

Temporal and Object Quantification Networks

Jiayuan Mao^{1*}, Zhezeng Luo^{1*}, Chuang Gan², Joshua B. Tenenbaum¹, Jiajun Wu³,
Leslie Pack Kaelbling¹ and Tomer D. Ullman⁴

¹Massachusetts Institute of Technology

²MIT-IBM Watson AI Lab

³Stanford University

⁴Harvard University

Abstract

We present *Temporal and Object Quantification Networks* (TOQ-Nets), a new class of neuro-symbolic networks with a structural bias that enables them to learn to recognize complex relational-temporal events. This is done by including reasoning layers that implement finite-domain quantification over objects and time. The structure allows them to generalize directly to input instances with varying numbers of objects in temporal sequences of varying lengths. We evaluate TOQ-Nets on input domains that require recognizing event-types in terms of complex temporal relational patterns. We demonstrate that TOQ-Nets can generalize from small amounts of data to scenarios containing more objects than were present during training and to temporal warpings of input sequences.

1 Introduction

Every day, people interpret events and actions in terms of concepts, defined over temporally evolving relations among agents and objects [Zacks *et al.*, 2007; Stränger and Hommel, 1996]. For example, in a soccer game, people can easily recognize when one player has control of the ball, when a player *passes* the ball to another player, or when a player is *offsides*. Although it requires reasoning about intricate relationships among sets of objects over time, this cognitive act is effortless, intuitive, and fast. It also generalizes directly over different numbers and arrangements of players, and detailed timings and trajectories. In contrast, most computational representations of sequential concepts are based on fixed windows of space and time, and lack the ability to perform relational generalization.

In this paper, we develop generalizable representations for learning complex activities in time sequences from realistic data. As illustrated in Fig. 1, we can describe complex events with a first-order linear temporal logic [FO-LTL; Pnueli 1977] formula, which allows us to flexibly decompose an input sequence into stages that satisfy different criteria over time.

Object quantifiers (\forall and \exists) are used to specify conditions on sets of objects that define each stage. Such representations immediately support generalization to situations with a varying number of objects, and sequences with different time warpings.

More concretely, the variety of complicated spatio-temporal trajectories that *high pass* can refer to in a soccer game can be described in these terms: in a high pass from player A to teammate B , A is *close* to the ball ($distance(A, ball) < \theta_1$) and moving ($velocity(A) > \theta_2$) *until* the ball moves over the ground ($z_{position}(ball) > \theta_3$), which is in turn *until* teammate B gets control of the ball ($teammate(A, B) \wedge distance(B, ball) < \theta_1$). Beyond modeling human actions in physical environments, these structures can be applied to events in any time sequence of relational states, e.g., characterizing particular offensive or defensive maneuvers in board games such as chess or in actual conflicts, or detecting a process of money-laundering amidst financial transaction records.

In this paper, we propose a neuro-symbolic approach to learning to recognize temporal relational patterns, called *Temporal and Object Quantification Networks* (TOQ-Nets), in which we design structured neural networks with an explicit bias that represents finite-domain quantification over both entities and time. A TOQ-Net is a multi-layer neural network whose inputs are the properties of agents and objects and their relationships, which may change over time. Each layer in the TOQ-Net performs either *object* or *temporal* quantification.

The key idea of TOQ-Nets is to use *tensors* to represent relations between objects, and to use tensor pooling operations over different dimensions to realize temporal and object quantifiers (G and \forall). For example, the colorful matrix in Fig. 1(a) can be understood as representing a unary color property of a set of 8 entities over a sequence of 8 time steps. Representing a sequence of relations among objects over time would require a 3-dimensional tensor. Crucially, the design of TOQ-Nets allows *the same network weights* to be applied to domains with different numbers of objects and time sequences of different lengths. By stacking object and temporal quantification operations, TOQ-Nets can easily learn to represent higher-level sequential concepts based on the relations between entities over time, starting from low-level sensory input and supervised with only high-level class labels.

There are traditional symbolic learning or logic synthesis methods that construct first-order or linear temporal logic

*: equal contribution; order determined by a coin toss.

Email: {jiayuanm, ezzluo}@mit.edu

Project page: <http://toqnet.csail.mit.edu>

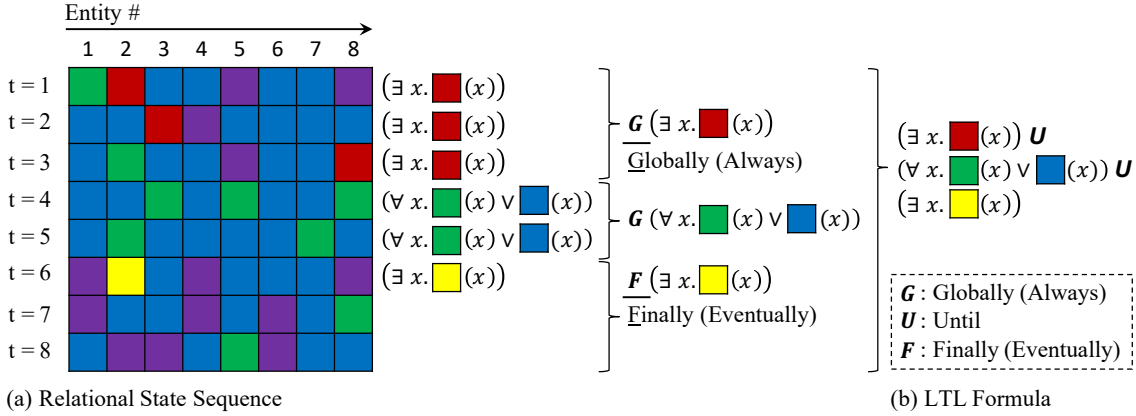


Figure 1: (a) An input sequence composed of relational states: each column represents the state of an entity that changes over time. A logic formula describes a complex concept or feature that is true of this temporal sequence using object and temporal quantification. The sequence is segmented into three stages: throughout the first stage, $\color{red}\blacksquare$ holds for at least one entity, until the second stage, in which each entity is always either $\color{blue}\blacksquare$ or $\color{green}\blacksquare$, until the third stage, in which $\color{yellow}\blacksquare$ eventually becomes true for at least one of the entities. (b) Such events can be described using first-order linear temporal logic expressions.

expressions from accurate symbolic data [Neider and Gavran, 2018; Camacho *et al.*, 2018; Chou *et al.*, 2020]. TOQ-Nets take a different approach and can learn from noisy data by backpropagating gradients, which allows them to start with a general perceptual processing layer that is directly fed into logical layers for further processing.

We evaluate TOQ-Nets on two perceptually and conceptually different benchmarks: trajectory-based sport event detection and human activity recognition, demonstrating several important contributions. First, TOQ-Nets outperform both convolutional and recurrent baselines for modeling temporal-relational concepts across benchmarks. Second, by exploiting temporal-relational features learned through supervised learning, TOQ-Nets achieve strong few-shot generalization to novel actions. Finally, TOQ-Nets exhibit strong generalization to scenarios with more entities than were present during training and are robust w.r.t. time warped input trajectories. These results illustrate the power of combining symbolic representational structures with learned continuous-parameter representations to achieve robust, generalizable interpretation of complex relational-temporal events.

2 TOQ-Nets

The input to a TOQ-Net is a tensor representation of the properties of all entities at each moment in time. For example, in a soccer game, the input encodes the position of each player and the ball, as well as their team membership at each step of an extended time interval. The output is a label of the category of the sequence, such as the type of soccer play it contains.

The first layer of a TOQ-Net (Fig. 2 (i)) extracts temporal features for each entity with an *input feature extractor* that focuses on entity features within a fixed and local time window. These features may be computed, e.g., by a convolutional neural network or a bank of parametric feature templates. The output of this step is a collection of nullary, unary, and binary relational features over time for all entities. Throughout the paper we will assume that all output tensors of this layer are

binary-valued, but it can be extended directly to real-valued functions. This input feature extractor is task-specific and is not the focus of this paper.

Second, these temporal-relational features go through several *relational reasoning* layers (RRLs), detailed in Section 2.2, each of which performs linear transformations, sigmoid activation, and *object quantification* operations. The linear and sigmoid functions allow the network to realize learned Boolean logical functions, and the object quantification operators can realize quantifiers. Additional RRLs enable deeper nesting of quantified expressions, as illustrated in Fig. 2. All operations in these layers are performed for all time steps in parallel.

Next, the RRLs perform a final quantification, computing for each time step a set of a nullary features that are passed to the *temporal reasoning layers* (TRLs), as detailed in Section 2.3. Each TRL performs linear transformations, sigmoid activation, and *temporal quantification*, allowing the model to realize a subset of linear temporal logic [Pnueli, 1977]. As with RRLs, adding more TRLs enables the network to realize logical forms with more deeply nested temporal quantifiers.

In the last layer, all object and time information is projected into a set of features of the initial time step, which summarize the temporal-relational properties of the entire trajectory (e.g., “the kicker eventually scores”), and fed into a final softmax unit to obtain classification probabilities for the sequence.

It is important to understand the representational power of this model. The *input transformation layer* learns basic predicates and relations that will be useful for defining more complex concepts, but no specific predicates or relations are built into the network in advance. The relational reasoning layers build quantified expressions over these basic properties and relations, and might construct expressions that could be interpretable as “the player is close to the ball.” Finally, the temporal reasoning layer applies temporal operations to these complex expressions, such as “the player is close to the ball until the ball moves with high speed.” Critically, *none* of the symbolic properties or predicates are hand defined—they are all constructed by the initial layer in order to enable the

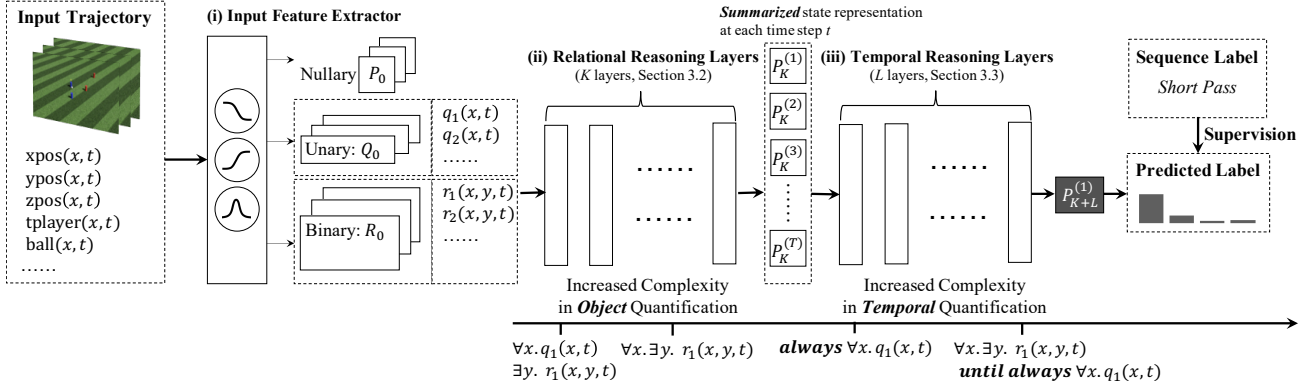


Figure 2: A TOQ-Net contains three modules: (i) an input feature extractor, (ii) relational reasoning layers, and (iii) temporal reasoning layers. To illustrate the model’s representational power, we show that logical forms of increasing complexity can be realized by stacking multiple layers.

network to express the concept it is being trained on.

TOQ-Nets are not fully first-order: all quantifiers operate only over the finite domain of the input instance, and can be seen as “short-hand” for finite disjunctions or conjunctions over objects or time points. In addition, the depth of the logical forms it can learn is determined by the fixed depth of the network. However, our goal is not to fully replicate temporal logic, but to bring ideas of object and temporal quantification into neural networks, and to use them as structural inductive biases to build models that generalize better from small amounts of data to situations with varying numbers of objects and time courses.

2.1 Temporal-Relational Feature Representation

TOQ-Nets use tensors as internal representations between layers; they represent, all at once, the values of all predicates and relations grounded on all objects at all time points. The operations in a TOQ-Net are vectorized, operating in parallel on all objects and times, sometimes expanding the dimensionality via outer products, and then re-projecting into smaller dimensions via max-pooling. This processing style is analogous to representing an entire graph using an adjacency matrix and using matrix operations to compute properties of the nodes or of the entire graph. In TOQ-Nets, the input to the network, as well as the feature output of intermediate layers, is represented as a tuple of three tensors.

Specifically, we use a vector of dimension D_0 to represent aspects of the state that are global and do not depend on any specific object at each time t . We use a matrix of shape $N \times D_1$ to represent the unary properties of each entity at time t , where N is the number of entities and D_1 is the hidden dimension size. Similarly, we use a tensor of shape $N \times N \times D_2$ to represent the relations between each pair of entities at time step t . As a concrete example, illustrated in Fig. 2, the number of entities N is the total number of players plus one (the ball). For each entity x and each time step, the inputs are their 3D position, type (ball or player) and team membership. The TOQ-Net outputs the action performed by the target player. Since there are only entity features, the input trajectory is encoded with a “unary” tensor of shape $T \times N \times D_1$, where T is the length of the trajectory. That is, there are no nullary

or binary inputs in this case.

2.2 Relational Reasoning Layers

Our Relational reasoning layers (RRLs) follow prior work on Neural Logic Machines [Dong *et al.*, 2019], illustrated in Fig. 3 (i). Consider a specific time step t . At each layer l , the input to a neural logic layer is a 3-tuple $(P_{l-1}, Q_{l-1}, R_{l-1})$, which corresponds to nullary, unary, and binary features respectively. Their shapes are D_0 , $N \times D_1$, and $N \times N \times D_2$. The output is another 3-tuple (P_l, Q_l, R_l) , given by

$$\begin{aligned}
 P_l &= \text{NN}_P(\text{Concat}[P_{l-1}; \max(Q_{l-1}, \text{dim} = 0)]), \\
 Q_l &= \text{NN}_Q(\text{Concat}[Q_{l-1}; \max(R_{l-1}, \text{dim} = 0); \\
 &\quad \max(R_{l-1}, \text{dim} = 1); \text{expand}(P_{l-1}, \text{dim} = 1)]), \\
 R_l &= \text{NN}_R(\text{Concat}[R_{l-1}; \text{expand}(Q_{l-1}, \text{dim} = 0); \\
 &\quad \text{expand}(Q_{l-1}, \text{dim} = 1)]).
 \end{aligned}$$

where NN_* are single fully-connected layers with sigmoid activations. For unary and binary features, NN_Q and NN_R are applied along the feature dimension. That is, we apply the same linear transformation to the unary features of all entities. A different linear transformation is applied to the binary features of all entity pairs. $\text{Concat}[\cdot; \cdot]$ is the concatenation operation, applied to the last dimension of the tensors (the feature dimension). \max , also called the “reduced” max operation, takes the maximum value along the given axis of a tensor. The expand operation, also called “broadcast,” will duplicate the input tensor N times and stack them together along the given axis. RRLs are applied identically to the input features at every time step t . That is, we use the same neural network weights in a RRL for all time steps in parallel.

RRLs are motivated by relational logic rules in a finite and fully grounded universe. The max reduction operations implement a differentiable version of an existential quantifier over the finite universe of individuals, given that the truth values of the propositions are represented as values in $(0.0, 1.0)$. Because preceding and subsequent RRLs can negate propositions as needed, we omit explicit implementation of finite-domain universal quantification, although it could be added by including analogous min reductions. Thus, as illustrated in Fig. 3 (i),

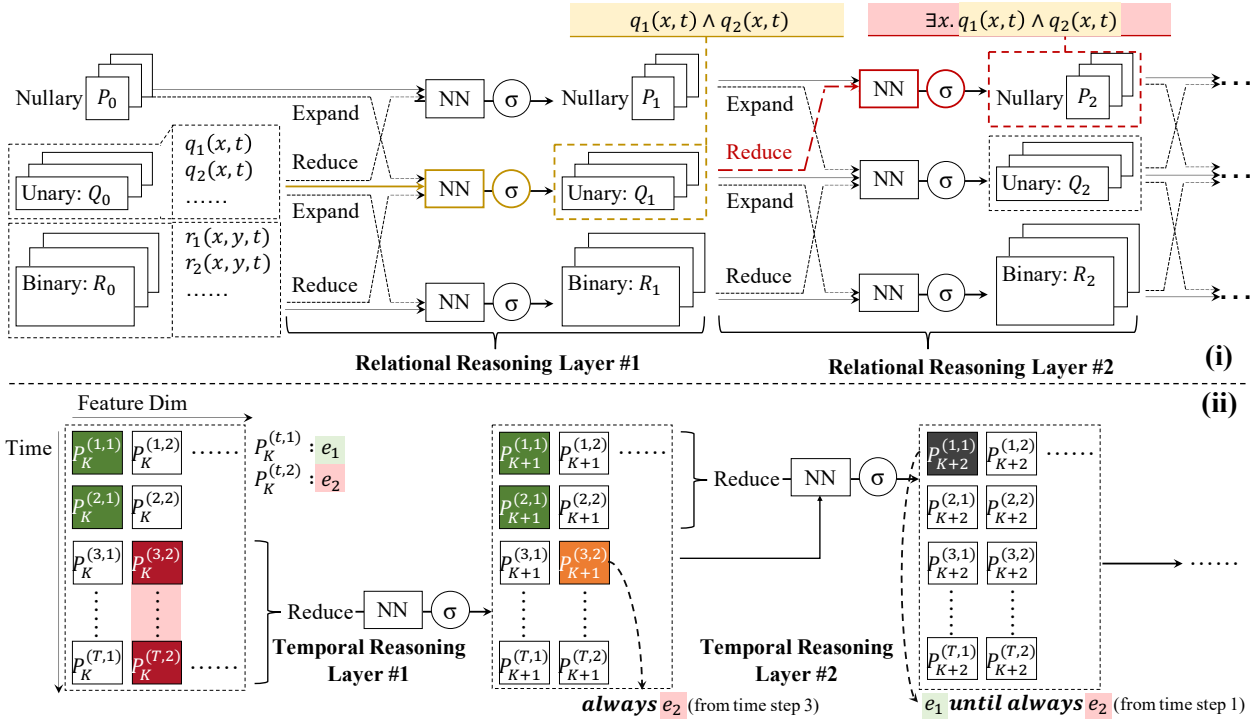


Figure 3: Illustration of (i) relational reasoning layers and (ii) temporal reasoning layers. We provide two illustrative running traces. (i) The first relational reasoning layer takes unary predicates q_1 and q_2 as input and its output Q_1 is able to represent $q_1 \wedge q_2$. The $\max(Q_1, \text{dim} = 0)$ in layer 2 can represent $\exists x. q_1(x, t) \wedge q_2(x, t)$. (ii) Assume P_K encodes the occurrence of events e_1 and e_2 at each time step. The first temporal reasoning layer can realize *always* e_2 with a temporal pooling from time step 3 to time step T . In the second temporal reasoning layer, the temporal pooling summarizes that e_1 holds true from time step 1 to 2. Thus, the NN should be able to realize e_1 *until* (*always* e_2).

given input features $q_1(x, t)$ and $q_2(x, t)$, we can realize the formula $\exists x. q_1(x, t) \wedge q_2(x, t)$ by stacking two such layers.

Throughout the paper we have been using only nullary, unary, and binary features, but the proposed framework itself can be easily extended to higher-order relational features. From a graph network [Bruna *et al.*, 2014; Kipf and Welling, 2017; Battaglia *et al.*, 2018] point of view, one can treat these features as the node and edge embeddings of a fully-connected graph and the relational reasoning layers as specialized graph neural network layers for realizing object quantifiers.

2.3 Temporal Reasoning Layers

Temporal reasoning layers (TRLs) perform *quantification* operations similar to relational reasoning layers, but along the *temporal* rather than the object dimension. The first TRL takes as input the summarized event representation produced by the K -th relational reasoning layer, $P_K^{(t)}$ for at all time steps t , as a matrix of shape $T \times D$. Each TRL is computed as

$$P_{K+l}^{(t)} = \max_{t' > t} \text{NN}_l \left(\text{Concat} \left[P_{K+l-1}^{(t')}; \max_{t \leq t'' < t'} P_{K+l-1}^{(t'')} \right] \right). \quad (1)$$

We will walk through the subexpressions of this formula.

1. P_{K+l-1} is the output tensor of the previous temporal reasoning layer, of shape $T \times C$, where T is the number of time steps, and C is the number of feature channels. Each entry in this tensor $P_{K+l-1}^{(t)}[i]$ can be interpreted as: event i happens at time t .

2. $\left(\max_{t \leq t'' < t'} P_{K+l-1}^{(t'')} \right)$, abbreviated as $Q_{K+l-1}^{(t, t')}$ in the following text, is a vector of shape C . Its entry $Q_{K+l-1}^{(t, t')}[i]$, where $t \leq t'$, represents the concept that event i happens some time between t and t' . **Importantly**, together with the preceding and subsequent neural network operations, which can realize negation operations, it also allows us to describe the event that i holds true for all time steps between t and t' .
3. NN_l is a fully connected neural network with sigmoidal activation, which gets uniformly applied to all time steps t and a future time step $t' > t$. Its input is composed of two parts: the events that happen at t' , i.e., $P_{K+l-1}^{(t')}$, and the events that happen between t and t' , summarized with temporal quantification operations, i.e., $\left(\max_{t \leq t'' < t'} P_{K+l-1}^{(t'')} \right)$.
4. The outer-most max pooling operation $\max_{t' > t}$ enumerates over all $t' > t$, and test whether the condition specified by NN_l holds for at least one such t' .

A special case is the first temporal reasoning layer. It takes P_K as its input, which is the output of last relational reasoning layer. Thus, the first temporal reasoning layer implements:

$$P_{K+1}^{(t)} = \text{NN}_1 \left(\max_{t \leq t'' < T} P_K^{(t'')} \right),$$

where T is the sequence length. Note that there is no enumer-

t	Input		1-st Layer				2-nd Layer
	$p(t)$	$q(t)$	Gp	Fp	Gq	Fq	$pU(Gq)$
1	T	T	F	T	F	T	T
2	T	F	F	T	F	T	T
3	F	T	F	F	T	T	F
4	F	T	F	F	T	T	F

Table 1: A running example of different temporal quantification formulas that TOQ-Nets can realize. For clarity, we use T and F for True/False. In the actual computation, they are “soft” Boolean values ranges in $[0, 1]$. **G** means *always*; **F** means *Eventually*; **U** means *until*.

Algorithm 1 An example temporal structure that the second temporal reasoning layer can recognize.

Input: $p(t)$, $q(t)$, $(Gp)(t)$, $(Fp)(t)$, $(Gq)(t)$, and $(Fq)(t)$
Output: $(pU(Gq))(t)$, which is *true* if p holds *true* from time step t until q becomes always *true*.

```

1: for  $t \leftarrow 1$  to  $T$  do
2:   for  $t' \leftarrow t + 1$  to  $T$  do
3:     if  $\forall t'' \in [t, t']$ .  $p(t'')$  and  $(Gq)(t')$  then
4:        $(pU(Gq))(t) \leftarrow true$ 
5:     end if
6:   end for
7: end for

```

ation for a future time step $t' > t$ involved. In addition, for all temporal layers, we add residual connections by concatenating their inputs with the outputs.

Next, let’s consider a running example illustrating how a TOQ-Net can recognize the event that: event p holds true until event q becomes always true. In LTL, this can be written as $pU(Gq)$. Using the plain first-order logic (FOL) language, we can describe it as: $\exists t. [\forall t'. (0 \leq t' < t) \implies p(t')] \wedge (\forall t'. (t' \geq t) \implies q(t'))$.

For simplicity, we consider a tensor representation for two events $p(t)$ and $q(t)$, where $p(t) = 1$ if it happens at time step t and $p(t) = 0$ otherwise. Given the input sequence of length 4 in Table 1 ($p(t)$ and $q(t)$), the first layer is capable of computing the following four properties for each time step t : Gp (*always p*), which is true if p holds true for all future time steps starting from t , Fp (*eventually p*), which is true if p is true for at least one time step starting from t , and similarly, Gq (*always q*) and Fq (*eventually q*). Overall, together with residual connections, the first layer can recognize six useful events: $p(t)$, $q(t)$ (from residual connection), Gp , Fp , Gq , and Fq (by temporal quantification, i.e. pooling operations along the temporal dimension).

The second layer can realize the computation depicted in Algorithm 1. For every time step t , it enumerates all $t' > t$ and computes the output based on

- the events at t' (represented as $P_{K+l-1}^{(t')}$ in Equation 1, concretely the $(Gq)(t')$ in the Algorithm 1 example), and
- the state of another event between t and t' (represented as $(\max_{t \leq t'' < t'} P_{K+l-1}^{(t'')})$ in Equation 1, concretely the $\forall t'' \in [t, t']$. $p(t'')$ in the Algorithm 1 example).

From the perspective of First-Order Linear Temporal Logic (FO-LTL), stacking multiple temporal reasoning layers enables us to realize FO-LTL formulas such as:

$$p_1 U p_2 U p_3 U \dots U p_k,$$

which is interpreted as p_1 holds true until p_2 becomes true and p_2 holds true from that until p_3 becomes true \dots , and

$$Fp_1 XF p_2 XF p_3 XF \dots XF p_k,$$

which is interpreted as p_1 eventually becomes true and after that p_2 eventually becomes true and after that \dots , and in addition, formulas with interleaved until and eventually quantifiers. Here, XF is a composition of the neXt operator and the FInally operator in LTL. Meanwhile, as described so far, TOQ-Nets can only nest object quantification inside temporal quantification, so it can represent *always* $\exists x. q_1(x) \wedge q_2(x)$, but not $\exists x$. *always* $q_1(x) \wedge q_2(x)$. This can be solved by interleaving relational and temporal reasoning layers.

It is important to notice that, by using object and temporal pooling operations together with trainable neural networks to realize logic formulas with object and temporal quantifiers, the idea itself generalizes to a broader set of FO-LTL formulas. We design TOQ-Nets to model only a subset of FO-LTL formulas, because they can be computed efficiently (with only $O(T^2)$ space) and they are expressive enough for the type of data we are trying to model.

3 Experiments

We compare our model with other approaches to object-centric temporal event detection in this section, and include an application of TOQ-Nets to concept learning over robot object-manipulation trajectories in the supplementary material. The setups and metrics focus on data efficiency and generalization.

3.1 Baseline Approaches

We compare TOQ-Nets against five baselines. The first two are spatial-temporal graph convolutional neural networks [STGAN; Yan *et al.* 2018] and its variant STGCN-MAX, which models entity relationships with graph neural networks and models temporal structure with temporal-domain convolutions. The third is STGCN-LSTM, which uses STGCN layers for entity relations but LSTM [Hochreiter and Schmidhuber, 1997] for temporal structures. The last two baselines are based on space-time graphs: Space-Time Graph [Wang and Gupta, 2018] and Non-Local networks [Wang *et al.*, 2018]. We provide details about our implementation and how we choose the model configurations in the supplementary material.

3.2 Trajectory-Based Soccer Event Detection

We start our evaluation with an event-detection task in soccer games. The task is to recognize the action performed by a specific player at specific time step in a soccer game trajectory.

Dataset and setup. We collect training and evaluation datasets based on the gfootball simulator*, which provides a physics-based 3D football simulation. It also provides AI

*<https://research-football.dev/>

Model	Reg.	Few-Shot	Full
STGCN	73.2 \pm 1.6	26.0 \pm 5.7	62.8 \pm 0.6
STGCN-MAX	73.6 \pm 1.5	28.6 \pm 5.0	63.6 \pm 0.7
STGCN-LSTM	72.7 \pm 1.4	23.8 \pm 5.9	61.9 \pm 0.6
Space-Time	74.8 \pm 1.5	31.7 \pm 6.1	65.2 \pm 0.6
Non-Local	76.5 \pm 2.4	45.0 \pm 6.3	69.5 \pm 2.4
TOQ-Net (ours)	87.7 \pm 1.3	52.2 \pm 6.3	79.8 \pm 0.8

Table 2: Results on the soccer event dataset. Different columns correspond to different action sets (the regular, few-shot, and full action sets). The performance is measured by per-action (macro) accuracy, averaged over nine few-shot splits. The \pm values indicate standard errors. TOQ-Net significantly outperforms all baseline methods on the few-shot action set.

agents that can be used to generate random plays. The simulator provides the 3D coordinates of the ball and the players as well as the action each player is performing at each time step. There are in total 13 actions defined in the simulator, including *movement*, *ball_control*, *trap*, *short_pass*, *long_pass*, *high_pass*, *header*, *shot*, *deflect*, *catch*, *interfere*, *trip* and *sliding*. We exclude *header* and *catch* actions, as they never appear in AI games. We also exclude *ball_control* and *movement*, since they just mean the agent is moving (with or without the ball). Thus, in total, we have nine action categories. We run the simulator with AI-controlled players to generate plays, and formulate the task as classifying the action (9-way classification) of a specific player at a specific time step given a temporal context (25 frames). For each action, we have generated 5,000 videos, except for *sliding*, for which we generated 4,000 videos because it is rare in the AI games. Among the generated examples, 62% (2,462 or 3,077) are used for training, 15% are used for validation, and 23% are used for testing. Each trajectory is an 8-fps replay clip that contains 17 frames (about two seconds). There is a single “target” player in each trajectory. The action label of the trajectory is the action performed by this target player at frame #9. We randomly split all actions into two categories: seven “regular” actions, for which all game plays are available, and two “few-shot” actions, for which only 50 clips are available during training.

Input features. Each trajectory is represented with 7 time-varying unary predicates, including the 3D coordinate of each player and the ball and four extra predicates defining the type of each entity x : $\text{IsBall}(x)$, $\text{IsTargetPlayer}(x)$, $\text{SameTeam}(x)$, $\text{OpponentTeam}(x)$, where $\text{SameTeam}(x)$ and $\text{OpponentTeam}(x)$ indicates whether x is of the same team as the target player. We also add a temporal indicator function which is a Gaussian function centered at frame of interest with variance $\sigma^2 = 25$.

Results. Table 2 shows the result. Our model significantly outperforms all the baselines in all three action settings, suggesting that our model is able to discover a set of useful features at both input and intermediate levels and use them to compose new action classifiers from only a few examples.

Generalization to more players. Due to their object-centric design, TOQ-Nets can generalize to soccer games with a varying number of agents. After training on 6v6 soccer games (i.e., 6 players on each team), we evaluate the perfor-

mance of different models on games with different numbers of players: 3v3, 4v4, 8v8, and 11v11. For each action we have generated, on average, 1,500 examples for testing. Table 3 summarizes the result and the full results are provided in the supplementary material. Comparing the columns highlighted in yellow, we notice a significant performance drop for all baselines while TOQ-Net performs the best. By visualizing data and predictions, we found that misclassifications of instances of *shot* as *short pass* contribute most to the performance degradation of our model when we have more players. Specifically, the recall of *shot* drops from 97% to 60%. In soccer plays with many agents, a shot is usually unsuccessful and a player from another team steals the ball in the end. In such scenarios, TOQ-Net tends to misclassify such trajectories as a *short pass*. Ideally, this issue should be addressed by understanding actions based on agents’ goals instead of the actual outcome [Intille and Bobick, 2001]. We leave this extension as a future direction.

Generalization to temporally warped trajectories. Another crucial property of TOQ-Nets is to recognize actions based on their sequential order in the input trajectory, instead of binding features to specific time steps in the trajectory. To show this, we test the performance of different models on time warped trajectories. Each test trajectory has a length of 25, and each trajectory is labeled by the action performed by the target player at any time step between the 6-th and 19-th frame. We ensure that the target player performs only one action during the entire input trajectory. Thus, the label is unambiguous. The results are shown in Table 3. Specifically, our test set consists of 25-frame trajectories, and the action may occur at anytime between the 6th and the 19th frame. By comparing rows with and without time warping, we notice a 60% performance drop for STGCN, STGCN-MAX, and STGCN-LSTM. In contrast, TOQ-Nets still have reasonable performance. Note that Space-Time and Non-Local model have almost no performance drop against time warping because they are completely agnostic to temporal ordering.

3.3 Extension to Real-World Datasets

The proposed TOQ-Net can also be extended to other real-world datasets. These examples further illustrate the robustness of TOQ-Net to temporal variations in activities.

Toyota Smarthome

Toyota Smarthome [Das *et al.*, 2019] is a dataset that contains videos of humans performing everyday activities such as “walk”, “take pills”, and “use laptop”. It also comes with 3D-skeleton detections. There are around 16.1k videos in the dataset, and 19 activity classes. The videos’ length varies between a few seconds to 3 minutes. We subsample 30 frames for each video. We split frames into training (9.9k), validation (2.5k), and testing (3.6k). We treat human joints as entities. The input is then the position of joints, the velocity of joints, limb lengths, and joint angles. We evaluated our model and STGCN on a 19-way classification task. We also test model performance on time-warped sequences by accelerating the trajectories by two times.

Our model achieves a comparable accuracy to STGCN on the standard classification task: (42.0% vs. 43.0%). Impor-

Model	3v3	4v4	6v6	6v6 (Time Warp)	8v8	11v11
STGCN	40.7 \pm 1.0 (-40.4%)	63.2 \pm 4.9 (-7.4%)	68.2 \pm 2.8	52.8 \pm 7.0 (-22.6%)	55.4 \pm 3.3 (-18.8%)	44.4 \pm 2.1 (-34.9%)
STGCN-MAX	47.4 \pm 3.2 (-33.7%)	68.8 \pm 2.0 (-3.8%)	71.5 \pm 1.9	56.5 \pm 4.5 (-21.0%)	59.1 \pm 0.7 (-17.3%)	45.6 \pm 2.5 (-36.2%)
STGCN-LSTM	39.7 \pm 1.1 (-43.1%)	60.4 \pm 0.2 (-13.5%)	69.8 \pm 0.1	30.6 \pm 0.6 (-56.1%)	55.8 \pm 2.0 (-20.0%)	44.1 \pm 0.7 (-36.8%)
Space-Time	29.0 \pm 1.6 (-60.4%)	53.5 \pm 3.2 (-27.0%)	73.3 \pm 0.3	70.7 \pm 0.3 (-3.5%)	33.9 \pm 2.8 (-53.7%)	15.2 \pm 1.8 (-79.3%)
Non-Local	45.9 \pm 5.1 (-41.2%)	70.7 \pm 5.3 (-9.5%)	78.1 \pm 5.8	77.7 \pm 5.0 (-0.5%)	58.5 \pm 10.8 (-25.1%)	41.8 \pm 13.6 (-46.5%)
TOQ-Net	77.4\pm3.5 (-12.4%)	88.3\pm0.7 (-0.0%)	88.4\pm0.6	86.9\pm0.4 (-1.7%)	81.3\pm1.7 (-8.0%)	77.1\pm1.7 (-12.8%)

Table 3: Results on generalization to scenarios with more agents and temporally warped trajectories on the soccer event dataset. All models are trained only on 6v6 games. The standard errors indicated by the \pm signs are computed with three random seeds.

tantly, on the generalization test to time-warped sequences, our model has only a 0.8% performance drop (41.2%), while STGCN drops 10.7% (32.3%). This indicates that the temporal structures learned by TOQ-Net improve model generalization to varying time courses.

Volleyball Activity

The volleyball dataset [Ibrahim *et al.*, 2016] contains 4830 video clips collected from 55 youtube volleyball videos. They are labeled with 8 group activities (e.g. “left spike” and “right pass”). Each video contains 20 frames with the labeled group activity performed at the 10-th frame. The dataset also includes annotations for players, including the bounding box, the indicator of whether the player is involved in the group activity, and the individual action such as “setting”, “digging”, and “spiking”. We use the manual annotations (processed by an MLP) as the input features. We train models to classify the video into one of the eight group activities, following the original split, i.e., 24, 15, and 16 of 55 videos are used for training, validation, and testing.

On the standard classification task, TOQ-Net achieves a comparable performance with STGCN (73.3% vs. 73.6%). When we perform time warping on the input sequences, STGCN’s performance drops by more than 25.0% (39.5% on temporally shifted trajectories and 48.6% on $2\times$ quick motion trajectories), while our model drops only 3% (70.3% on temporally shifted trajectories and 70.7% on quick motion trajectories). This again shows the generalization ability of TOQ-Net w.r.t. varying time courses, and the robustness of learned temporal structures.

4 Related Work

Action concept representations and learning. First-order and linear temporal logics [LTL; Pnueli 1977] have been used for analyzing sporting events [Intille and Bobick, 1999; Intille and Bobick, 2001] and activities of daily living [Tran and Davis, 2008; Brendel *et al.*, 2011] in logic-based reasoning frameworks. However, these frameworks require extra knowledge to annotate relationships between low-level, primitive actions and complex ones, or performing search in a large combinatorial space for candidate temporal logic rules [Penning *et al.*, 2011; Lamb *et al.*, 2007]. By contrast, TOQ-Nets enable end-to-end learning of complex action descriptions from sensory input with only high-level action-class labels.

Temporal and relational reasoning. This paper is also related to work on using data-driven models for modeling relational and temporal structure, such as LTL [Neider and

Gavran, 2018; Camacho *et al.*, 2018; Chou *et al.*, 2020], Logical Neural Networks [Riegel *et al.*, 2020], ADL description languages [Intille and Bobick, 1999], and hidden Markov models [Tang *et al.*, 2012]. These models need human-annotated action descriptions and symbolic state variables (e.g., *pick up x* means a state transition from not *holding x* to *holding x*), and dedicated inference algorithms such as graph structure learning. In contrast, TOQ-Nets have an end-to-end design, and can be integrated with other neural networks. People have also used structural representations to model object-centric temporal concepts with graph neural networks [GNNs; Yan *et al.* 2018], recurrent neural networks [RNNs; Ibrahim *et al.* 2016], and integrated GNN-RNN architectures [Deng *et al.*, 2016; Qi *et al.*, 2018]. TOQ-Nets use a similar relational representation, but different models for temporal structures.

5 Conclusion and Discussion

The design of TOQ-Nets suggests multiple research directions. For example, the generalization of the acquired action concepts to novel object kinds, such as from *opening fridges* to *opening bottles*, needs further exploration. Meanwhile, TOQ-Nets are based on physical properties, e.g. 6D poses. Incorporating representations of mental variables such as goals, intentions, and beliefs can aid in action and event recognition [Baker *et al.*, 2017; Zacks *et al.*, 2001; Vallacher and Wegner, 1987].

In summary, we have presented TOQ-Nets, a neuro-symbolic architecture for learning to describe complex state sequences with quantification over both entities and time. TOQ-Nets use tensors to represent the time-varying properties and relations of different entities, and use tensor pooling operations over different dimensions to realize temporal and object quantifiers. TOQ-Nets generalize well to scenarios with varying numbers of entities and time courses.

Acknowledgements

We thank Tomás Lozano-Pérez for helpful discussions and suggestions. We gratefully acknowledge support from the Center for Brains, Minds and Machines (NSF STC award CCF1231216), ONR MURI N00014-16-1-2007, NSF grant 1723381, AFOSR grant FA9550-17-1-0165, ONR grant N00014-18-1-2847, the Honda Research Institute, MIT-IBM Watson Lab, SUTD Temasek Laboratories, the Samsung Global Research Outreach (GRO) Program, and the Stanford Institute for Human-Centered AI (HAI). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our sponsors.

References

- [Baker *et al.*, 2017] Chris L Baker, Julian Jara-Ettinger, Rebecca Saxe, and Joshua B Tenenbaum. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4):1–10, 2017.
- [Battaglia *et al.*, 2018] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv:1806.01261*, 2018.
- [Brendel *et al.*, 2011] William Brendel, Alan Fern, and Sinisa Todorovic. Probabilistic event logic for interval-based event recognition. In *CVPR*, 2011.
- [Bruna *et al.*, 2014] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. In *ICLR*, 2014.
- [Camacho *et al.*, 2018] Alberto Camacho, Jorge Baier, Christian Muise, and Sheila McIlraith. Finite ltl synthesis as planning. In *ICASP*, 2018.
- [Chou *et al.*, 2020] Glen Chou, Necmiye Ozay, and Dmitry Berenson. Explaining multi-stage tasks by learning temporal logic formulas from suboptimal demonstrations. In *RSS*, 2020.
- [Das *et al.*, 2019] Srijan Das, Rui Dai, Michal Koperski, Luca Minciullo, Lorenzo Garattoni, Francois Bremond, and Gianpiero Francesca. Toyota smarhome: Real-world activities of daily living. In *ICCV*, 2019.
- [Deng *et al.*, 2016] Zhiwei Deng, Arash Vahdat, Hexiang Hu, and Greg Mori. Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. In *CVPR*, 2016.
- [Dong *et al.*, 2019] Honghua Dong, Jiayuan Mao, Tian Lin, Chong Wang, Lihong Li, and Denny Zhou. Neural logic machines. In *ICLR*, 2019.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.
- [Ibrahim *et al.*, 2016] Mostafa S. Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori. A hierarchical deep temporal model for group activity recognition. In *CVPR*, 2016.
- [Intille and Bobick, 1999] Stephen S Intille and Aaron F Bobick. A framework for recognizing multi-agent action from visual evidence. In *AAAI*, 1999.
- [Intille and Bobick, 2001] Stephen S Intille and Aaron F Bobick. Recognizing planned, multiperson action. *CVIU*, 81(3):414–445, 2001.
- [Kipf and Welling, 2017] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- [Lamb *et al.*, 2007] Luis C. Lamb, Rafael V. Borges, and A. d’Avila Garcez. A connectionist cognitive model for temporal synchronisation and learning. In *AAAI*, 2007.
- [Neider and Gavran, 2018] Daniel Neider and Ivan Gavran. Learning linear temporal properties. In *FMCAD*, 2018.
- [Penning *et al.*, 2011] Leo de Penning, A. d’Avila Garcez, Luís C. Lamb, and John-Jules C. Meyer. A neural-symbolic cognitive agent for online learning and reasoning. In *IJCAI*, 2011.
- [Pnueli, 1977] Amir Pnueli. The temporal logic of programs. In *FOCS*, 1977.
- [Qi *et al.*, 2018] Mengshi Qi, Jie Qin, Annan Li, Yunhong Wang, Jiebo Luo, and Luc Van Gool. stagnet: An attentive semantic rnn for group activity recognition. In *ECCV*, 2018.
- [Riegel *et al.*, 2020] Ryan Riegel, Alexander Gray, Francois Luus, Naweed Khan, Ndivhuwo Makondo, Ismail Yunus Akhalwaya, Haifeng Qian, Ronald Fagin, Francisco Barahona, Udit Sharma, Shajith Ikbal, Hima Karanam, Sumit Neelam, Ankita Likhyan, and Santosh Srivastava. Logical neural networks. In *NeurIPS*, 2020.
- [Stränger and Hommel, 1996] Jürgen Stränger and Bernhard Hommel. The perception of action and movement. In *Handbook of perception and action*, volume 1, pages 397–451. Elsevier, 1996.
- [Tang *et al.*, 2012] Kevin Tang, Li Fei-Fei, and Daphne Koller. Learning latent temporal structure for complex event detection. In *CVPR*, 2012.
- [Tran and Davis, 2008] Son D Tran and Larry S Davis. Event modeling and recognition using markov logic networks. In *ECCV*, 2008.
- [Vallacher and Wegner, 1987] Robin R Vallacher and Daniel M Wegner. What do people think they’re doing? action identification and human behavior. *Psychol. Rev.*, 94(1):3, 1987.
- [Wang and Gupta, 2018] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [Yan *et al.*, 2018] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018.
- [Zacks *et al.*, 2001] Jeffrey M Zacks, Barbara Tversky, and Gowri Iyer. Perceiving, remembering, and communicating structure in events. *J. Exp. Psychol. Gen.*, 130(1):29, 2001.
- [Zacks *et al.*, 2007] Jeffrey M Zacks, Nicole K Speer, Khena M Swallow, Todd S Braver, and Jeremy R Reynolds. Event perception: a mind-brain perspective. *Psychol. Bull.*, 133(2):273, 2007.

Supplementary Material for Temporal and Object Quantification Networks

Jiayuan Mao^{1*}, Zhezheng Luo^{1*}, Chuang Gan², Joshua B. Tenenbaum¹, Jiajun Wu³,
Leslie Pack Kaelbling¹ and Tomer D. Ullman⁴

¹Massachusetts Institute of Technology

²MIT-IBM Watson AI Lab

³Stanford University

⁴Harvard University

We start from introducing a new dataset we collected for model evaluation, namely the RL Bench dataset (Appendix A). We then provide implementation details of different models in Appendix B. Finally, in Appendix C, we present additional ablation studies and visualizations.

A Manipulation Concept Learning from 6D Poses

Structural action representations can also be usefully applied to other domains. Here we show the result in a robotic environment, where the goal is to classify the action performed by a robotic arm.

Dataset and setup. We generated a dataset based on the RL Bench simulator [James *et al.*, 2020], which contains a set of robotic object-manipulation actions in a tabletop environment. We use 24 actions from the dataset, including *CloseBox*, *CloseDrawer*, *CloseFridge*, *CloseGrill*, *CloseJar*, *CloseLaptopLid*, *CloseMicrowave*, *GetIceFromFridge*, *OpenBox*, *OpenFridge*, *OpenMicrowave*, *OpenWineBottle*, *PickUpCup*, *PressSwitch*, *PushButtons*, *PutGroceriesInCupboard*, *PutItemInDrawer*, *PutRubbishInBin*, *PutTrayInOven*, *ScoopWithSpatula*, *SetTheTable*, *SlideCabinetOpenAndPlaceCups*, *TakePlateOffColoredDishRack*, and *TakeToiletRollOffStand*. We randomly initialize the position and orientation of different objects in the scene and use the built-in motion planner to generate robot arm trajectories. Depending on the task and the initial poses of different objects, trajectories may have different lengths, varying from 30 to 1,000. Most of the trajectories have ~ 50 frames.

For each action category, we have generated 100 trajectories. Among the generated examples, 60% of the examples are used for training, 15% are used for validation, and the other 25% are used for testing.

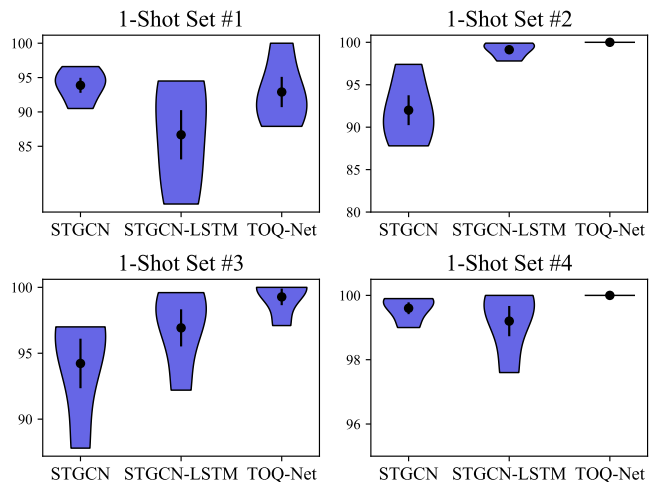
Input. Each trajectory contains the poses of different objects, at each time step. Each object except for the robot arm is represented as 13 unary predicates, including the 3D position, the 3D orientation represented as quaternions (4D vector), and the 3D bounding box (axis-aligned, 6D vector). The robot arm is represented as 8 unary predicates, including the 3D position, the 3D orientation, and a binary predicate that indicates whether the gripper is open or closed. Note that deformable objects are split into several pieces so that the

learner can infer their state (open or closed). For example, the lid of the box is a separate object. We also input the 6D pose and the state of the robot gripper. During training, we used on average 65 robot trajectories for regular actions.

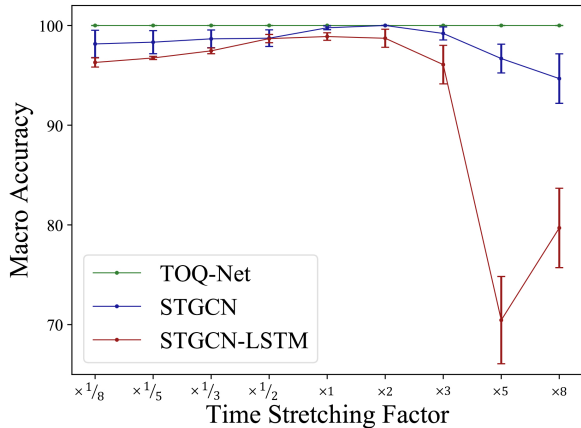
Input features. All methods take the raw trajectories as input. In TOQ-Nets, we use an extra single linear layer with sigmoid activation as the input feature extractor. To evaluate temporal modeling with recurrent neural networks, STGCN-LSTM uses kernel size 1 in its (STGCN) temporal convolution layers. Ablation studies can be found in Appendix B.

Model	Reg.	1-Shot	Full
STGCN	99.89 \pm 0.05	94.92 \pm 1.03	98.79 \pm 0.23
STGCN-LSTM	99.92 \pm 0.03	95.48 \pm 1.67	98.86 \pm 0.43
TOQ-Net	99.96 \pm 0.02	98.04 \pm 0.97	99.48 \pm 0.24

Table 1: Few-shot learning on the RL Bench dataset, measured by per-action (macro) accuracy and averaged of four 1-shot splits and four random seeds per split. The \pm values indicate standard errors. On the right we show the sampled performance of different models on each individual 1-shot split.



Results. We start with evaluating different models on the standard 24-way classification task. We summarize results



1

Figure 1: Comparing different models with different time stretching factors on the RL Bench dataset.

in Fig. 1. It also illustrates each model’s performance under different time stretching factors of the input trajectory, from $1/8\times$ to $8\times$. Similar to the soccer dataset, the TOQ-Net outperforms both convolutional and recurrent models across all time stretching factors.

Generalization to new action concepts. We also evaluate how well different models can generalize the learned knowledge about opening and closing other objects to novel objects. Specifically, we hold out 50% of all *Open X* actions and 50% of the *Close X* actions to form the 1-shot action learning set. For example, the learner may only see a single instance of the opening of a box during training. Actions with more examples are “regular.” All models are tested on a 3-way classification task: *Open-X*, *Close-X*, and *Other*. Results are summarized in Table 1. There are several held-out sets that are noticeably more “difficult” than others, such as the Set #1 in Table 1. Our visualization shows that actions *close jar* and *close drawer* are harder generalization targets, compared to other actions. This is possibly because the motion of *close jar* involves some rotation, and the motion of *close drawer* looks more like pushing than other *close-X* actions.

B Implementation Details

In this section, we present the implementation detail of our model, the TOQ-Nets, and five baselines (STGCN, STGCN-LSTM, STGCN-MAX, Space-Time and Non-local), including the model architecture, input features, and the training methods.

B.1 TOQ-Nets

In the soccer event detection task, we use three object quantification layers and three temporal quantification layers. Each object quantification layer has a hidden dimension of 16 (i.e., all nullary, unary, and binary tensors have the same hidden dimension of 16). Each temporal quantification layer has a hidden dimension of 48.

In the manipulation concept learning task, we use three object quantification layers and three temporal quantification

layers. Each object quantification layer has a hidden dimension of 48. Each temporal quantification layer has a hidden dimension of 144.

Input feature extractor. We use the following physics-based input feature extractor for soccer event detection. Given the trajectories of all entities, including all players and the ball, we first compute the following physical properties: ground speed (i.e., the velocity on the xy -plane), vertical speed (i.e., the velocity along the z direction), and acceleration. We also compute the distance between each pair of entities.

After extracting the physical properties, we use the following feature templates to generate the input features. For each physical property, X , we first normalize X across all training trajectories, and then we create $c = 5$ features:

$$\sigma\left(\frac{X - \theta_i}{\beta}\right),$$

where σ is the sigmoid function, β is a scalar hyperparameter, $\theta_i, i = 1, 2, 3, 4, 5$ are trainable parameters. These feature templates can be interpreted as differentiable implementations of $\mathbf{1}[X > \theta_i]$, where $\mathbf{1}[\cdot]$ is the indicator function. During training, we make β decay exponentially from 1 to 0.001.

In the manipulation concept learning task, we use a single fully-connected layer with sigmoid activation as the feature extractor. The hidden dimension of the layer is 64.

B.2 STGCN

We use the same architecture as Yan et al. [Yan et al., 2018]. Table 5 summarizes the hidden dimensions and the kernel sizes.

The STGCN model’s output is a tensor of size $(T/4) \times N \times 256$, where T is the length of the input trajectory, and N is the number of entities. Following Yan et al. [Yan et al., 2018], we perform an average pooling over the temporal and the entity dimension and get a feature vector of size 256. We apply a softmax layer on top of the latent feature to classify the action.

Input feature extractor. In the soccer event detection task, we do not input the distance between each pair of entities as we do for the TOQ-Nets, as the STGCN model does not support binary predicate inputs (see STGCN-2D for an ablation study). In the manipulation concept learning task, STGCN uses the same input format as the TOQ-Nets. Specifically, we represent the state of each entity with 13 unary predicates, including 3D position, 3D orientation, and bounding boxes.

B.3 STGCN-MAX

The STGCN-MAX model is a variation of STGCN. During graph propagation, for each node, instead of taking average of all propagated information from all neighbors, we take their maximum value over all feature dimensions, and update the hidden state of each node with this. This baseline replicates a very similar propagation rule as our relational reasoning layer. However, it still uses temporal convolutions to model temporal structures.

B.4 STGCN-2D

The original STGCN does not support input features of higher dimensions, so we extend STGCN to STGCN-2D to add bi-

	Reg.	New	All
TOQ-Net	87.7	52.2	79.8
STGCN-2D	84.1	46.9	75.8

Table 2: Ablation study of the STGCN-2D model on the soccer event dataset, evaluated by its few-shot learning performance.

	9-way	3v3	11v11	6v6(Time warp)
TOQ-Net	88.4	77.4	77.1	86.9
STGCN-2D	84.5	17.5	16.6	39.7

Table 3: Ablation study of the STGCN-2D model on the soccer event dataset, evaluated by its generalization to different number of players and time-warped trajectories.

nary inputs such as the distance. This extension slightly improves the model performance on the standard 9-way classification test. It also helps in the few-shot setting. Specifically, on the few-shot learning setting:

Meanwhile, adding extra binary inputs does not improve the generalization of the network to games with a different number of agents or to time-warped sequences. This following table extends the Table 2 in the main text. Specifically, we train STGCN-2D on 6v6 soccer games and test it on scenarios with a different number of agents (3v3 and 11v11) and also temporally warped trajectories. Our model shows a significant advantage.

B.5 STGCN-LSTM

The STGCN-LSTM model first encodes the input trajectory with an STGCN module. The output of the STGCN module is a tensor of shape $T \times N \times 256$, where T is the length of the input trajectory, and N is the number of entities. We perform an average pooling over the entity dimension and get a tensor of shape $T \times 256$. Next, we encode this feature with a 2-layer bidirectional LSTM module. The hidden dimension for the LSTM module is 256. We use a single softmax layer on top of the LSTM module to classify the action.

STGCN-LSTM uses the same architecture as STGCN for both tasks, except for the kernel size and the stride, because our goal is to evaluate the performance of recurrent neural networks (the LSTM model) on modeling temporal structures. In the soccer event detection task, we use kernel size 3 and stride 1 so that the STGCN module will have the representational power to encode local physical properties. In the manipulation concept learning task, we use a temporal kernel size of 1 and stride 1. Meanwhile, Table 4 shows the ablation study of different kernel sizes, evaluated by the few-shot learning performance on the RL Bench dataset. Different kernel sizes (1 and 3) have a similar performance.

B.6 Space-Time Region Graph Networks

We use the same architecture for each Space-Time layer as Wang et al. proposed in Space-Time Graph [Wang and Gupta, 2018]. Table 6 summarises the detailed parameters of all layers.

Model	Reg.	1-Shot	Full
STGCN	99.89 \pm 0.05	94.92 \pm 1.03	98.79 \pm 0.23
STGCN-LSTM (kernel size 3)	99.89 \pm 0.05	96.16 \pm 1.39	99.04 \pm 0.31
STGCN-LSTM (kernel size 1)	99.92 \pm 0.03	95.48 \pm 1.67	98.86 \pm 0.43
TOQ-Net	99.96 \pm 0.02	98.04 \pm 0.97	99.48 \pm 0.24

Table 4: Ablation study of different kernel sizes for the STGCN-LSTM model on the RL Bench dataset, measured by per-action (macro) accuracy and averaged of four 1-shot splits and four random seeds per split. The \pm values indicate standard errors.

B.7 Non-Local Neural Networks

We use the same architecture for each Non-local layer as Wang et al. proposed in Non-Local networks [Wang et al., 2018]. The original non-local network works on pixels. Here instead, we apply it to the space-time graph constructed in the same way as Wang and Gupta 2018. The major difference between the Space-Time Graph and the Non-local Network is that Space-Time Graph uses graph convolutions while Non-local uses transformer-like attentions over all nodes in the graph. Table 7 summarises the detailed parameters of all layers.

Soccer event detection. We use the identical training strategy for our model and both baselines. Each training batch contains 128 examples, which are sampled as following: we first uniformly sample an action category, and then uniformly sample a trajectory labeled as this action category.

In all the soccer event tasks, we optimized our model using Adam [Kingma and Ba, 2015], with the learning rate initialized to $\eta = 0.002$. The learning rate is fixed for the first 50 epochs. After that, we decrease the learning rate η by a factor of 0.9 if the best validation loss has not been updated during the last 6 epochs.

All models are trained for 200 epochs. In the task of generalization to different input scales, we first train different models on the original dataset for 200 epochs (\sim 40,000 iterations) and finetune them on the new dataset with few samples for 10,000 iterations.

Manipulation concept learning. We optimized different models using Adam with the initial learning rate $\eta = 0.003$. We applied the same learning rate schedule as in the soccer event detection task. The batch size for the manipulation concept learning task is 32 for models with STGCN backbone and is 4 for TOQ-Nets. All models are trained for 150 epochs.

C Additional Ablation Study

C.1 Network Architecture

In general, TOQ-Net has a smaller number of weights than baselines. We add additional comparisons to models of different #params. Specifically, model_S is a smaller version of the model and model_T is the tiny version of the model whose #params are about equal to or smaller than #params of TOQ-Net. To obtain smaller models we typically reduce the number

Model	Input Dim.	Output Dim.	Kernel Size	Stride	Dropout	Residual
STGCN	-	64	7	1	0	False
	64	64	7	1	0.5	True
	64	64	7	1	0.5	True
	64	64	7	1	0.5	True
	64	128	7	2	0.5	True
	128	128	7	1	0.5	True
	128	128	7	1	0.5	True
	128	256	7	2	0.5	True
	256	256	7	1	0.5	True
STGCN _S	-	16	7	1	0	False
	16	16	7	1	0.5	True
	16	32	7	2	0.5	True
	32	32	7	1	0.5	True
	32	64	7	2	0.5	True
STGCN _T	-	8	7	1	0	False
	8	12	7	2	0.5	True
	12	64	7	1	0	True

Table 5: The STGCN architecture used in the paper.

of layers and the hidden state dimensions. On the few-shot task of the soccer event dataset, we test all baselines with all combination of features and capacities. Table 8 summarizes the results. In general, the “small” variations are the best for all models (except for STGCN-LSTM). Across all experiments presented in the main paper, we use the best architecture in Table 8.

C.2 Generalization to different input scales

Intuitively, after learning how to detect events in a soccer game, the learner should also generalize concepts to analogs with the same basic structure but enacted on proportionately larger or smaller spatial or temporal scales—as in the soccer-inspired game futsal. To evaluate generalization to different input scales, we first train all models on the original dataset with 9-way classification supervision. After that, we finetune each model on a new dataset where all input values are doubled (i.e., the court is now double the size, and the players now move at four times the speed), but with only a small number of examples. Thus, time flows twice fast, and the player moves four times the speed compared with the original dataset. Results are summarized in Fig. 2. While all baselines can benefit from the pre-training on the original dataset, our TOQ-Net outperforms all baselines by a significant margin when the finetuning dataset is small.

C.3 Full Results for the Time-Warping Generalization

In the paper we have evaluated the generalization performance of different models to time-warped sequences but with the same number of players. Table 9 shows the results on time-warped sequences with different number of agents.

C.4 #Layers of TOQ-Net

We study the TOQ-Net performance over varying number of relation reasoning layers and temporal reasoning layers on the 9-way classification task of the soccer event dataset (See figure Fig. 3), and decide to use the combination of 3 relational layers and 4 temporal layers. Both relational layers and temporal layers play important roles in the performance. As a generalization test, we also visualize the accuracy on temporally warped trajectories. We can see that temporal reasoning layers are important for the robustness against temporal warping.

C.5 Interpretability

Though interpretability is not the main focus of our paper, in Fig. 4, we show how TOQ-Net composes low-level concepts into high-level concepts: from inputs, to intermediate features at different temporal layers, and the final output labels. See the figure caption for detailed analysis.

Model	Input Dim.	Output Dim.	Stride	Residual
Space-Time	-	64	1	False
	64	64	1	True
	64	64	1	True
	64	64	1	True
	64	128	2	True
	128	128	1	True
	128	128	1	True
	128	256	2	True
	256	256	1	True
Space-Time _S	-	16	1	False
	16	16	1	True
	16	32	2	True
	32	64	2	True
	64	128	1	True
Space-Time _T	-	32	1	False
	32	64	2	True
	64	64	2	True

Table 6: The Space-Time Region Graph architecture used in the paper.

Model	Input Dim.	Output Dim.	Stride	Residual
Non-Local	-	64	1	False
	64	64	1	True
	64	64	1	True
	64	64	1	True
	64	128	2	True
	128	128	1	True
	128	128	1	True
	128	256	2	True
	256	256	1	True
	256	256	1	True
Non-Local _S	-	32	1	False
	32	32	2	True
	32	64	2	True
	64	128	1	True
Non-Local _T	-	32	1	False
	32	32	2	True
	32	64	2	True

Table 7: The Non-local Neural Networks architecture used in the paper.

Model	#params	Reg.	Few-Shot	Full
STGCN	3.42M	72.0 \pm 1.5	22.6 \pm 5.2	61.0 \pm 0.8
STGCN _S	263K	73.0 \pm 1.5	22.5 \pm 4.9	61.7 \pm 0.5
STGCN_T	34K	73.2 \pm 1.6	26.0 \pm 5.7	62.8 \pm 0.6
STGCN-LSTM	2.08M	72.5 \pm 1.6	20.2 \pm 5.0	60.9 \pm 0.7
STGCN-LSTM _S	233K	72.1 \pm 1.7	20.5 \pm 5.9	60.6 \pm 0.7
STGCN-LSTM_T	33K	72.7 \pm 1.4	23.8 \pm 5.9	61.9 \pm 0.6
STGCN-MAX	3.42M	73.5 \pm 4.8	20.2 \pm 6.0	61.7 \pm 4.1
STGCN-MAX_S	263K	73.6 \pm 1.5	28.6 \pm 5.0	63.6 \pm 0.7
STGCN-MAX _T	34K	74.5 \pm 1.2	25.4 \pm 6.1	63.6 \pm 0.7
Space-Time	263K	71.3 \pm 1.4	26.9 \pm 7.1	61.4 \pm 1.2
Space-Time _S	24K	74.7 \pm 1.4	30.7 \pm 6.8	64.9 \pm 0.6
Space-Time_T	14K	74.8 \pm 1.5	31.7 \pm 6.1	65.2 \pm 0.6
Non-Local	1.23M	74.5 \pm 4.2	44.3 \pm 7.0	67.8 \pm 3.7
Non-Local_S	108K	76.5 \pm 2.4	45.0 \pm 6.3	69.5 \pm 2.4
Non-Local _T	32K	75.6 \pm 1.2	44.8 \pm 6.4	68.8 \pm 1.0
TOQ-Net	35K	87.7 \pm 1.3	52.2 \pm 6.3	79.8 \pm 0.8

Table 8: Results on the soccer event dataset for baselines with different capacities, where model_S and model_T denote the small and the tiny variant for each model. The performance is measured by per-action (macro) accuracy, averaged over nine few-shot splits. The \pm values indicate standard errors. For each baseline we showed the best performance over three levels of capacities. In fact, performances of most of the models are not affected much by the capacity. We highlight the best-performing variants of all baselines, and use them in all comparisons in the main text.

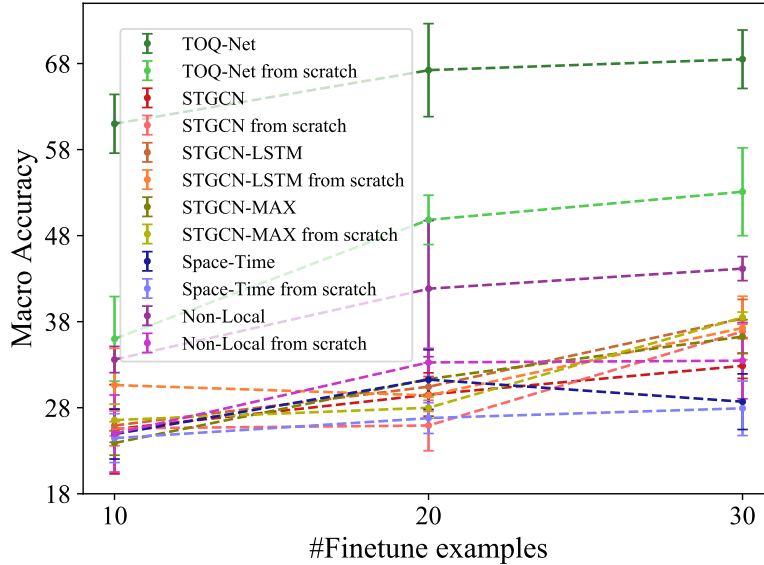


Figure 2: Generalization to soccer environments with a different court size and agent speeds. The standard errors are computed based on three random seeds.

Model	Time Warp	3v3	4v4	6v6	8v8	11v11
STGCN	N	40.7 \pm 1.0 (-40.4%)	63.2 \pm 4.9 (-7.4%)	68.2 \pm 2.8 (0.0%)	55.4 \pm 3.3 (-18.8%)	44.4 \pm 2.1 (-34.9%)
STGCN	Y	32.6 \pm 2.8 (-52.3%)	50.2 \pm 4.4 (-26.5%)	52.8 \pm 7.0 (-22.6%)	43.2 \pm 4.9 (-36.7%)	34.0 \pm 3.7 (-50.2%)
STGCN-MAX	N	47.4 \pm 3.2 (-33.7%)	68.8 \pm 2.0 (-3.8%)	71.5 \pm 1.9 (0.0%)	59.1 \pm 0.7 (-17.3%)	45.6 \pm 2.5 (-36.2%)
STGCN-MAX	Y	37.5 \pm 7.2 (-47.5%)	52.3 \pm 4.2 (-26.9%)	56.5 \pm 4.5 (-21.0%)	46.6 \pm 3.7 (-34.8%)	36.9 \pm 1.9 (-48.4%)
STGCN-LSTM	N	39.7 \pm 1.1 (-43.1%)	60.4 \pm 0.2 (-13.5%)	69.8 \pm 0.1 (0.0%)	55.8 \pm 2.0 (-20.0%)	44.1 \pm 0.7 (-36.8%)
STGCN-LSTM	Y	21.8 \pm 0.8 (-68.8%)	27.8 \pm 1.3 (-60.2%)	30.6 \pm 0.6 (-56.1%)	25.8 \pm 1.0 (-63.1%)	22.6 \pm 0.8 (-67.6%)
Space-Time	N	29.0 \pm 1.6 (-60.4%)	53.5 \pm 3.2 (-27.0%)	73.3 \pm 0.3 (0.0%)	33.9 \pm 2.8 (-53.7%)	15.2 \pm 1.8 (-79.3%)
Space-Time	Y	29.7 \pm 3.1 (-59.5%)	51.6 \pm 2.5 (-29.6%)	70.7 \pm 0.3 (-3.5%)	33.8 \pm 2.2 (-53.9%)	14.9 \pm 1.6 (-79.7%)
Non-Local	N	45.9 \pm 5.1 (-41.2%)	70.7 \pm 5.3 (-9.5%)	78.1 \pm 5.8 (0.0%)	58.5 \pm 10.8 (-25.1%)	41.8 \pm 13.6 (-46.5%)
Non-Local	Y	46.7 \pm 4.1 (-40.2%)	69.9 \pm 4.7 (-10.5%)	77.7 \pm 5.0 (-0.5%)	58.7 \pm 12.6 (-24.9%)	41.3 \pm 13.6 (-47.1%)
TOQ-Net	N	77.4 \pm 3.5 (-12.4%)	88.3 \pm 0.7 (-0.0%)	88.4 \pm 0.6 (0.0%)	81.3 \pm 1.7 (-8.0%)	77.1 \pm 1.7 (-12.8%)
TOQ-Net	Y	76.0 \pm 2.6 (-14.0%)	87.7 \pm 1.4 (-0.8%)	86.9 \pm 0.4 (-1.7%)	80.3 \pm 1.1 (-9.1%)	74.9 \pm 2.3 (-15.2%)

Table 9: Full results on generalization to scenarios with all combinations of #agents and temporally warping on the soccer event dataset. The standard error of all values are smaller than 2.5%, computed based on three random seeds.

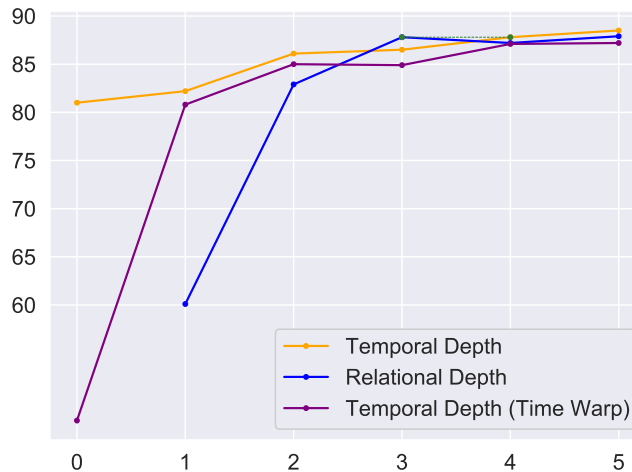


Figure 3: Comparing # of reasoning layers. Accuracy is tested on 6v6 9-way classification. When # of relational reasoning layers vary (blue), temporal reasoning layers are fixed at 4. When # of temporal reasoning layers vary (orange), relational reasoning layers are fixed at 3. When we have 0 temporal reasoning layers, we predict the sequence label based on the feature of the frame of interest. The purple line shows the performance on temporally warped trajectories when we have the # of relational reasoning layers fixed and vary the # of temporal reasoning layers. Relational layers are required for the prediction task (# of relational reasoning layer must be greater than 0), because there is no nullary feature input to the network: we need at least one relational reasoning layer to gather information across the entities. Overall, adding reasoning layers improves results, and 3 relational layers + 4 temporal layers is a good balance between computation and performance.

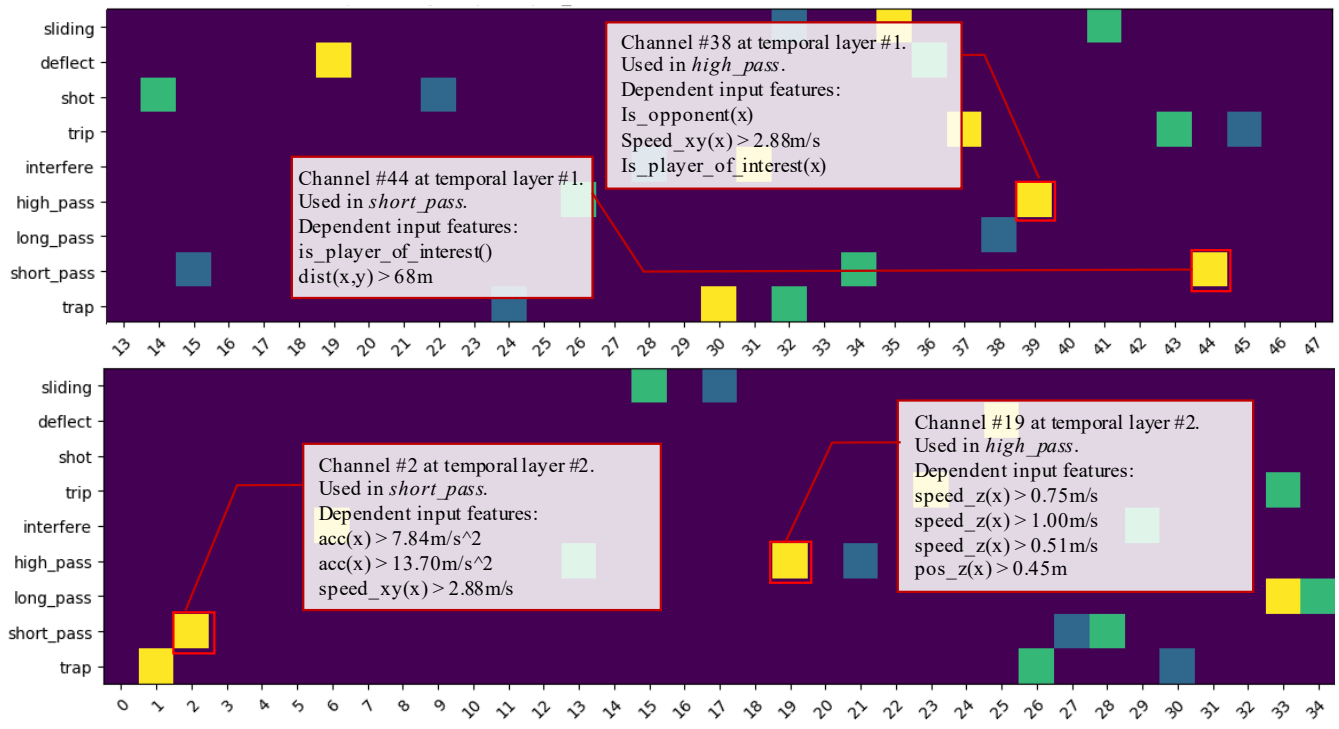


Figure 4: Relevant features in temporal layers. Feature dependencies are computed by gradient. These dependencies and thresholds are learned end-to-end from data. Insets detail features for events in the first and the second stage of the *high pass* and *short pass*.

References

- [James *et al.*, 2020] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark & learning environment. *RA-L*, 2020.
- [Kingma and Ba, 2015] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [Wang and Gupta, 2018] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [Yan *et al.*, 2018] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018.