representations with much larger dimensionality. The evolution also allows us to obtain multiple different (i.e., diverse) architectures instead of a single architecture, enabling the construction of a powerful ensemble.

4.1. Search space and base architecture

We evolve our architecture to have heterogeneous modules (or layers), as opposed to repeating one identical module multiple times. This is in contrast to many prior architecture search works on image-based CNNs (e.g., NasNet [40] and AmoebaNet [21]). Having heterogeneous modules is motivated by the recent observations that video CNN architectures may need differently sized temporal filters at different layers (e.g., bottom-heavy vs. top-heavy [37]), and this is important for representing videos. In order to keep the entire search space manageable while evolving modules heterogeneously, we use the Inception meta-architecture with a fixed number of total modules. We also constrain the complexity of the connections between the layers within a module so that space-time layers are parallel, while making the evolution to explore temporal aspects of the modules.

Our architecture evolution particularly focuses on providing answers to the following questions: 1) Which space-time convolutional layer should the architecture use and where, to best benefit recognition? 2) How should we modify the modules (composed of multiple layers) originally designed for image classification for video recognition? 3) How many frames should each convolutional filter consider, depending on its layer type and location? That is, what are the optimal temporal kernel sizes (i.e., L) for each layer?

The Inception meta-architecture we use is illustrated in Figure 3. This meta architecture was also used in I3D and S3D, and is composed of three initial (space-time) convolutional layers (i.e., the 'stem') followed by nine Inception modules. Each Inception module originally has a two space-time convolutional layers as well as a pooling layer and multiple 1x1 convolutional layers (Figure 4). Our evolutionary search starts from the initial pool of architectures with such Inception modules, which is later modified based
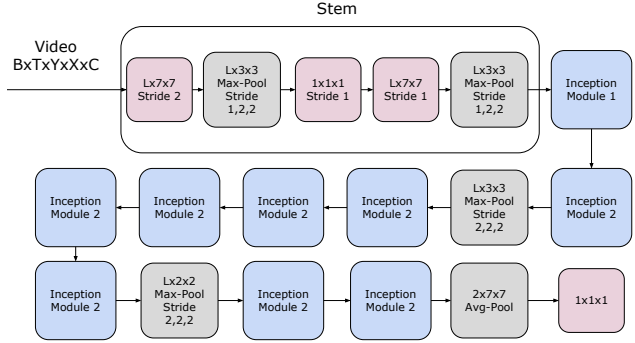

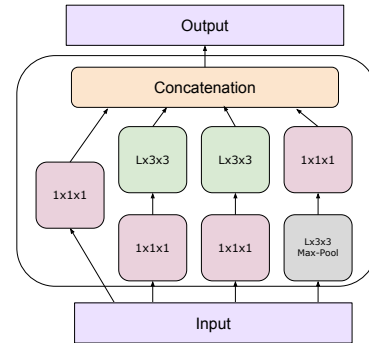Figure 3. Inception meta-architecture.

Figure 4. Structure of the base Inception module

on the mutations described in the following subsection.

The three convolutional layers mentioned in Section 3 are the layer types we consider in our architecture evolution, and we use $\{1, 3, 5, 7, 9, 11\}$ as the set of possible temporal kernel sizes. As a result, the architecture search space size becomes $O((3 \times 6 + 1)^{3+B \times N} + (6 + 1)^{D \times N})$ where B and D are the maximum of number of space-time conv and pooling layers we allow in each module and $N = 9$ is the number of Inception modules in the architecture. There are 3 individual layers (often also called a 'stem') before the $N = 9$ Inception modules. Each space-time conv. layer has 3×6 possible options and each space-time pooling has has 6 options. Also, there could be an option to omit the layer, making the total number of choices to be $3 \times 6 + 1$ and $6 + 1$. We fix the spatial size of the kernels to be 3×3.

Although the search space is very big, the idea is that an exhaustive search is not necessary and it is possible to find good local optima by evolving from various initial samples (i.e., architectures).

4.2. Evolutionary algorithm

We summarize our evolutionary algorithm in Algorithm 1. We follow the standard genetic algorithm setting: maintaining a population of size P, where each individual in the population is a particular architecture in our case. Architectures in the initial population are obtained by randomly sampling from our large search space (Section 4.1), encour-

aging diversity and exploration.

Algorithm 1 Evolutionary search algorithm

function SEARCH
    Randomly initialize the population, P
    Evaluate each individual in P
    for $i <$ number of evolutionary rounds **do**
        $S =$ random sample of 25 individuals
        $parent =$ the most fit individual in S
        $child = parent$
        for $\max(\lceil d - \frac{i}{r} \rceil, 1)$ **do**
            $child = mutate(child)$
        end for
        evaluate $child$ and add to population
        remove least fit individual from population
    end for
end function

At each round of the evolution (which can also be parallelized), our algorithm randomly selects a S number of samples from the entire population and compares their recognition performance. The architectures with the highest fitness (i.e., validation accuracy) becomes the 'parent', and mutation operators are applied to the selected parent to generate a new 'child' architecture to be added to the population. Whenever a new architecture is added, it is trained with the training set for a number of iterations, and is evaluated with a separate validation set (different from the actual test and validation sets) to measure the recognition accuracy. This performance becomes the 'fitness' of the architecture.

Having S where $1 < S \leq P$ controls the randomness in the parent selection. It avoids the algorithm repeatedly selecting the same parent, which might already be at a local maximum. It also allows the algorithm to be more easily parallelized by avoiding such situations. Our algorithm could be parallelized by distributing the loop over multiple workers, each responsible selecting S random samples from the population, deciding the parent, and mutating it to get the child.

The mutation operators modify the parent architecture, which we describe more in the subsubsection below. Importantly, we design the mutation in our algorithm to happen by applying multiple randomly chosen mutation operators. In order to encourage more diverse architectures, we developed the strategy of applying multiple mutation operators in the early stage of the evolution while reducing their numbers at the later stages, which is analogous to controlling the learning rate in a CNN model learning. As described in Algorithm 1, we apply $\max(d - \frac{i}{r}, 1)$ number of mutation operators where d is the maximum number of operators we want to apply in the beginning, and r controls how quickly we want to decrease their numbers linearly.

Once a child architecture is added to the population, in order to maintain the size of the population to P, the evolu-
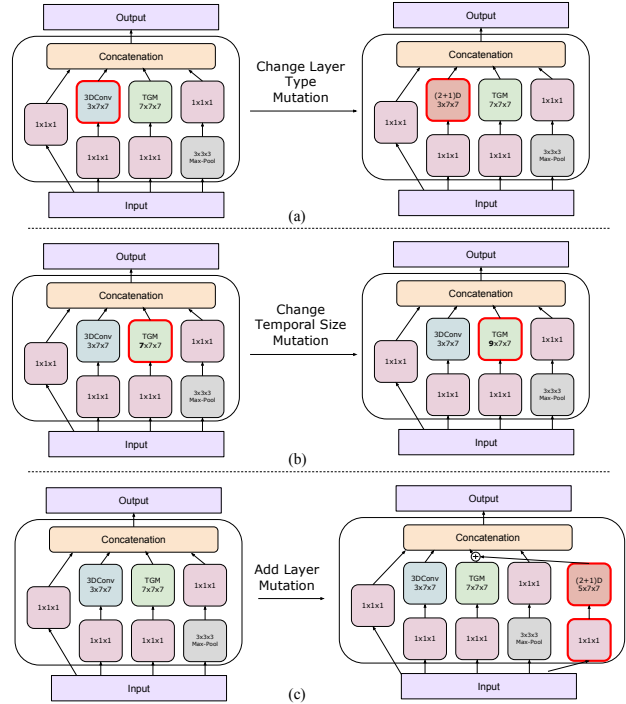

Figure 5. Example mutations applied to a module, including (a) layer type change, (b) filter length change, and (c) layer addition.

tionary algorithm selects an individual to discard from the pool. We tried different criteria including the lowest fitness as well as discarding the oldest architecture (i.e., [21]), which did not make much difference in our case. P and S are the hyper-parameters of the algorithm, $P = 100$ and $S = 25$ in the experiments.

Mutations. As described in Algorithm 1, once the parent architecture is selected from the pool of architectures, our approach applies 'mutations' to the parent architecture to generate a new child architecture. In order to explore the architecture search space we describe in Section 4.1 efficiently, we consider the following three mutation operators: (i) Select a space-time conv. layer within the parent architecture, and change its 'type'. (ii) Select a space time conv. layer or a pooling layer, and change the temporal size (i.e., L) of its filters. (iii) Select a module from the parent architecture, and add/remove a space-time conv. layer within it. We constrain the module so that its space-time convolutional layers are applied in parallel, and combine the parallel layers by averaging. Figure 5 illustrates examples of our mutation operators applied to layers of a module.

Ensemble. We obtain a number of top performing architectures after the evolutionary search is completed. Because of the design of our evolutionary algorithm to promote the population to maintain diverse individual architectures, unlike reinforcement learning-based search methods [40] that optimize policies to generate the architecture. Thus a con-

struction of a good ensemble is very likely. We do late fusion of the top few architectures (e.g., 3 or 5), adding the outputs of their softmax layers: $F^*(x) = \sum_i F_i(x)$ where x is the input video and F_i are the top CNN models. In the experiments, we found our approach obtains very diverse top performing architectures, and further ensembling allows us to improve the overall recognition meaningfully.

We named our final ensemble network as EvaNet (Evolved Video Architecture).

5. Experiments

Although our evolutionary architecture search is applicable for various different video understanding tasks, we particularly focus on the problem of human activity recognition in our experiments. That is, our video CNN architectures are evolved using public activity recognition video datasets including HMDB [12] and Kinetics [10], with the architecture fitness being the activity recognition accuracy on the validation set (a held-out portion of the training set). In all our experiments, our evolutionary algorithm never had any access to the actual test set during its training and evolution. Fitness of the architectures during evolution was measured with the separate set generated from the training set.

5.1. Datasets

HMDB: HMDB [12] is a public video classification dataset with ∼7000 videos of 51 action classes.

Kinetics: We use Kinetics-400 dataset [10] with the videos available on YouTube in November 2018. Note that this version has less training videos than the initial version of Kinetics-400, making the dataset more difficult. This version has 225,946 training, 18,584 validation and 36,500 test videos, about 25k videos smaller than the original Kinetics-400 dataset (i.e., missing about 8% of train/val/test data). We report our numbers on this version, and compare ours with other works tested on the new version as well as the reported numbers on the original version.

Charades: Charades [25] is an activity recognition dataset with ∼10K videos, whose durations are 30 seconds on average. We use Charades particularly to confirm whether our architecture evolution finds structures different from those found with shorter videos like Kinetics. We use the classification evaluation.

5.2. Experimental setup and details

Our architecture evolution was done with multiple parallel workers. A smaller input size and a fewer number of iterations were used when training each architecture to measure their relative performances, which become the sample 'fitness'. Please check the appendix in the supplementary material for more details.

We continued our evolution for 2000 rounds, newly generating, training/evaluating, and discarding 2000 CNN ar-

chitectures. Note that ∼300 rounds were often sufficient to find good architectures (Figure 7). Once the architecture evolution is complete and the top performing CNNs are found, we fully trained the models with higher resolution inputs (64x224x224 for HMDB/Kinetics and 128x224x224 for Charades) and with more training iterations (120k for Kinetics and 64k for Charades).

ImageNet pre-training. ImageNet pretraining is standard practice for video recognition [1]. Taking advantage of 2D filters learned from ImageNet allows a model to start its training with more reasonable initial 3D filters. This is done by uniformly inflating the ImageNet-trained 2D filters to 3D in the case of standard 3D conv. layers (i.e., as in [1]), and is done by initializing the 2D kernel part in (2+1)D and inflated TGM layers with such filters (i.e., as in [37]). This not only allows the training of models to converge faster but also provides better final performances, benefiting from millions of images.

Although our video CNN evolution approach itself is general and it allows mutations generating various architectures/modules without any constraint, two constraints were used in our experiments to make our evolved architectures directly benefit from ImageNet pre-trained weights more easily: (1) We constrained the space-time convolutional layers and pooling layers in our modules to be parallel. We also always inserted a 1x1x1 conv. layer before each space-time conv. layer and after each pooling layer, just as was done in the basic Inception module (Fig. 4). (2) We forced each of our module to have at least two space-time convolutional layers and one space-time pooling layer as in the basic Inception module. The outputs of multiple space-time conv. or pooling layers were combined by element-wise averaging (i.e., Fig. 5(c)), essentially allowing our modules to have the same input/output channel size as the modules in the pre-trained basic Inception model.

5.3. Baselines

We compare our results to the state-of-the-art on activity recognition, and further implemented those following the same Inception meta-architecture for direct comparison. We trained (1) the original I3D with standard 3D convolutional layers. We also trained the Inception models with: (2) 3D conv. layers with $L = 3$, (3) (2+1)D conv. layers, and (4) the proposed iTGM layers. The difference between (1) and (2) is that (1) uses $L = 7$ in the first 3D conv. layer and $L = 3$ in all the other 3D layers (which was a handcrafted design), while (2) uses $L = 3$ in all its layers.

5.4. Results

The architecture evolution was conducted with each dataset individually, finding the optimal architecture for each dataset. Below summarizes the classification accura-

Table 1. HMDB split 1 comparison to baselines, with and without Kinetics pre-training. The models were all initialized with ImageNet weights. '3D-Ensemble' was formed by combining 3 3D Conv baselines (trained separately), and 'iTGM-Ensemble' was formed by combining 3 3D Conv baselines and 2 iTGM baselines.

	HMDB			HMDB(pre-train)		
	RGB	Flow	RGB+F	RGB	Flow	RGB+F
Baselines						
I3D	49.5	61.9	66.4	74.8	77.1	80.1
3D Conv	47.4	60.5	65.9	74.3	76.8	79.9
(2+1)D Conv	27.8	56.4	51.8	74.4	76.5	79.9
iTGM Conv	56.5	62.5	68.2	74.6	76.7	79.9
3D-Ensemble			67.6			80.4
iTGM-Ensemble			69.5			80.6
Top individual models from evolution						
Top 1	60.7	63.2	70.3	74.4	78.5	**81.2**
Top 2	63.4	62.5	**71.2**	75.6	78.4	80.5
Top 3	60.5	63.1	70.5	75.5	78.6	79.8
Top 4	59.0	60.9	69.8	-	-	-
Top 5	61.6	62.5	70.3	-	-	-
EvaNet			**72.8**			**82.4**

Table 2. HMDB performances averaged over the 3 splits.

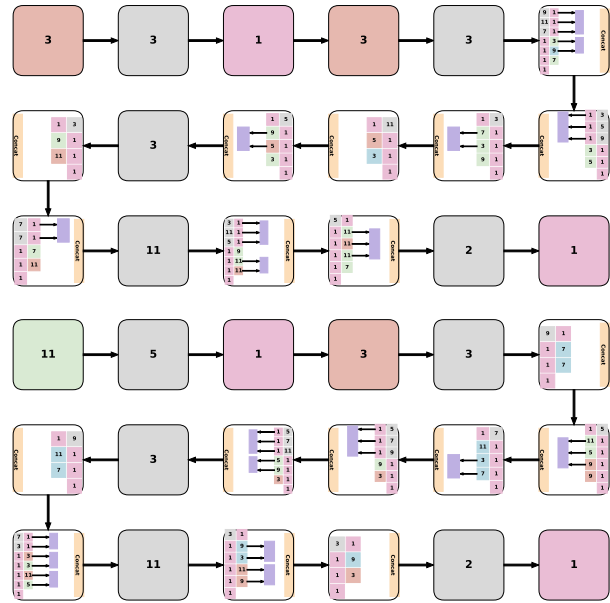|	
Two-stream [26]	59.4
Two-stream+IDT [3]	69.2
R(2+1)D [32]	78.7
Two-stream I3D [1]	80.9
PoTion [2]	80.9
Dicrim. Pooling [34]	81.3
DSP [33]	81.5
3D-Ensemble (our baseline)	79.9
iTGM-Ensemble (our baseline)	80.1
EvaNet (ours)	**82.1**

Figure 6. Example architectures: Kinetics RGB (top); Charades (bottom).

Table 3. Kinetics performance comparison to baselines, all initialized with ImageNet weights.

	RGB	Flow	RGB+F
Baselines			
3D Conv	70.6	62.1	72.6
(2+1)D Conv	71.1	62.5	74.3
iTGM Conv	71.2	62.8	74.4
3D-Ensemble			74.6
iTGM-Ensemble			74.7
Top individual models from evolution			
Top 1	71.9	63.8	**75.8**
Top 2	71.7	64.9	75.2
Top 3	72.9	64.8	75.4
EvaNet			**76.8**

cies of the video CNN architectures our evolution automatically found and their ensembles.

HMDB: Table 1 shows the accuracy of our evolved CNNs compared to the baseline architectures following the same Inception meta-architecture. Furthermore, because of the diverse nature of the evolved architectures, we were able to form their ensemble and increase their performance further. As we can confirm, our heterogeneous ensemble was superior to the ensembles obtained by training the same architecture (e.g., I3D) multiple times. Table 2 compares our performance with the previous state-of-the-arts. To our knowledge, we are obtaining the highest HMDB video classification accuracy: 82.1%

Kinetics: Table 3 shows the classification accuracy of our evolved CNNs on Kinetics-400. Similar to Table 1 we are able to confirm that our architecture evolution finds better performing models than any prior models with the same meta-architecture. Table 4 compares our algorithm with the

state-of-the-arts. Here too we can see that the evolved architectures outperform the previous ones. Note that the performances reported on the older version of Kinetics-400 are not directly comparable to our numbers.

Charades: We also test our algorithms on the popular Charades dataset. Table 5 compares our approach against the previously reported performances. As shown, we outperform the state-of-the-art and establish a new one. Our CNNs only use RGB (i.e., one-stream) in this experiment.

Smaller search space (i.e., fewer mutation ops). We conducted an experiment with a simpler version of our evolution with fewer mutation operators. Reducing the number of mutation operators restricts possible architectures, essentially making the architecture search space to be smaller.

Table 4. Kinetics-400 accuracy. Note that * are the reported numbers on the initial Kinetics dataset, which is no longer available. We report the numbers based on the new Kinetics version from Nov 2018. The new version has 8% less training/validation videos.

Method	Kinetics-400 new	old
Two-stream I3D [1]	72.6	74.1*
Two-stream (2+1)D [32]	-	75.4*
Two-stream S3D-G [37]	76.2	77.2*
Non-local NN [35]	-	77.7*
EvaNet (ours)	76.8	-

Table 5. Charades classification results against state-of-the-arts. Similar to [35], we use Kinetics pre-training.

	mAP
Two-Stream [24]	18.6
Two-Stream + LSTM [24]	17.8
Async-TF [24]	22.4
TRN [39]	25.2
Dicrim. Pooling [34]	26.7
Non-local NN [35]	37.5
3D-Ensemble (our baseline)	35.2
iTGM-Ensemble (our baseline)	35.7
EvaNet (ours)	**38.1**

We implemented a version of our architecture evolution only using the mutation operators (i) and (ii), while also constraining each module to share the same space-time conv. type and temporal filter length across its layers. On Kinetics, this simpler evolution gave us the performance of 75.5, which is lower than our 76.8. The result suggests that it is meaningful to have our 'layer addition mutation'.

Evolution vs. random search. We compared our architecture evolution with the random architecture search (Figure 7). We observe that both the evolution and the random search accuracies improve as they explore more samples (benefiting from our search space designed). However, our architecture evolution obtains much higher accuracy and much more quickly with few initial rounds of evolution, suggesting our mutations are being effective. The evolution was done on a smaller subset of Kinetics-400 dataset and for 1000 iterations.

Architecture findings. Figures 1 and 6 show examples of the architectures found. The grey blocks are max-pooling, blue is 3D conv, red is (2+1)D conv, green is iTGM, pink is 1x1x1, and purple is element-wise averaging. The numbers represent the temporal size (L) of the kernel. More architecture examples can also be found in Appendix in the supplementary material. Although the top architectures found by our evolutionary algorithm are similarly good in terms of classification accuracy, they turned out to be extremely di-

verse (e.g., Fig. 1 vs. Fig. 6-bottom, both from Charades). This directly suggest forming their ensemble would benefit the recognition even further, which is exactly what we observed in Tables 1 and 3.

Interesting substructures discovered include: (1) modules combining multiple space-time pooling layers with different temporal intervals and (2) modules heavily relying on Inflated TGM or (2+1)D conv. layers instead of standard 3D conv. layers. Such modules were commonly observed at the middle-level and the bottom-level of the architectures. The top-most module more commonly relied on 3D conv. layers, suggesting that they can be used to capture subtle spatio-temporal patterns once the representation becomes more abstract and their dimensionality decrease

Video CNN architectures evolve differently depending on the datasets. This is very natural, and we were able to explicitly confirm this. The architectures have many more layers with longer space-time filters (e.g., 9 or 11) when evolved for Charades, while they only had a small number of them when evolved for HMDB or Kinetics. An average activity duration in Charades videos are around 12 seconds, while HMDB and Kinetics videos are on the average of 3 to 5 seconds. Different architectures are needed for different datsets/tasks, and we are providing an evolutionary approach to automate the architecture design.

6. Conclusion

We presented the new approach of evolving video CNN architectures. Our evolutionary algorithm starts with a pool of randomly generated architectures, and they are evolved over multiple rounds using the mutation operators we introduced. We particularly focused on making the algorithm explore architectures capturing space-time information in videos using combinations of standard 3D conv. layers, (2+1)D conv. layers, and newly introduced inflated TGM layers. The resulting architecture outperformed handcrafted architectures, and we also found interesting substructures including the modules with parallel space-time layers focusing on different temporal intervals.

References

[1] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[2] V. Choutas, P. Weinzaepfel, J. Revaud, and C. Schmid. Potion: Pose motion representation for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[3] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1933–1941, 2016.
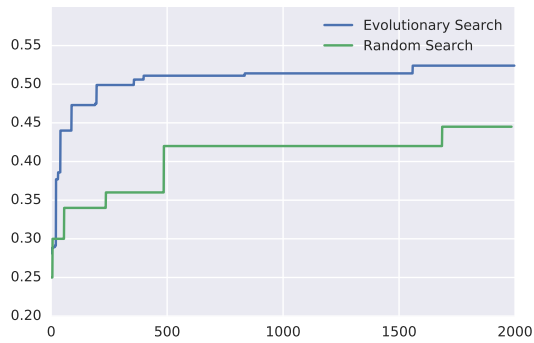
Figure 7. Random search vs. our evolutionary algorithm.

[4] K. Hara, H. Kataoka, and Y. Satoh. Learning spatio-temporal features with 3d residual networks for action recognition. In *Proceedings of the ICCV Workshop on Action, Gesture, and Emotion Recognition*, volume 2, page 4, 2017.

[5] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[6] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[7] S. Jegou, M. Drozdzal, D. Vazquez, A. Romero, and Y. Bengio. One hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. 2016.

[8] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. In *International Conference on Machine Learning (ICML)*, pages 495–502, 2010.

[9] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1725–1732, 2014.

[10] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.

[11] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. 2012.

[12] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2011.

[13] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. 1998.

[14] M. Lin, Q. Chen, and S. Yan. Network in network. 2013.

[15] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar. Focal loss for dense object detection. 2017.

[16] W. Luo, B. Yang, and R. Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. 2018.

[17] E. Mansimov, N. Srivastava, and R. Salakhutdinov. Initialization strategies of spatio-temporal convolutional neural networks. *arXiv preprint arXiv:1503.07274*, 2015.

[18] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4694–4702. IEEE, 2015.

[19] A. Piergiovanni, C. Fan, and M. S. Ryoo. Learning latent sub-events in activity videos using temporal attention filters. In *Proceedings of the American Association for Artificial Intelligence (AAAI)*, 2017.

[20] A. Piergiovanni and M. S. Ryoo. Temporal gaussian mixture layer for videos. *arXiv preprint arXiv:1803.06316*, 2018.

[21] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le. Regularized evolution for image classifier architecture search. *arXiv preprint arXiv:1802.01548*, 2018.

[22] E. Real, S. Moore, A. Selle, Y. L. S. Saurabh Saxena, Q. Le, and A. Kurakin. Large-scale evolution of image classifiers. In *International Conference on Machine Learning (ICML)*, 2017.

[23] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. 2015.

[24] G. A. Sigurdsson, S. Divvala, A. Farhadi, and A. Gupta. Asynchronous temporal fields for action recognition. *arXiv preprint arXiv:1612.06371*, 2016.

[25] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016.

[26] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems (NIPS)*, pages 568–576, 2014.

[27] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016.

[28] M. Tan, B. Chen, R. Pang, V. Vasudevan, and Q. V. Le. Mnasnet: Platform-aware neural architecture search for mobile. *arXiv preprint arXiv:1807.11626*, 2018.

[29] D. Tran, L. D. Bourdev, R. Fergus, L. Torresani, and M. Paluri. C3d: generic features for video analysis. *CoRR, abs/1412.0767*, 2(7):8, 2014.

[30] D. Tran, J. Ray, Z. Shou, S.-F. Chang, and M. Paluri. Convnet architecture search for spatiotemporal feature learning. *arXiv preprint arXiv:1708.05038*, 2017.

[31] D. Tran, J. Ray, Z. Shou, S.-F. Chang, and M. Paluri. Convnet architecture search for spatiotemporal feature learning. *arXiv preprint arXiv:1708.05038*, 2017.

[32] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[33] J. Wang and A. Cherian. Learning discriminative video representations using adversarial perturbations. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2018.

[34] J. Wang, A. Cherian, F. Porikli, and S. Gould. Video representation learning using discriminative pooling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1149–1158, 2018.

[35] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[36] S. Xie, R. Girshick, P. Dollar, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[37] S. Xie, C. Sun, J. Huang, Z. Tu, and K. Murphy. Rethinking spatiotemporal feature learning for video understanding, 2018.

[38] S. Yeung, O. Russakovsky, N. Jin, M. Andriluka, G. Mori, and L. Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision (IJCV)*, pages 1–15, 2015.

[39] B. Zhou, A. Andonian, and A. Torralba. Temporal relational reasoning in videos. *arXiv preprint arXiv:1711.08496*, 2017.

[40] B. Zoph and Q. Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2017.

[41] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

A. Evolution training details

Our architecture evolution was done with 50 parallel workers. Each worker selects $S = 25$ random samples from the population to generate one new child architecture based on the individual with the highest fitness (i.e., the parent). The architecture is trained using 12 GPUs on the training data. As training video CNNs is computationally expensive, during the search, we train the models with video segments of size $32 \times 176 \times 176$ (for HMDB and Kinetics) or $64 \times 176 \times 176$ (for Charades) where 32 and 64 are the number of frames. We use the batch size of 144 (12 per GPU). Each newly generated child architecture is trained for 1000 iterations (i.e., it looks at 144000 samples), then evaluated with a separate validation set of 1000 examples. The classification accuracy measured using the validation becomes the 'fitness' used in our algorithm. We observed that relative recognition performances of the models (on the validation set) is stable after training for 1000 iterations, and we used this setting in our evolutionary algorithm to reduce the model training time necessary for the architecture evaluation.

B. Mutation Rate

In Fig. 8, we compare the architecture evolution done with a constant mutation rate of 1 or 3 (per round) and our annealed mutation rate. As we described in the main section of the paper, our evolutionary algorithm applies a set of
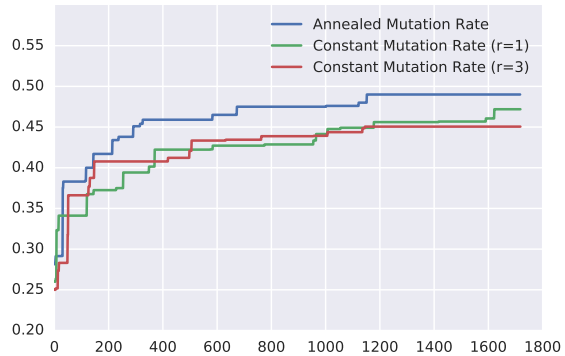

Figure 8. Comparison of the architecture search with various mutation rates. We observe that the constant rate takes longer to reach higher performance while the higher mutation rate initially learns faster, but plateaus at a lower value. Annealing the mutate rate based on the number of architectures evaluated provides the best performance.

random mutation operators at each round. In our annealed mutation rate strategy, the number of the mutation operators to apply is decided based on the evolution round i: it starts with $d = 7$ mutations initially and it is linearly decreased by $\lfloor i/r \rfloor$ where r is 100 in our experimental setting. That is, at the ith round, a total of $max(\lceil d - i/r \rceil, 1)$ random mutations were applied to the parent. We find that the annealed mutation rate performs the best. Our strategy allows the search to explore more diverse architectures based on the best initial models, but then refine the top performing models after many evolution rounds.

C. Discovered Architectures

We illustrate the architectures found in the following figures. The color of each layer corresponds to a specific layer type, as illustrated in Fig. 9. More specifically, blue, green, red are for space-time 3D conv layer, (2+1)D layer, and iTGM layer. Pink is for 1x1x1 conv layer, and grey is for max pooling. Average pooling layer and channel concatenation layer are indicated with purple and organge colors.

In Figures 10, 11, and 12, we show the architectures found when searching on Kinetics using RGB inputs. We observe that the networks learn quite different architectures. For example, the third inception module is quite different in all three networks. In Figures 13, 14, and 15, we illustrate the models found when searching on Kinetics using optical flow as input. When using optical flow as input, we observe that the architectures perfer to use layers with shorter temporal durations, using very few layers with size 11 and 9 when compared to the RGB networks. (2+1)D conv layers and iTGM layers were used much more commonly in both RGB and optical flow architectures. Parallel space-time conv and pooling layers with different temporal
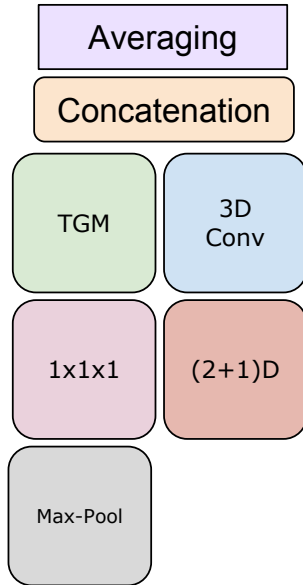
Figure 9. The color for each layer type.

lengths were also very commonly observed.

In Figures 16, 17, and 18, we illustrate the architectures found when searching on Charades. We observe that on Charades, the architectures generally capture longer temporal intervals (e.g., the first layer has size 11) and many layers contain longer kernels (i.e., 9 and 11) especially compared to the architectures found on Kinetics.
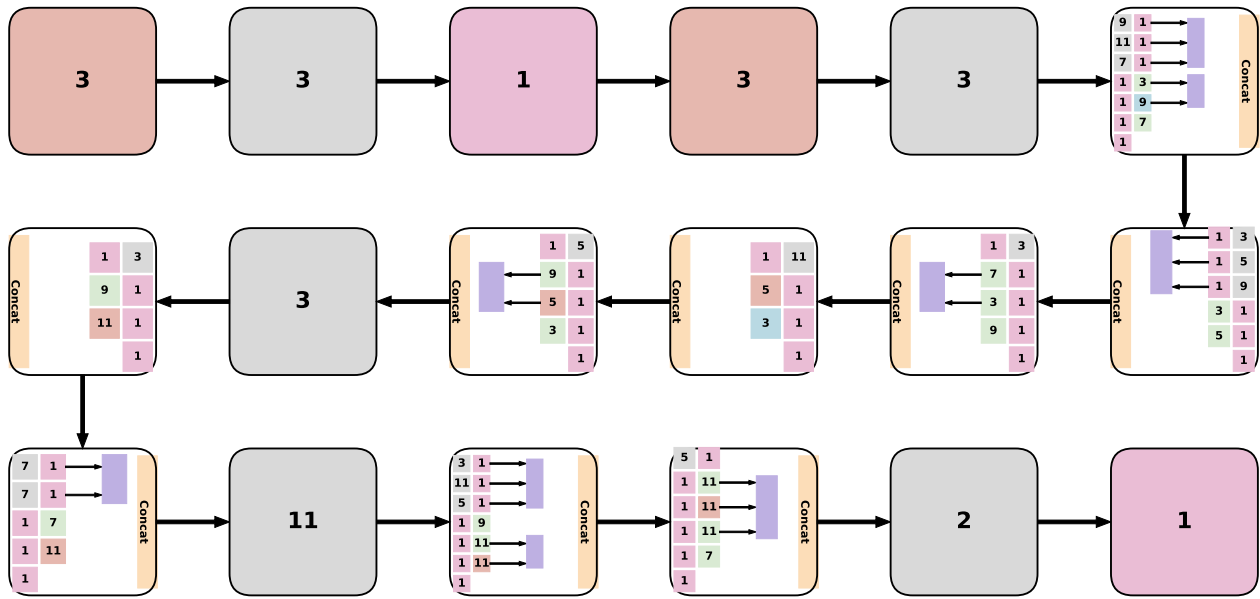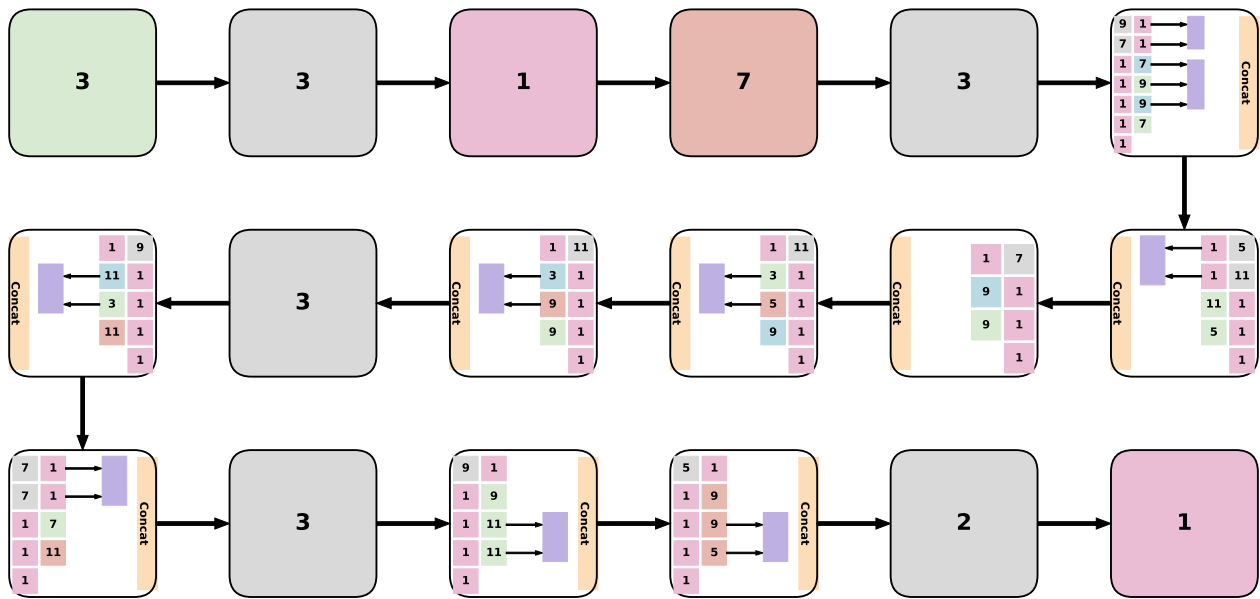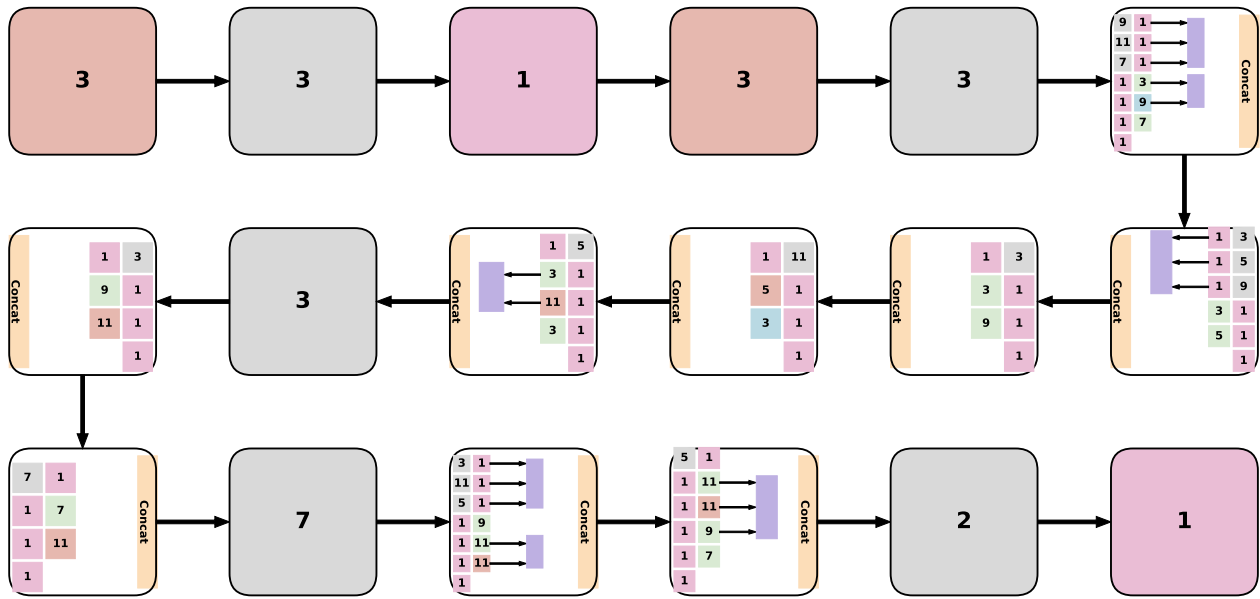
Figure 10. Kinetics RGB Model 1.

Figure 11. Kinetics RGB Model 2.

Figure 12. Kinetics RGB Model 3.
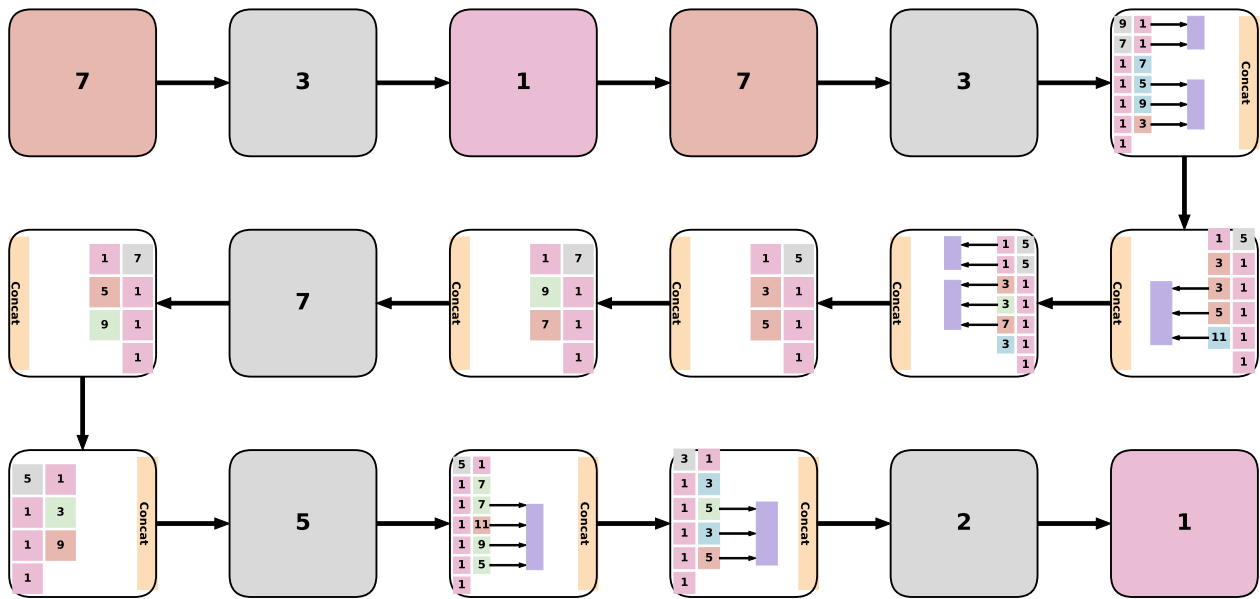
Figure 13. Kinetics optical flow Model 1.
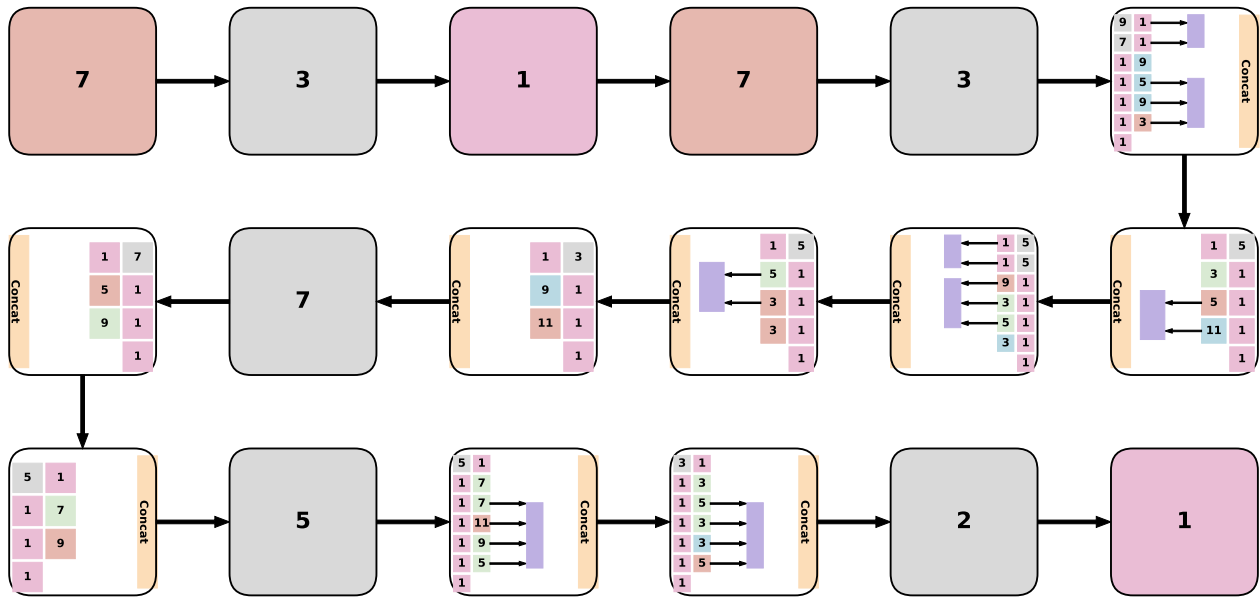
Figure 14. Kinetics optical flow Model 2.
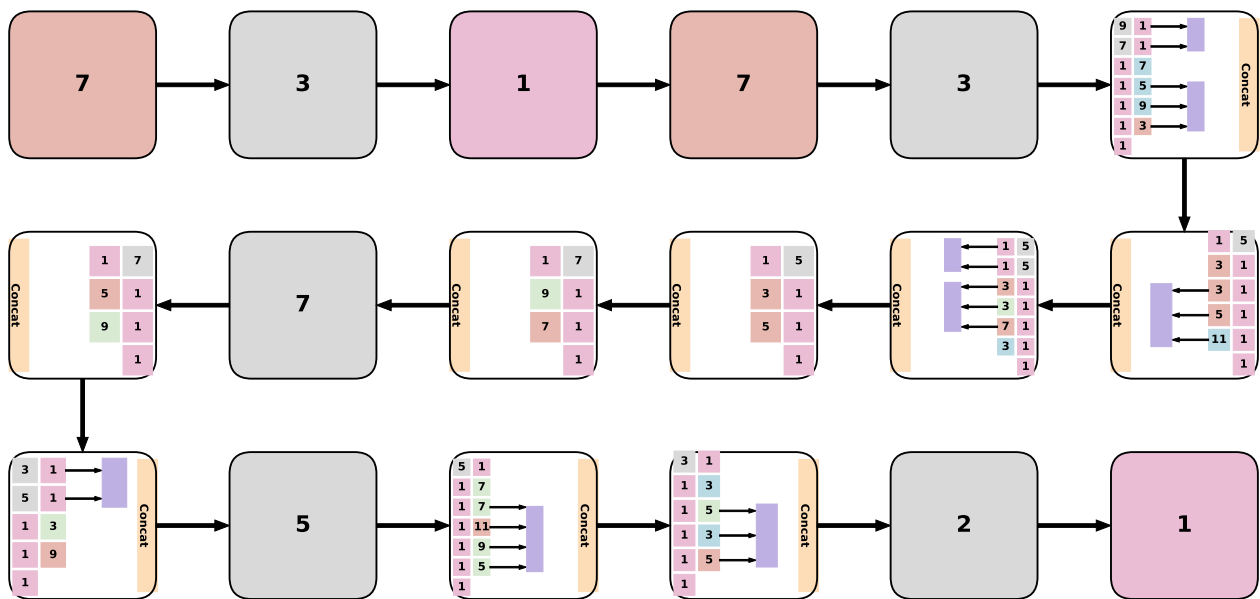
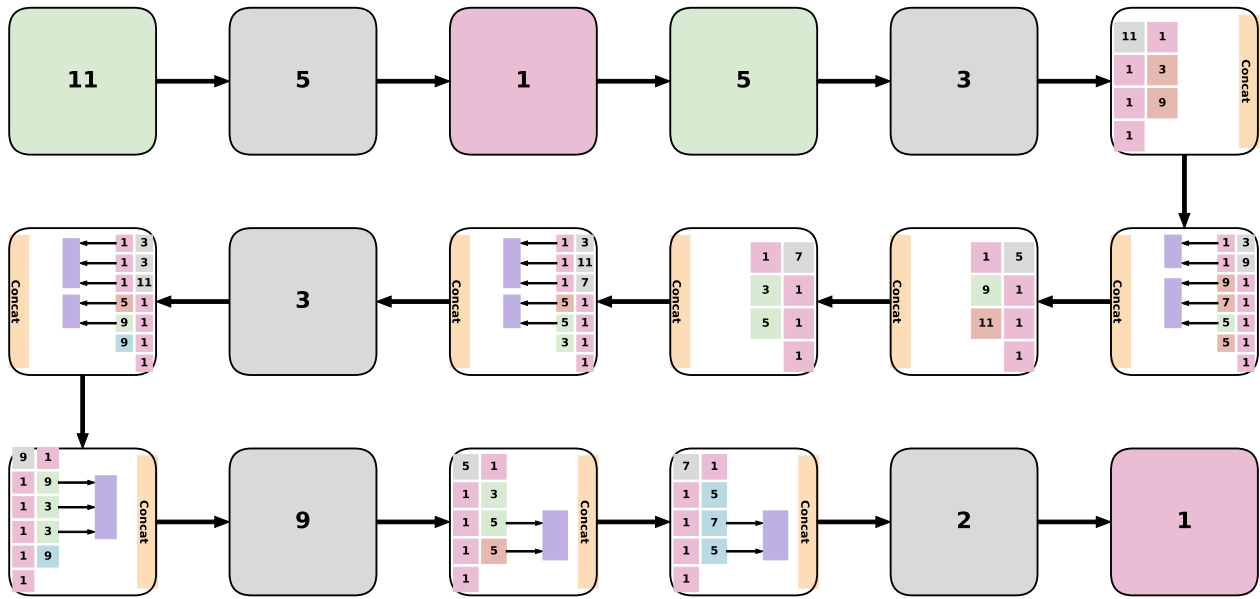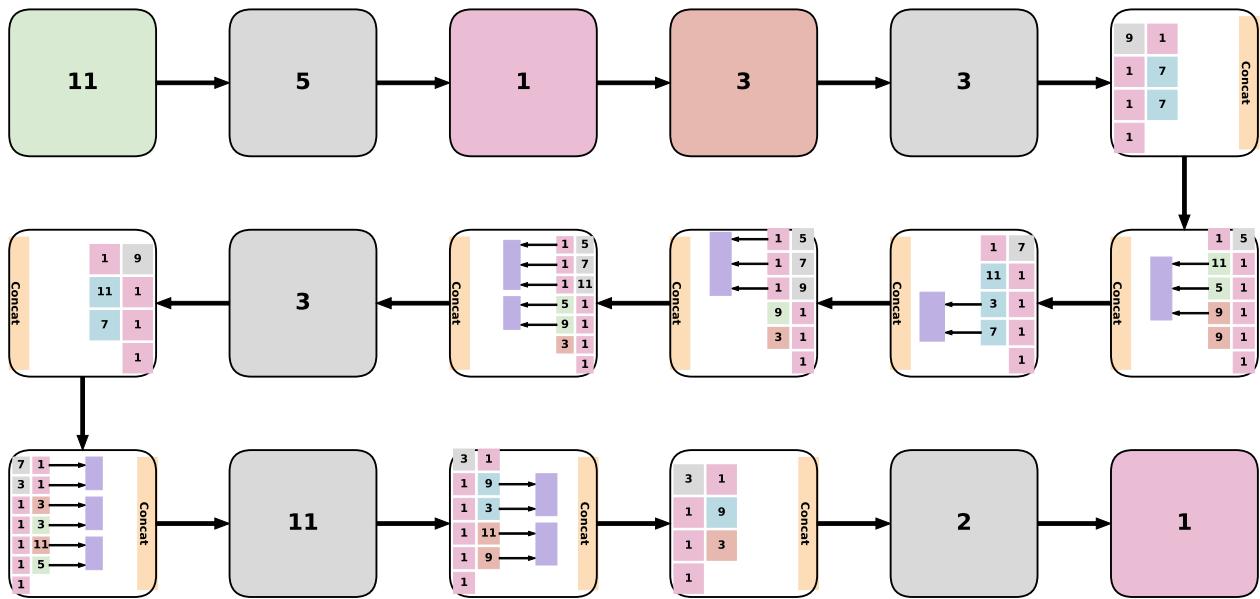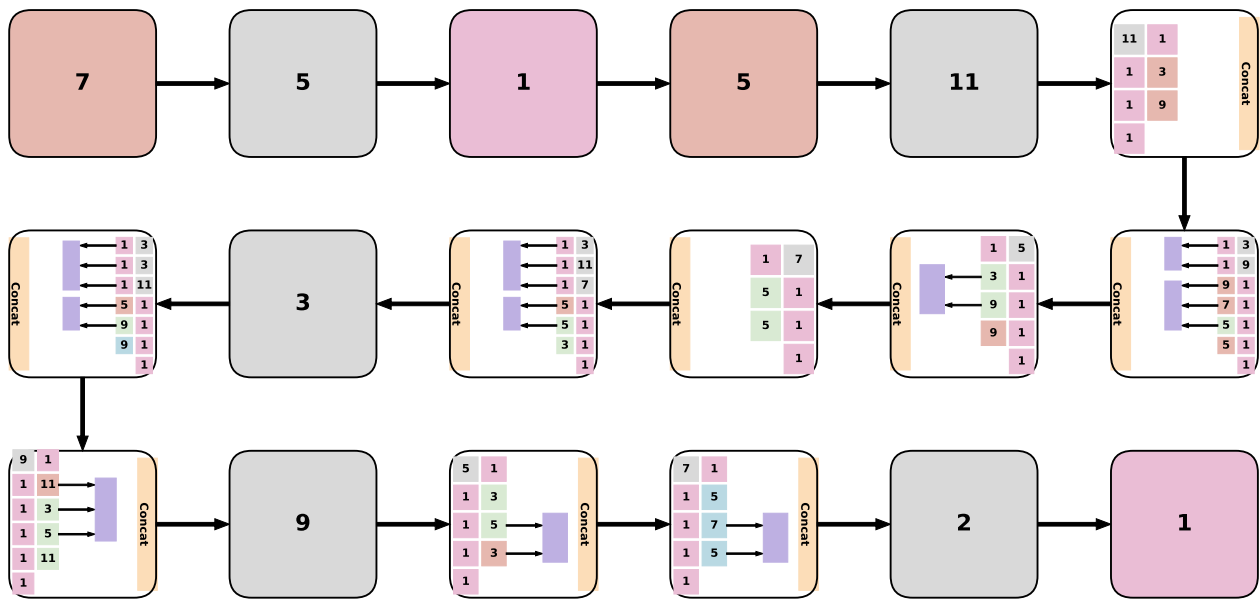Figure 15. Kinetics optical flow Model 3.

Figure 16. Charades RGB Model 1.

Figure 17. Charades RGB Model 2.

Figure 18. Charades RGB Model 3.