Learning Driving Behavior By Timed Syntactic Pattern Recognition

Sicco Verwer, Mathijs de Weerdt, and Cees Witteveen

Given time series data,
discretize but include time information

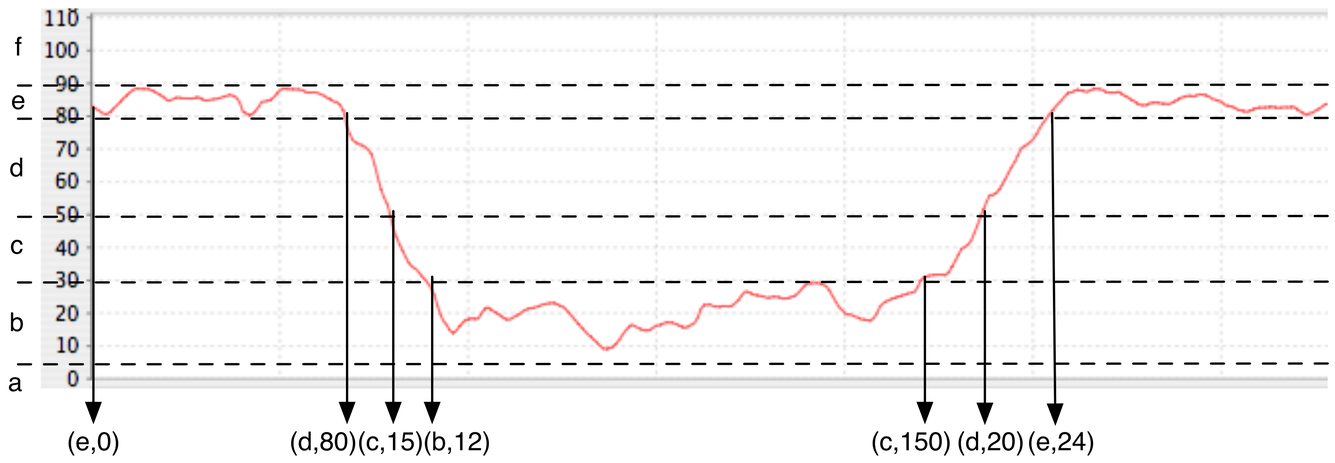

in order to
learn a timed automaton model

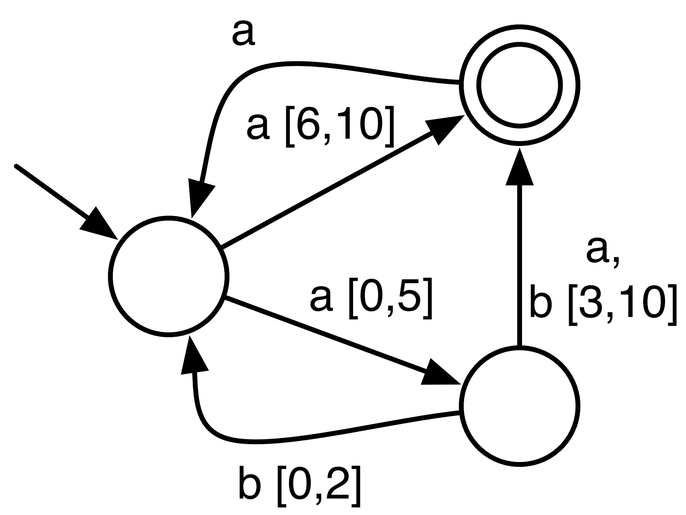

Real-time automata succinctly capture
timed relations between events

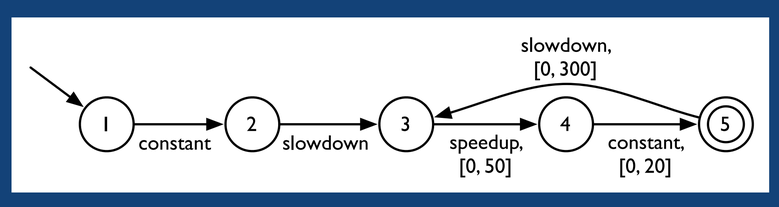

A deterministic real-time automaton that models `harmonica driving'. This often occurs when a truck is driving at a somewhat higher speed than the vehicle directly in front of it. The driver slows down a bit, waits until there is enough distance between him and the vehicle in front (this takes at most 50 seconds), and then speeds up again and closes in on the vehicle (for at most 20 seconds). This whole process often repeats itself a couple of times (within 5 minutes) before the driver finally adjusts the speed of the truck to match the vehicle in front of him.

The result of this whole process is unnecessary fuel consumption.

Use a deterministic real-time automaton to
model the causal states of a process

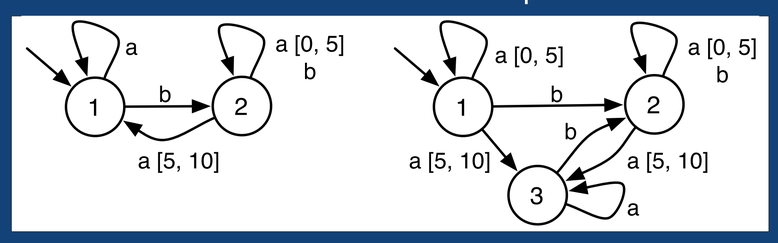

The causal states is an automaton model that represents behavior identical to a non-deterministic version of the original process model. This non-deterministic version has an identical structure as the original model, but has the complete set of states as possible start states.

Left: an original deterministic real-time automaton model.
Right: the determinized causal states of the original model.

RTI+ learns a model for the causal states
from random subsequences of a single
time-stamped event sequence

We obtain historical data consisting of 50,000 event occurrences for the truck speed, and about 100,000 occurrences for the truck engine and fuel sensors. From this data, we take 10,000 random subsequences of length 20 and use these to learn models for the causal states of each sensor using the RTI+ algorithm.

Results:

	size	AIC	AIC single state model
speed	2,305	871,550	1,472,840
fuel	960	842,536	1,102,270
engine	1,009	697,190	1,076,750

Semi-supervised learning of real-time automata classifiers
by labelling the learned causal states

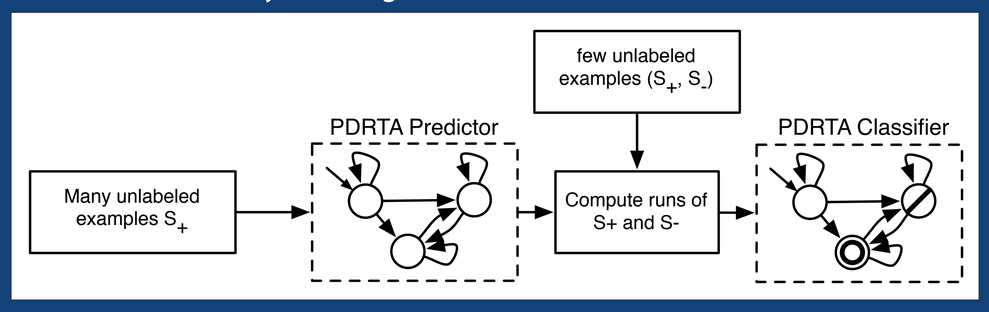

We test our approach on a type of driving behavior that results in unnecessary fuel usage: accelerating too quickly. We obtained some labeled examples for this behavior by driving a short test lap with many traffic lights. After a full stop, the truck driver accelerated too fast about half of the time, and the other half he accelerated normally. We labeled the timed-stamped event sequences that occurred during the first 1.5 minutes from the start of the acceleration.

The labeled examples for the speed sensor:

accelerating too fast	normal acceleration
$(a, 42)(b, 80)(c, 12)(d, 10)(c, 7)$	$(a, 10)(b, 14)(c, 21)(b, 42)$
$(a, 13)(b, 65)(c, 6)(d, 8)$	$(a, 11)(b, 34)(c, 13)(b, 18)$
$(a, 6)(b, 39)(c, 10)(d, 10)$	$(a, 8)(b, 10)(c, 14)(b, 45)$
$(a, 7)(b, 15)(c, 8)(d, 7)$	$(a, 25)(b, 5)(c, 18)(b, 27)$
$(a, 7)(b, 2)(c, 7)(d, 7)$	$(a, 11)(b, 6)(c, 20)(b, 14)$
$(c, 22)(b, 4)(c, 12)(d, 7)(c, 8)$	$(a, 12)(b, 18)(c, 15)(b, 20)$
$(a, 7)(b, 11)(c, 6)(d, 7)(c, 16)$	$(a, 8)(b, 35)(c, 24)(b, 20)$
$(a, 8)(b, 29)(c, 7)(d, 8)(c, 8)$	$(a, 12)(b, 32)(c, 18)(d, 13)$
$(a, 11)(b, 28)(c, 7)(b, 12)$	$(a, 7)(b, 62)(c, 21)(b, 28)$
$(a, 10)(b, 320)(c, 12)(d, 7)(c, 34)$	$(a, 12)(b, 8)(c, 17)$
$(a, 10)(b, 8)(c, 8)(d, 8)(c, 31)$	$(a, 10)(b, 27)(c, 19)(d, 16)$
$(a, 8)(b, 24)(c, 9)(d, 8)(c, 12)$	

Labeled causal states for the speed sensor

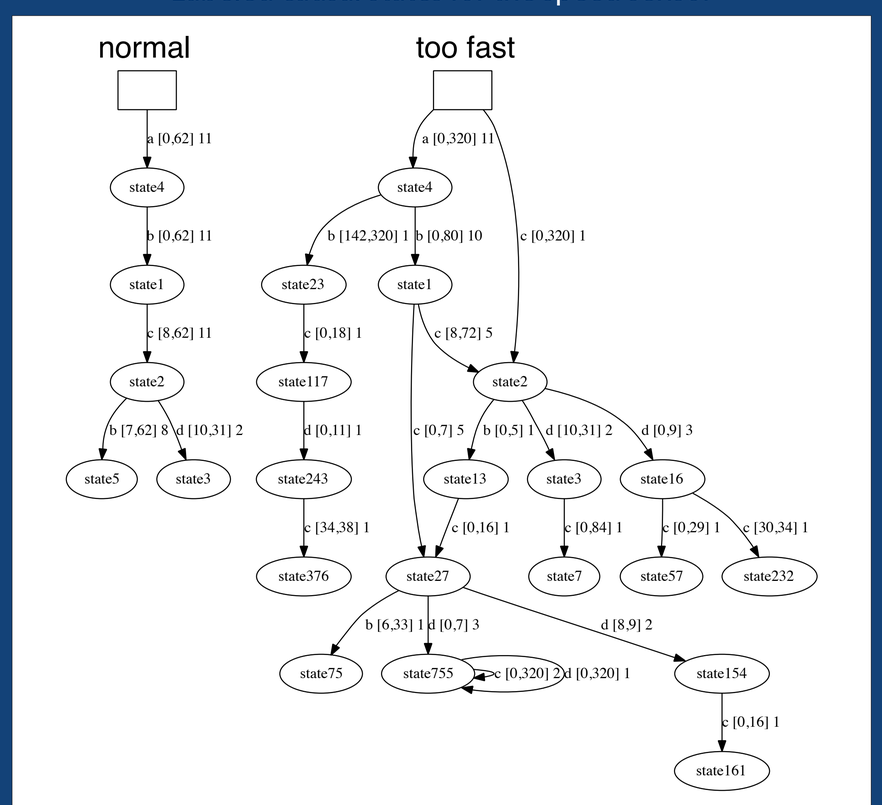

Classifier results

We compare the labels given by a rule provided by domain experts to the labels assigned by the labeled causal states.

These rules match in 79% of all 1500 cases. This is high considering that only 23 example sequences were used to label the causal states.

Result table:

		decision rule	
		quick	normal
classifier	quick	920	61
	normal	262	289

KATHOLIEKE UNIVERSITEIT LEUVEN

TU Delft