Using the Principles of Classical Conditioning to Learn Event Sequences

Timothy A. Furze [1] and **Brandon Bennett** [2]

Abstract. In order for an autonomous agent to interact rationally within its environment, it must have knowledge of that environment. Given that the wealth of knowledge that even small children evidently quickly acquire, it is infeasible for an agent to be directly encoded with much, if any, knowledge about the real world. This means that it would be best to instead imbue the agent with the ability to learn the knowledge for itself. Given the non-triviality of the problem of programming an agent with this ability, this paper looks at a system that qualitatively replicates one of the main psychological processes that biological agents use to learn about their environment, that of classical conditioning. Initial testing of the system shows results that are inconclusive but are encouraging. This leads to the conclusion that further work is needed to ascertain the utility of the approach.

1 INTRODUCTION

Classical Conditioning is a phenomenon of learning that begins during an early stage of development, according to Piaget's theory of cognitive development [18]. Due to its prevalence within animals it can be argued to be central to any agent's development of its understanding of its environment. The theory of classical conditioning, primarily introduced by Pavlov [17], allows for an agent to passively learn about its environment. The principal mechanism of classical conditioning is that of an agent learning to associate two stimuli that the agent observes as repeatedly occurring in pairs. The pair of stimuli is usually one stimulus that causes a reflex action in the agent and another stimulus that, if encountered in isolation prior to any pairing, would not cause any reflex.

By considering examples of stimuli pairings that would become associated through classical conditioning in a natural environment of a biological agent, the utility of such a mechanism to the agent can be seen. The smell of a particular food pairing with its taste and the sight of fire pairing with the sensation of heat are two examples of pairs of stimuli that a biological agent could conceivably learn to associate with one another in the course of its development in a natural environment. These sorts of examples suggest that classical conditioning can be seen as a mechanism to infer relationships between stimuli that can be treated as two aspects of the same, more complex, stimulus without the agent having any prior knowledge.

With this conception of classical conditioning in mind, it suggests that the mechanisms of classical conditioning could be used to infer relationships between pairs of events and so allowing the construction of patterns and sequences of events in an unsupervised manner with no prior knowledge. This paper introduces a system that uses a model of classical conditioning in order for an agent to learn to recognise increasingly complex sequences of events starting from a limited set of observed geometrical changes within its environment.

The system that was developed to test and expand on this theory, is comprised of three sub-systems that each provide data to one another forming a feedback loop to allow the system to find increasingly complex sets of patterns. The first sub-system reads a stream of events that describe simple geometrical changes within the observed scene and recognises patterns of those events that occur in its database of event patterns. The second sub-system takes both the base events and the instances of the recognised patterns and provides pairings of event instances that satisfy temporal and event complexity measures. The third sub-system takes the pairs of event instances and provides a list of those event pairs that should be considered significant to the first sub-system to use as its database of event patterns. The third sub-system uses a model of classical conditioning to decide which of the event pairs is significant.

The system was applied to the domain of visual extrinsic object motion (i.e. object tracking) in order to evaluate the system. The test that was done was that of a video of a person throwing a ball in the air. The prediction was that given the data derived from this scene, the system would infer that when the ball went up, that it would expect that the ball would later come down. This would be evidence of the system having developed a simplistic account of gravity.

This paper is structured as follows. Section 2 covers the background of the phenomena of classical conditioning and previous work in the learning and recognition of event sequences. Section 3 looks at the workings of the system that learns the event sequences and how this is done by modelling the mechanisms of classical conditioning. The work done to evaluate the system is presented in section 4. Concluding remarks and potential future directions for this work is then covered in section 5.

2 BACKGROUND

2.1 Classical conditioning

This theory is also known as Pavlovian Conditioning, named after Ivan Pavlov, one of the primary people who introduced the theory. Pavlov's widely-known experiments with dogs, first published in English in 1927 [17] were among the first experiments to demonstrate the collection of phenomena that are now collectively known as classical conditioning. Pavlov's famous experiment conducted with dogs was to create an audible tone (mostly a bell or metronome) immediately prior to the dogs having a substance directly placed into their mouth that would cause the reflex action of salivation (usually meat powder or a weak acid). This was done multiple times. The same audible tone was then presented to the dogs without the presentation of

[1] University of Leeds, United Kingdom, email: phy3taf@leeds.ac.uk
[2] University of Leeds, United Kingdom, email: B.Bennett@leeds.ac.uk

the substance. The result was that the dogs' salivary response was observable with the tone even when substance was not presented. This salivary response without the substance correlated with the number of presentations of the tone where the substance was jointly presented. Pavlov used this experiment and others like it to derive a theory of animal learning.

The derived theory of animal learning from this is that an arbitrary neutral stimulus can become associated with any non-neutral stimulus, (i.e. a stimulus that triggers a reflex response) based on their similar co-occurrence in time. Thus when the neutral stimulus is presented alone, the subject gives a similar response to the unconditioned response, as it has come to expect that the non-neutral stimulus will follow. In the literature around classical conditioning, the names of the stimulus and the responses have particular names. The neutral stimulus is known as the conditioned stimulus (CS) which in Pavlov's experiment corresponds to the generated tone. The non-neutral stimulus is termed the unconditioned stimulus (US) which in Pavlov's experiment corresponds to the substance placed in the dogs' mouths. The response to the non-neutral stimulus is called the unconditioned response (UR) which in Pavlov's experiment corresponds to the salivary reflex the dogs had to the substance. The response to the neutral stimulus after the association had been formed is the conditioned response (CR) which in Pavlov's experiment corresponds to the salivary response the dogs had to the tone when the substance was not present.

There are several phenomena that have been observed in the interaction of CSs and USs. The most notable of these are: Acquisition, Extinction, Reacquisition, Blocking, Secondary Conditioning, The Inter-Stimulus Interval, Intermittent Stimulus Facilitation and Conditioned Inhibition.

- **Acquisition** – Acquisition is the process whereby the CS becomes associated with the US and thus the CR. This is the phenomenon that was discussed above. The strength of the association (e.g. measured by the amount of saliva produced) is a sigmoid-like function of the number of reinforcements of the CS (i.e. the number of presentations of the CS where the US follows).
- **Extinction** – Extinction is the process whereby a CS that is already associated with the US is repeatedly and consistently presented to the subject without the US. The strength of the association is weakened and eventually returns to the same level of association as observed prior to acquisition.
- **Reacquisition** – Reacquisition is the name given to the phenomenon where a previously extinguished CS-US association is acquired again. During reacquisition, it takes a fewer number of reinforcements to re-acquire the same strength association than it did the previous time that association was acquired.
- **Blocking** – Blocking is where a previously conditioned CS stops a second CS from acquiring an association with the US (i.e. demonstrating a CR) when the two CSs are reinforced simultaneously.
- **Secondary Conditioning** – Secondary Conditioning is where a secondary CS can be conditioned to elicit a CR through reinforcement only with a primary CS (where the primary CS has been reinforced with the US). This effect is typically weak as the extinction of the primary CS will happen while the secondary CS is being conditioned.
- **The Inter-Stimulus Interval** – The inter-stimulus interval is the time between the start of the CS and the start of the US. This time gives rise to several situations that affect the acquisition process. This leads to two modes of acquisition, Delay and Trace conditioning. Delay conditioning is where the CS overlaps or finishes immediately before the US appears. Trace conditioning is where the CS finishes with a period of inactivity before the US appears. The inter-stimulus interval affects the rate of acquisition of a CS-US association. The rate follows a curve where small intervals are negligible, it then rapidly moves up to a peak and then gently decays, similar to the curve of a log-normal distribution. The difference between delay and trace conditioning is that the latter has a much faster decay after the peak.
- **Intermittent Stimulus Facilitation** – During conditioning, a longer inter-stimulus interval gives a weaker CR. If a second CS is presented between the first CS and the US, the CR of the first CS is stronger.
- **Conditioned Inhibition** – Conditioned Inhibition refers to an effect where a CS can be made to create an inhibitory effect on a CS-US association. This can be demonstrated in the following experiment: two CSs, CS_1 and CS_2 are conditioned separately to associate with the US. A third CS, CS_0, is then non-reinforced simultaneously with CS_1. Presenting CS_0 simultaneously with CS_2 will then not elicit a CR.

Ever since classical conditioning became widespread in the discourse of psychology, there has been numerous models of classical conditioning that vary in complexity and fidelity. The most well-known model is Rescorla and Wagner's model that was presented in 1972 [22]. This model has served as the basis of later models [12, 30]. The Rescorla-Wagner model works by calculating a difference between the current association strength and what the new trial implies it should be. The rate of learning is based on the salience of both the CS and the US. More recently, there has been a trend to use artificial neural networks to model classical conditioning [26, 25, 10, 7]. Balkenius and Morén [2] presented a comparative study of a number of modern models, including artificial neural network based models, those based on Rescorla and Wagner's model, among others.

2.2 Event sequence learning

Research into learning patterns of event sequences mainly comes from two different fields of computer science research, namely data mining and computer vision. Data mining applies the algorithms that learn event sequences to discover important frequent sequences of events from data that has a temporal component. For example, within the domain of shopping, finding rules that state that certain items have a tendency to be bought at the same time during particular points in the day, or customers who bought one specific item later return to buy another specific item. The main work in the area of mining rules of association (independent of a temporal context) is the work by Agrawal [1]. Work more directly involved with mining associations in a temporal domain is the work of Mannila [13], among others [31, 8, 19]. This work looks to mine sequential patterns of events that appear frequently. This area of data mining as a whole looks more on optimising time, space and I/O write complexity rather than trying to optimise the output rules themselves. Therefore this area, while being relevant in that it attempts to find the same sort of output, is not fully relevant to the work of this paper as the emphasis of the field is more on optimising computational resources rather than trying to have the rules more closely match that of human experience.

Computer vision research in this area more looks at optimising the output itself against a calculated ground-truth with computational efficiency as a secondary goal. One of the influential works in this field, though looks at recognition rather than direct learning is that of Ivanov and Bobick [9] who presented the idea of finding patterns as being akin to parsing a stochastic variant of a context-free grammar. This allowed the powerful idea of looking for events at different levels of abstraction, which is used in this current work. Another important work in the area is that of Stauffer and Grimson [28], who extended their seminal work in object tracking [27] to learn classifications of activity sequences by applying statistical methods to determine co-occurrences.

One approach that has been particularly successful in learning event sequences is to use Inductive Logic Programming. Inductive Logic Programming [15], or ILP, is a branch of machine learning that, through a variety techniques, attempts to find generalised logical rules that explain a set of specific relations. Typically, the rules are expressed as first-order horn clauses. While the technique has been used in the data mining aspect of event sequence learning [19], it has had a larger impact on the computer vision aspect. There have been two prominent works that have used the ideas of ILP to learn event sequences. The first of these is Needham et al. [16], in which the system presented is able to learn from observation only, the rules to a number of simple games, such as paper-scissors-stone. This was accomplished by using an ILP system (PROGOL) [15] to learn generic rules that state the required action given a particular game state. The second work is that of Fern et al. [6], which does not directly use an ILP system, as the authors came to the conclusion that first-order logic horn clauses was a poor representation to use for learning temporal event sequences. However, many of the ideas of ILP were used on a language specifically developed by the authors to represent temporal events. This event system was then used to learn to recognise a variety of verbs from the system being presented a video of that action.

While the system does use first-order logic as to represent its events, the system presented by this paper does not use ILP. The reason for this is two-fold. Firstly, ILP systems in wide use are batch-based programs, where the learning happens in a separate phase to the recognition and all the data the system is required to learn from is required before any recognition can be done. The second reason ILP could not be used by the current system is that ILP requires examples to be labelled as either positive or negative examples of a particular concept, meaning ILP methods are supervised learning methods. The system presented by this paper is an unsupervised system.

3 THE SYSTEM

The purpose of the system is to find sequences of events that are temporally associated with each other by utilising the theories of classical conditioning. This utilisation of the theory of classical conditioning makes one important divergence from most theories of classical conditioning, namely that this system does not assume the need for there to be a reflex-causing stimulus at all, and that a neutral stimulus can gain association with another neutral stimulus via the same mechanism. The reflex response to particular stimuli and the response of stimuli conditioned to them allows for the effects of this association to be measured.

There is evidence that supports this particular divergence. The first piece of evidence is in the phenomenon of classical conditioning known as secondary conditioning, as described in the previous section. This supports the divergence by showing that an association can occur between two conditioned stimuli and that there is nothing inherent in the nature of non-neutral stimuli that causes this association effect to happen. Another piece of supporting evidence is in Rescorla's substantiation of the S-S interpretation of conditioning [21]. The S-S (stimulus-stimulus) interpretation of conditioning states that the CS becomes associated with the US, as opposed to the S-R (stimulus-response) interpretation where the CS becomes associated with the UR. This supports the divergence as it shows that it is not a direct back-propagation of the response when two stimuli become associated.

The remainder of this section describes how the system operates. The system comprises of three component sub-systems that feed data between each other. Figure 1 shows the modules and the data that is passed between them.

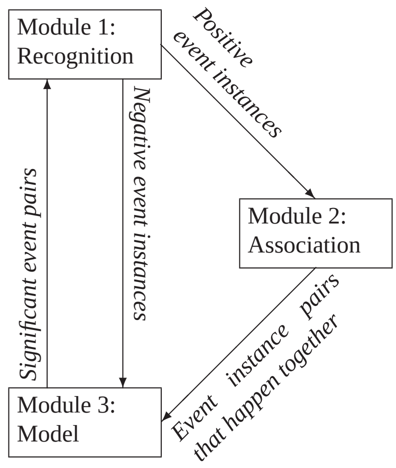

Figure 1. The three sub-systems and the data flows between them.

As input to the system, the system takes a series of time-ordered bounding boxes for each object of interest in the observed scene. In the case of the experiment, this input data would be the bounding boxes for the ball and the person. This is then processed to find geometrical changes, which are used as events to be passed to the first sub-system.

The first sub-system takes the stream of basic events and compares these events with a database comprising of event patterns to be recognised as more complex events. The sub-system recognises both positive and negative instances of these complex events. Positive complex event instances are those event patterns where a pattern is observed. Negative complex event instances are those events where the first half of a pattern is observed, but the latter half of the pattern does not follow.

The second sub-system then identifies and outputs pairings of the positive events whose temporal relationship satisfies a set of criteria such that they can be said to happen together. Only those events that have an equal pattern length are compared for reasons of efficiency.

The third sub-system has two input sources. The first input source is the instances of the identified event pairings with the second input source being the negative complex event instances. To these inputs, the sub-system applies a functional model of classical conditioning. In the model, instances from first input source are treated as positive reinforcements and the instances from the second input are treated as negative reinforcements. This results in a list of pairs of types event instances (both complex and basic) together with a measure of their association strength. The pairs that have a high association strength measure are then fed back to the first sub-system. In the first sub-system, these pairings are treated as a single composite event and are added to the database of events that the first sub-system recognises. Should the association strength of an event pairing that has been allocated a composite event subsequently weaken such that it is no longer considered to have a high association strength, its corresponding composite event is removed from that same list of events.

3.1 The recognition sub-system

The recognition system recognises two types of event instance, atomic events and composite events. Atomic event instances are generated through an analysis of processed sensor data provided as input. The recognition system for atomic events is an expansion of the system presented by dos Santos et al. [5]. Composite event instances are generated by matching their component events against the list of generated events. Composite events may have either atomic events or composite events as their component events but for reasons of computational efficiency, both component events of a composite event must be of equal recursive depth. In other words, the depth of atomic events is zero; the depth of composite events comprising of two atomic events is one; and the depth of a composite event comprising of two composite events that each both comprise of two atomic events is two.

The sub-system outputs positive and negative instances of events. Positive event instances are instances of event pairings that have been observed to happen. Negative event instances are instances of event pairings that were expected to happen but did not. An event is expected to happen when the first component event of a composite event happens, but the second was not observed to happen. By those definitions, all atomic events are positive event instances. The positive event instances are passed to the association sub-system whereas the negative event instances are passed directly to the model of classical conditioning.

The external input to the recognition system used within this paper is data that represents the extrinsic motion of objects within the agent's field of view (i.e. objects moving around a scene, rather than the movement of sub-components of the object while the object itself is static). This data is split into temporal frames. In each frame each object is represented as a bounding box labelled with an identifier unique to that object. It is expected that the system is general enough to be applicable to different domains.

For each frame, a set of state information regarding the objects present within the frame is generated. The set of state variables initially includes the x and y position of the centre of each box. The remainder of the state variables are based on each pair of objects. Each pair of objects has four state variables that describe their relationship. The first state variable is the distance between the centres of each box. The next state variable represents one of the mutually exclusive possible states of *"A is coalescent with B"* (which means that the boxes of the two objects A and B overlap to the extent that the two objects cannot be reliably distinguished), *"A is externally connected with B"* (the two boxes are touching but do not significantly overlap) or *"A is disconnected with B"* (the two boxes are distinctly separate). These three possible states are based on a variant of the region connection calculus [20, 24]. The third state variable represents one of the possible mutually exclusive states *"A is to the left of B"*, *"B is to the left of A"* or *"Both A and B are in-line in the X axis"*. The final state variable represents one of the possible mutually exclusive states *"A is above B"*, *"B is above A"* or *"Both A and B are in-line in the Y axis"*.

After these states have been generated for a frame, they are compared with the states of the previous frame. Based on the changes in each state type, multiple atomic events are generated based on each atomic event's logical definition encoded within the system. Figure 2 lists the names and English definitions of the events that can be generated. Note that the last two events are generated by comparing the lists of objects present in a frame rather than from any of the states generated. These atomic events are those that have been identified as being pertinent to the test domain of the extrinsic motion of objects.

- `Approaching(X,Y)` – `X` and `Y` are approaching each other.
- `Receding(X,Y)` – `X` and `Y` are receding from each other.
- `Static(X,Y)` – The distance separating `X` and `Y` does not change.
- `MergeR(X,Y)` – `X` is merging with `Y` on the right of `Y`.
- `MergeL(X,Y)` – `X` is merging with `Y` on the left of `Y`.
- `MergeT(X,Y)` – `X` is merging with `Y` on the top of `Y`.
- `MergeB(X,Y)` – `X` is merging with `Y` on the bottom of `Y`.
- `EmergeR(X,Y)` – `X` is emerging from `Y` on the right of `Y`.
- `EmergeL(X,Y)` – `X` is emerging from `Y` on the left of `Y`.
- `EmergeT(X,Y)` – `X` is emerging from `Y` on the top of `Y`.
- `EmergeB(X,Y)` – `X` is emerging from `Y` on the bottom of `Y`.
- `MakeCR(X,Y)` – `X` has made contact with `Y` on the right of `Y`.
- `MakeCL(X,Y)` – `X` has made contact with `Y` on the left of `Y`.
- `MakeCT(X,Y)` – `X` has made contact with `Y` on the top of `Y`.
- `MakeCB(X,Y)` – `X` has made contact with `Y` on the bottom of `Y`.
- `BreakCR(X,Y)` – `X` has broken contact with `Y` on the right of `Y`.
- `BreakCL(X,Y)` – `X` has broken contact with `Y` on the left of `Y`.
- `BreakCT(X,Y)` – `X` has broken contact with `Y` on the top of `Y`.
- `BreakCB(X,Y)` – `X` has broken contact with `Y` on the bottom of `Y`.
- `MoveRight(X)` – `X` has moved right.
- `MoveLeft(X)` – `X` has moved left.
- `MoveUp(X)` – `X` has moved up.
- `MoveDown(X)` – `X` has moved down.
- `Lost(X)` – Object `X` has ceased to be detected.
- `Found(X)` – Object `X` has been newly detected.

Figure 2. The event types that the system uses to describe the transition between the states of one frame and the states of the next for the domain of extrinsic motion.

The list of atomic events is then compared with the list of composite event types (which is initially empty and is grown by the feedback from the model of classical conditioning sub-system). Where an atomic event is the first event of a composite event that appears in the list, an event instance of the type of the matched composite event is generated and is marked as being a potential event (as the second sub-event has yet to be observed). A potential event is an event that is believed to be currently ongoing but there is not the evidence to know for sure. The generated potential event is then recursively compared with the list of composite event types to generate further potential events of increasing complexity.

After the potential events have been generated, they need to be grown to so they can represent their true observed duration. The set of both the atomic and potential events of the time in-between the current and previous frame are compared with the events that were generated when the now-previous frame was the current frame. Where the same event has been generated in both consecutive frames, the event token of the event in the previous frame is extended to cover the current time frame of the event and the duplicate newly generated event instance is removed.

At the next stage, the list of potential events that are within a predetermined window of time before the current frame is compared to the list of atomic events that were generated during the current frame. If any of the atomic events are the second event of a potential event, the potential event instance in its entirety is replaced with an actual event as that potential event has now been confirmed. The set of newly confirmed potential events is then recursively compared with the list of potential events to generate further confirmed events of increasing complexity.

Where a potential event has yet to see its second event, but the first event has finished and its finishing time was longer ago than the width of the predetermined window of time before the current frame, then the potential event is classed as a negative event instance and is passed as such to the model of classical conditioning.

These stages outlined above are repeated for every subsequent pair of frames provided as input. Note that the system has been designed to be able to be used in an on-line manner. This on-line nature was required so that the system may continuously learn new associations throughout the lifetime of the agent.

3.2 The association sub-system

The purpose of the association system is to systematically record each pairing of event instances that are temporally close enough together that, based on defined criteria, they can be said to happen together.

The criteria that define the notion of two events happening together is based on the modes of conditioning that are a part of the inter-stimulus interval phenomena of classical conditioning, namely delay and trace conditioning. Delay conditioning notes that the period of the conditioned stimulus can either stop at the start of, or overlap, the period of the unconditioned stimulus. Whereas trace conditioning shows that end of the conditioned stimulus can have a short gap before the start of the conditioned stimulus, though the longer the gap the slower any association is formed. These ideas suggest for criteria for the notion of two events happening together, either the two events must temporally overlap or that the first event must have its finishing point within a defined window of time before the beginning of the second.

Due to these criteria, the association sub-system calculates its pairings based on whether an event is starting, stopping or continuing. An event instance is considered to be starting if the event that was generated in the current frame but not in the previous frame. An event instance is considered to have stopped if was generated in the previous frame but not in the current frame. An event instance is considered to be continuing if it was generated in both the current frame and the previous frame.

For every starting event that the recognition system generates, the association system records the list of events that occurred within the defined window of time before the current frame including those that are ongoing. Where an event is ongoing, it is marked as so in the list.

As each event finishes, the association system looks for all the occurrences of that event in the list of event pairings and notes its finishing time against those listings, removing the marker that it is a continuing event. As the parings get to the stage where both events have finished, they are passed to the model of classical conditioning sub-system.

Figure 3 depicts the moving window and 12 intervals of events. Each interval is inclusive at both ends, so for instance, event interval 1 is over six time steps. The current time step is marked as t, meaning that in the diagram, event 12 has not started happening yet. In the diagram, event 8 is starting, event 7 is stopping and events 5 and 11 are continuing. W is the length of time of the window; again, this is inclusive at both ends so that the system would generate an pairing instance for events 8 and 9. In fact, for this diagram only 3 of all possible pairings of the events would not be generated, being 3&8, 3&12 and 9&12. The density of the events is for illustration purposes and in the practical example of the test case, the events are more sparse.

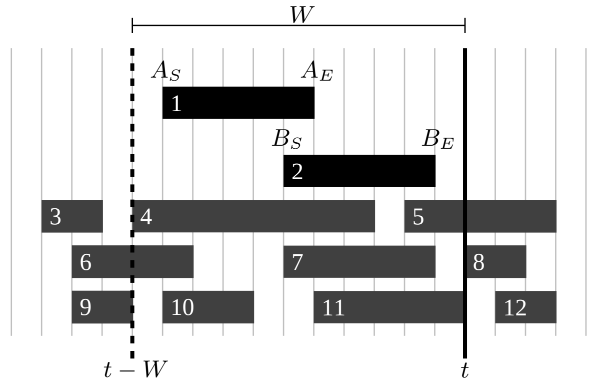

Figure 3. A demonstration of the window in relation to a series of events. Note that the vertical grouping of event intervals in the diagram is arbitrary.

3.3 The model of classical conditioning

The purpose of the model of classical conditioning as a sub-system is to create a mapping from a list of instances of event pairings and a list of negative event instances to a measure of the association strength for that pairing each event type present. While the other parts of the system also are responsible for modelling some of the phenomena of classical conditioning, it is this sub-system that attempts to model the main phenomena.

The model presented by this paper is a relatively simplistic model that does not claim to be able to compete neither on fidelity nor on complexity with those models that were developed as an exercise in of themselves. This raises the question of why the effort was undertaken to produce a new model at all, after all, if there are better models already in existence, why was one of these models not implemented instead? The reason is due to the divergence in the theory stated at the beginning of this section, that neutral stimuli can associate together without the presence of a reflex-causing stimulus. All of the models that have been encountered make the assumption that there is a natural strength of reflex of the reflex-causing stimulus that can be propagated across an association and is available to be factored into the calculation of the association strength. This assumption means that they cannot be used in this system. This is due to the very definition of being neutral stimulus, they cause no reflex action and so do not have any measure of the strength of reflex that could be propagated. So the model in this system attempts to be a proof-of-concept model.

The two inputs of the sub-system, the list of event pairings from the association sub-system and the list of negative event instances from the recognition sub-system, represent the twin notions of reinforcements and non-reinforcements. This sub-system treats them as such in the modelling.

The model was primarily developed through examining and attempting to approximate in a function, the response curves of the various phenomena as described in [2]. While this approach does not attempt to provide any explanatory power, it does allow for the desired responses. This approach has led to the production of three functions that determine different aspects of the association strength. All three functions are designed so that they perform in an iterative manner. In other words, the functions output the amount the current association strength (Y_N) should be changed by (δY), rather than calculating the new association strength (Y_{N+1}) directly. This means that only the current association strength needs to be stored in memory rather than retaining all the inputs to each of the functions. Note that the association strength is real-valued and constrained to the range $0 \leq Y \leq 1$. The new association strength is updated according to equation 1.

$$Y_{N+1} = Y_N + \delta Y \tag{1}$$

The first function, shown in equation 2, models the curve observed in the acquisition phenomenon. This function is applied for each reinforcement of an event pairing. As described previously in the background section, the acquisition phenomenon follows a sigmoid-like curve. In the equation, δX is the amount one reinforcement instance is to be counted (this is normally equal to 1), k_1 is a constant representing the learning rate of acquisition and Z is the output of the functions that model the effect of a change in the inter-stimulus interval.

$$\delta Y = Z \frac{(1 - Y_N)e^{k_1 \delta X} + Y_N - 1}{e^{k_1 \delta X} + Y_N + \frac{2}{Y_N} - 3} \tag{2}$$

The next function, shown in equation 3, models the effect of extinction. This function is applied for each non-reinforcement (i.e. a negative event instance) the sub-system receives for a given pair of events. Note that the functions that model the effect of changes in the inter-stimulus interval are not applied to the extinction function. This is because in the case of a negative instance, the size of the second event is not available, this means there is no inter-stimulus interval to be measured and so the functions cannot be applied. [2] did not provide any description of the extinction decay curve, however, Pavlov provided a small sample of data in lecture 4 of [17] that suggests a linear decay. In the equation, δX is the amount one non-reinforcement instance is to be counted (this is normally equal to 1) and k_2 is a constant representing the learning rate of extinction.

$$\delta Y = -k_2 \delta X \tag{3}$$

The final functions, are the functions that model the change in response due to changes in the inter-stimulus interval. These functions are only applied when dealing with reinforcements, as opposed to non-reinforcements. The reason that these functions have not been merged into equation 2 is due to the complexity of the equations. To allow for these functions to alter the output of acquisition function, its output is constrained to $0 \leq Z \leq 1$. As described previously in the background section, the phenomena due to changing the inter-stimulus interval suggests a curve similar to the curve of the log-normal distribution. In these equations, Z is the factor that the output of the acquisition function is to be multiplied by, I and J are intermediary values used to allow the function to be shown in a simpler form, A_S is the start time of the first event, A_E is the end time of the first event, B_S is the start time of the second event, B_E is the end time of the second event and W is the defined size of the moving window. Figure 3 shows these variables in relation to the pair of events 1 and 2 in that diagram.

$$I = \frac{1}{2} - \left(\frac{\max\left(0, \frac{A_E - B_S}{B_E - B_S}\right)}{2} \right) + \left(\frac{\max\left(0, \frac{B_S - A_E}{W}\right)}{2} \right) \tag{4}$$

$$J = \max\left(0, (|B_S - A_S| - 2I)\right) \tag{5}$$

$$Z = \frac{2(2 - I)e^{\frac{-2(\ln(J) - 1)^2}{(2+I)^2}}}{J(2 + I)\sqrt{\frac{\pi}{2}}} \tag{6}$$

When an association strength goes above a certain defined threshold, that pairing is added to the list of rules that the first sub-system uses to recognise as a composite event. If a pairing drops below the threshold through the extinction processes, that pairing is removed from the list.

This feedback of information is one of the central ideas of the system as it both allows for patterns of arbitrary length to be built up yet does not allow any combinatorial explosion to take place. It also has to be recognised that this can mean that the more complex the composite event the system needs to learn, the more examples it requires. This means that this list builds up simple representations first, creating the event representations that have a minimum description length before updating them with longer ones as required. With the ability to remove sequences that no longer have a strong enough evidence base, the system is able to retract locally maximal artifacts that are due to coincidences.

4 EVALUATING THE SYSTEM

The intention of the system was for it to passively learn about its presented environment without any initial data regarding that environment. To this end, the system was tested to see if it could find any patterns of events that can be argued to be semantically important with reference to the environment. The domain of extrinsic object motion was chosen due to its prevalence within computer vision and that it was the domain used by the work of dos Santos et al. [5] that formed one of the bases of this work. The domain also allows for the use of the principles of physical mechanics to form predictions.

The environment that was chosen was that of observing a person repeatedly throwing a ball in the air and catching it. The prediction was that the system would find a pattern of events that would represent the ball being thrown upwards followed by it falling downwards. This would mean that the system has come to expect (and through the system of potential events, generates expectations) that whenever the ball moves upwards, it will at some point come down again, this would be an expectation of gravity to enact on the ball.

Note that this application domain may appear to be similar to the application domain presented Bennett et al. [3]. However, this is not the case. The domain in Bennett et al. [3] used a basketball-like domain with multiple moving people as well as a moving ball as a source of complex movement to test the capabilities of the presented tracking system. The domain of the current paper uses a single, static person and a moving ball to allow for a domain simple enough to allow for a testable prediction to be made of what the system should learn.

An approximately 2:45 minute video (5006 frames at 30fps) was shot of a person throwing a ball in the air. This video was then hand-processed using the ViPER annotation tools [4, 14] so that the extrinsic motion of the relevant objects within the viewer could be extracted without need for an object tracking system, so that the inaccuracies of a tracking system could be avoided. The tracking data was then converted into a suitable format and input into to the system. The system was run with a window size of 30 frames, a rule association strength threshold of 0.85 and equation constants k_1 and k_2 set at 3 and 0.1 respectively.

The list of rules that had been generated by the system after it had completed processing every frame had 30 rules. None of these rules were compound rules. On inspection of the list of all association strengths, the majority of the associations were for compound events, and some were only marginally outside the threshold.

Figure 4 shows those pairings of events that were above the threshold. These are a mixture of encouraging results with a couple of anomalous results. For an effect that was reasonably expected, there is the tendency of groupings of related concepts. For example, results 2 to 5 indirectly imply that when static(A,B) holds that static(B,A) holds and that when an object A makes contact with the bottom of an object B, then object B has made contact with the top of object A. The knowledge of these implications is not coded into the system in any capacity as each atomic event is independently searched for and generated.

The main encouraging results given the prediction made, is that of 7 & 8 and 11 to 14. 7 & 8 show that the system is expecting for the ball to be receding from the person when it is moving up, and 11 to 14 show that the system expects that when the ball emerges from the bounding box of the person, that it also breaks contact with the box.

The majority of the anomalous results relate to the relations showing various types of stasis. A number of these can be explained by the nature of the recorded video. The video recording was of a relatively low quality, which included the movement being jerky in places. The prevalence of the static events could be attributed to this. This throws up the question of the utility of recording the stasis events at all.

One interesting and unexpected rule is number 29. It was unexpected as the person does not make many movements other than with the arms. This result is due to the person moving their arms up above their head to throw the ball up. This makes the bounding box of the person taller and so the centre point of the bounding box moves up.

5 CONCLUSIONS AND FURTHER WORK

The results found in testing the system presented in this paper appear to be inconclusive but encouraging. The best explanation that can be offered for the lack of composite rules is that the video used was too short to give the system the time that would be needed to see these rules gain a high enough association strength to be included. The results are encouraging though, as several parts that would be required for a full composite rule that would expect gravity to enact on the ball are present.

Further work in the short term would be to re-run the experiment for a longer period of footage that is recorded with higher quality equipment. From this, a more concrete conclusion could be formed.

Beyond that, the first area of improvement to the system would be to create a model of classical conditioning that models a greater number of the phenomena in better quality. For instance, reacquisition, blocking and inhibitory phenomena are not implemented in the model presented.

Within a wider field, the system could be adapted to also model operant (instrumental) conditioning, this could be done by adding in agent actions as events in the system along with reward and punishment events. The work by Touretzky et al. [29, 23] may be useful in assisting work towards this goal.

It can be observed that animals learn both passively and actively. It is argued that an effective agent must be able learn using both modes.

```
 1. staticX(personA), moveDown(personA)
 2. static(personA, ball), makeCB(personA,
    ball)
 3. static(ball, personA), makeCB(personA,
    ball)
 4. static(personA, ball), makeCT(ball,
    personA)
 5. static(ball, personA), makeCT(ball,
    personA)
 6. staticX(personA), moveRight(personA)
 7. moveUp(ball), receding(personA, ball)
 8. moveUp(ball), receding(ball, personA)
 9. moveLeft(ball), static(personA, ball)
10. moveLeft(ball), static(ball, personA)
11. emergeB(personA, ball), breakCB(personA,
    ball)
12. emergeT(ball, personA), breakCB(personA,
    ball)
13. emergeB(personA, ball), breakCT(ball,
    personA)
14. emergeT(ball, personA), breakCT(ball,
    personA)
15. moveLeft(ball), approaching(personA, ball)
16. moveLeft(ball), approaching(ball, personA)
17. static(personA, ball), mergeB(personA,
    ball)
18. static(ball, personA), mergeB(personA,
    ball)
19. static(personA, ball), mergeT(ball,
    personA)
20. static(ball, personA), mergeT(ball,
    personA)
21. staticX(ball), staticY(ball)
22. staticX(personA), staticY(ball)
23. staticX(personA), moveDown(ball)
24. staticY(personA), moveDown(ball)
25. staticX(personA), staticY(personA)
26. staticX(personA), moveRight(ball)
27. staticX(ball), moveRight(ball)
28. staticX(ball), moveDown(ball)
29. moveUp(personA), moveUp(ball)
30. staticX(personA), moveLeft(personA)
```

Figure 4. The resultant pairs of events that the system considered to be compound events after processing all the input data.

For instance, an animal can associate the sound of a rock slide with the sight of falling rocks. It can also be learn to actively avoid being hit by a rock. Only when both passive and active learning are together can the animal associate the sound of a rock slide with danger, without actually being caught in a rock slide. For another example, consider using a hairdryer to move a toy sailing ship. For a planning system to decide that course of action, the agent would need to have passively associated air currents with moving sailing ships and observed that the action of activating a hairdryer causes an air current.

During the development of the system, a question kept surfacing about randomised outcomes to event sequences. How should the system deal with event sequences where the outcome event is not deterministic but can be one of a set of outcomes? For an example, the rolling of a die; here there is a definite sequence of events leading up to the outcome. However there is not a single outcome but a definite set of outcomes. For example, one would not expect a seven to appear on a standard six-sided die. There are methods that do learn stochastic event sequences [11] but these operate in a batch manner. If it is possible for the system presented in this paper to learn stochastic events, then the system would be capable of adapting its existing hypotheses as new examples of the patterns of events are presented. This system, when combined with an extension to account for instrumental conditioning, could, in an unsupervised manner, dynamically learn about how an agent expects its environment to behave, in a way that allows adaptation to changes in that environment.

ACKNOWLEDGEMENTS

We would like to thank the referees for their comments which helped improve this paper.

REFERENCES

[1] Rakesh Agrawal and Ramakrishnan Srikant, 'Fast algorithms for mining association rules in large databases', in *Proceedings of the 20th International Conference on Very Large Data Bases*, eds., Jorge B. Bocca, Matthias Jarke, and Carlo Zaniolo, pp. 487–499, Santiago, Chile, (September 1994). Morgan Kaufmann Publishers Inc.

[2] Christian Balkenius and Jan Morén, 'Computational models of classical conditioning: A comparative study', Technical Report LUCS 62, Lund University Cognitive Studies, Lund, Sweden, (1998).

[3] Brandon Bennett, Derek R. Magee, Anthony G. Cohn, and David C. Hogg, 'Enhanced tracking and recognition of moving objects by reasoning about spatio-temporal continuity', *Image and Vision Computing*, **26**(1), 67–81, (January 2008).

[4] David Doermann and David Mihalcik, 'Tools and techniques for video performance evaluation', in *Proceedings of the 15th International Conference on Pattern Recognition*, eds., Alberto Sanfeliu and Juan J. Villanueva, volume 4, pp. 167–170, Barcelona, Spain, (September 2000). IEEE.

[5] Marcus V. dos Santos, Rod C. de Brito, Ho-Hyun Park, and Paulo E. Santos, 'Logic-based interpretation of geometrically observable changes occurring in dynamic scenes', *Applied Intelligence*, **31**(2), 161–179, (October 2009).

[6] Alan Fern, Robert Givan, and Jeffrey M. Siskind, 'Specific-to-general learning for temporal events with application to learning event definitions from video', *Journal of Artificial Intelligence Research*, **17**, 379–449, (December 2002).

[7] H. Hassan and M. Watany, 'On mathematical analysis of pavlovian conditioning learning process using artificial neural network model', in *Proceedings of the 10th Mediterranean Electrotechnical Conference*, eds., Constantinos S. Pattichis and Stathis Panis, volume 2, pp. 578–581, Cyprus, (May 2000). IEEE.

[8] Frank Höppner, 'Discovery of temporal patterns: Learning rules about the qualitative behaviour of time series', in *Proceedings of the Fifth European Conference on Principles of Data Mining and Knowledge Discovery*, eds., Luc De Raedt and Arno Siebes, pp. 192–203, Freiburg, Germany, (September 2001). Springer-Verlag.

[9] Yuri A. Ivanov and Aaron F. Bobick, 'Recognition of visual activities and interactions by stochastic parsing', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**(8), 852–872, (August 2000).

[10] Christopher Johansson and Anders Lansner, 'An associative neural network model of classical conditioning', Technical Report TRITA-NA-P0217, Department of Numerical Analysis and Computer Science, Royal Institute of Technology, Stockholm, Sweden, (2002).

[11] Stanley. Kok and Pedro Domingos, 'Learning markov logic network structure via hypergraph lifting', in *Proceedings of the 26th Annual International Conference on Machine Learning*, eds., Andrea P. Danyluk, Léon Bottou, and Michael L. Littman, pp. 505–512, Montreal, Canada, (June 2009). ACM.

[12] Nicholas J. Mackintosh, 'A theory of attention: Variations in the associability of stimuli with reinforcement', *Psychological Review*, **82**(4), 276–298, (July 1975).

[13] Heikki Mannila, Hannu Toivonen, and A. Inkeri Verkamo, 'Discovery of frequent episodes in event sequences', *Data Mining and Knowledge Discovery*, **1**(3), 259–289, (September 1997).

[14] Vladimir Y. Mariano, Junghye Min, Jin-Hyeong Park, Rangachar Kasturi, David Mihalcik, Huiping Li, David Doermann, and Thomas Drayer, 'Performance evaluation of object detection algorithms', in *Proceedings of the 16th International Conference on Pattern Recognition*, eds., Rangachar Kasturi, Denis Laurendeau, and Ching Y. Suen, volume 3, pp. 965–969, Quebec City, QC, Canada, (August 2002). IEEE.

[15] Stephen Muggleton, 'Inverse entailment and progol', *New Generation Computing*, **13**, 3–4, (December 1995).

[16] Chris J. Needham, Paulo E. Santos, Derek R. Magee, Vincent Devin, David C. Hogg, and Anthony G. Cohn, 'Protocols from perceptual observations', *Artificial Intelligence*, **167**(1–2), 103–136, (September 2005).

[17] Ivan P. Pavlov, *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*, Oxford University Press, 1927.

[18] Jean Piaget, 'The origins of intelligence in children', in *The Essential Piaget*, eds., Howard E. Gruber and J. Jacques Vonèche, 215–249, Basic Books, New York, New York, USA, (1977).

[19] Luboš Popelínský and Jan Blaťák, 'Toward mining of spatiotemporal maximal frequent patterns', in *Proceedings of Ninth European Conference on Principles and Practice of Knowledge Discovery in Databases Workshop on Mining Spatio-Temporal Data*, eds., Gennady Andrienko, Donato Malerba, Michael May, and Maguelonne Teisseire, Porto, Portugal, (October 2005).

[20] David A. Randell, Zhan Cui, and Anthony G. Cohn, 'A spatial logic based on regions and connection', in *Proceedings of 3rd International Conference on Knowledge Representation and Reasoning*, eds., Bernhard Nebel, Charles Rich, and William Swartout, pp. 165–176, Cambridge, Massachusetts, USA, (October 1992). Morgan Kaufmann.

[21] Robert A. Rescorla, 'Effect of us habituation following conditioning', *Journal of Comparative and Physiological Psychology*, **82**(1), 137–143, (January 1973).

[22] Robert A. Rescorla and Allan R. Wagner, 'A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement', in *Classical Conditioning II: Current Research and Theory*, eds., Abraham H. Black and William F. Prokasy, pp. 64–99, New York, New York, USA, (1972). Appleton Century-Crofts.

[23] Lisa M. Saksida, Scott M. Raymond, and David S. Touretzky, 'Shaping robot behavior using principles from instrumental conditioning', *Robotics and Autonomous Systems*, **22**(3–4), 231–249, (December 1997).

[24] Paulo Santos and Murray Shanahan, 'Hypothesising object relations from image transitions', in *15th European Conference on Artificial Intelligence*, ed., Frank van Harmelen, pp. 292–296, Lyon, France, (July 2002). IOS Press.

[25] Jos A. Schmajuk, Nestor A.; Larrauri, 'Experimental challenges to theories of classical conditioning: Application of an attentional model of storage and retrieval', *Journal of Experimental Psychology: Animal Behavior Processes*, **32**(1), 1–20, (January 2006).

[26] Nestor A. Schmajuk and James J. DiCarlo, 'A neural network approach to hippocampal function in classical conditioning', *Behavioral Neuroscience*, **105**(1), 82–110, (February 1991).

[27] Chris Stauffer and W. Eric L. Grimson, 'Adaptive background mixture models for real-time tracking', in *Proceedings of the 1999 Conference on Computer Vision and Pattern Recognition*, eds., Bruce A. Draper and Allen Hanson, volume 2, pp. 2246–2252, Fort Collins, CO , USA, (June 1999). IEEE.

[28] Chris Stauffer and W. Eric L. Grimson, 'Learning patterns of activity using real-time tracking', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**(8), 747–757, (August 2000).

[29] David S. Touretzky and Lisa M. Saksida, 'Operant conditioning in skinnerbots', *Adaptive Behavior*, **5**(3–4), 219–247, (January 1997).

[30] Linda J. van Hamme and Edward A. Wasserman, 'Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements', *Learning and Motivation*, **25**(2), 127–151, (May 1994).

[31] Mohammed J. Zaki, 'Spade: An efficient algorithm for mining frequent sequences', *Machine Learning*, **42**(1–2), 31–60, (January 2001).