Learning acceptable windows of contingency

Kevin Gold and Brian Scassellati
Department of Computer Science, Yale University, CT 06511 USA
email: kevin.gold@yale.edu, scaz@cs.yale.edu
tel: +1 (203) 432-1219
fax: +1 (203) 432-0593

ABSTRACT

By learning a range of possible times over which the effect of an action can take place, a robot can reason more effectively about causal and contingent relationships in the world. An algorithm is presented for learning the interval $[t1_{min}, t1_{max}]$ of possible times during which a response to an action can take place. The algorithm was implemented on a physical robot for the domains of visual self-recognition and auditory social-partner recognition. The environment model assumes that natural environments generate Poisson distributions of random events at all scales. A linear-time algorithm called *Poisson threshold learning* can generate a threshold T that provides an arbitrarily small rate of background events $\lambda(T)$ if such a threshold exists for the specified error rate.

Keywords

Contingency, poisson threshold learning, developmental robotics, causality, self-recognition, social partner recognition

1. INTRODUCTION

The identification of self-generated sensory information and the identification of social partners are both primarily problems of causality. If a motor command reliably causes a certain response in the visual field, a robot may correctly assume that it is seeing itself move. Similarly, social partner identification is not simply a matter of finding faces in the visual field; a judgment about whether a person is interacting is a judgment about causality.

To reason about causality, it is not enough to look at the order of events; one must also take into account when those events occurred. If one sends a letter in the mail and receives a letter from the recipient the very same day, it probably isn't a response. The various nodes running in parallel on a modern robot architecture must contend with exactly this problem; motor responses and visual or audio feedback may not be immediate, and assuming that they are can lead to mistakes. The round trip time between command and response may be quite large, for instance, in the case of a robot being controlled by a rack of computers that is not physically close to the robot's sensors and end-effectors.

Identifying a chain of causality does not necessarily require logical reasoning. Humans appear to have the ability to detect contingent responses from an early age. Infants at 12 months are willing to follow the 'gaze' of a faceless object that emits beeps and flashes contingent upon the infant's own behavior, but will not follow the object's gaze if the beeps and flashes are emitted at random [1]. Even 6-to-12-week-old infants become agitated if a closed-circuit video feed of their mothers interacting with them is replaced with a non-contingent recording [2]. Experiments with adults show that the illusion of animacy conferred by an object acting without visible cause is largely perceptual and automatic [3], suggesting that these decisions are fairly low-level, and therefore potentially easy to emulate.

Learning a window of acceptable times for an event is a very tractable machine learning problem. The problem of learning the minimum and maximum times between a cause and a particular effect is essentially that of learning an interval $[a, b]$ on the real number line. It is straightforward to prove that the Vapnik-Chervonenkis dimension, or VC-dimension, of this problem is 2, and that therefore the number of examples required to learn this interval with probability at least $1 - \delta$ and error at most ϵ is at most $\frac{13}{\epsilon}(2 \ln \frac{1}{\epsilon} + \ln \frac{1}{\delta})$ [4]. This result holds regardless of the distribution of the data, which is useful for dealing with potentially unpredictable events such as the actions of a human being. In a noiseless environment, once this many training examples have been presented, events falling outside the time window can be confidently discarded as not belonging to the class in question – whether the domain is self-generated motion or social responses.

In the real world, however, the robot is not provided with the clear positive and negative examples of an ideal theoretical environment. Instead, it is presented with streams of real-valued sensor data, with higher values that may reflect either a contingent reaction to the robot's actions, unrelated events occurring in the environment, or just sensor noise. If it merely finds the minimum and maximum of the observed values, it runs a high risk of choosing events generated by the environment, instead of its own actions, as extrema.

To solve this problem, this paper presents a method of adaptively setting a threshold on the input channel to guarantee a low rate of irrelevant threshold events. The method, called *Poisson threshold learning*, relies on the assumption that the onsets of environmental events are largely uncorrelated with each other to produce a model that generates a sensor threshold with an arbitrarily low rate of background noise. As long as the events of interest produce sensor values that are above this threshold, the model can successfully learn a time window for detecting those events contingently by discarding the expected number of false positives. Moreover, the algorithm for setting the threshold runs in time linear with the sample size, and is optimal given the Poisson assumptions described above.

This does not completely solve the problem of positively identifying instances of the class, however. While all events of class C may occur within a specific time window t, this does not logically imply that all events within time window t belong to class C. (This problem of 'affirming the consequent' is common when applying simple property detectors, such as using color to judge whether an image contains skin.) Whether this is a problem in practice depends on the size of the time window. If the domain is self-generated motion, then one may hope that the responses are quick enough and reliable enough to make the probability of other events occurring within the window very small. If the domain is highly unpredictable, as it is in identifying social responses, then it may be necessary either to bring in information from other sensors or to attempt to reason probabilistically.

Whether these solutions work in practice is an empirical question – one addressed in this paper for the domains of self-recognition and social partner identification. Other work from this lab has shown that a robot can learn to match the words 'I' and 'you' to the properties of 'speaker' and 'listener' by watching a game of catch in which the two participants comment on the action [5]. The studies presented here hope to associate with the robot's first-person understanding of these terms. By learning the grounded meanings of these terms, robots may become better able to produce and understand social behavior.

Though the Poisson threshold learning approach and the unification of self-recognition and social response detection are presented here for the first time, our approach is informed by previous approaches to these problems. Self-identification in [6] was performed by cross-correlating the repetitive motion of the robot's hand with proprioceptive feedback from the motors, identifying the zero-crossings of each. This method of using motion feedback is similar to the approach presented here for the visual domain. Cohen et al. [7] used Allen's six possible relationships between fluents [8], instead of the specific time windows, to characterize the effects of their robot's actions. Other recent approaches for learning about the self have included statistically reasoning about invariant visual properties [9] and mapping visual feedback to tactile feedback [10]. Much less work has been done on learning to detect social contingency, but [11] presented a robot programmed to do so; the probabilities it used were calculated using a supercomputer that ran a dynamic programming algorithm on human-collected experimental data. Finally, this work follows the general paradigm of a robot learning about its environment with very little preprogrammed knowledge – a methodology first proposed by Turing [12] and currently known as 'autonomous mental development' [13], 'epigenetic robotics' [14], or 'developmental robotics' (e.g., this issue of *Connection Science*).

The system presented here for detecting self-generated feedback in both the visual and auditory domains has been implented on Nico, a humanoid robot that has the same arm and head kinematics as a one-year-old infant. (Some of the work presented here for the visual domain first appeared in [15], as well as in the 2005 AAAI Spring Symposium on developmental robotics.) In the auditory domain, the system has been extended to handle the detection of social responses as well. Data is also presented on the resulting response time distributions, so as to inform the design of Bayesian algorithms that could decide whether events are self-generated, socially contingent, or neither.

2. POISSON THRESHOLD LEARNING

The general method for learning acceptable windows of contingency is as follows.

Assume that the robot is attempting to learn via a noisy sensor S in a complicated environment E. Given a module M that issues a command to an actuator at time 0, the onset of the feedback to the module M occurs at some time given by an unknown probability distribution $P(t)$. The time window $\{[t_{1_{min}}, t_{1_{max}}] : P(t < t_{1_{min}} \cup t > t_{1_{max}}) < \epsilon_1\}$ is the interval during which the onset of self-generated feedback is very likely (with probability $(1 - \epsilon)$).

In a noiseless environment, this would simply be a learning problem with a Vapnik-Chervonenkis dimension of 2, since any collection of two points cannot be inconsistently labeled, but a dataset of 3 exists that is impossible to consistently label – the case of two positive examples on either side of a negative example. (See [16] for a full explanation.) This would entail, by a theorem in [4], that $\frac{13}{\epsilon}(2 \ln \frac{1}{\epsilon} + \ln \frac{1}{\delta})$ trials are sufficient to learn the correct window with probability at least $(1 - \delta)$ and error at most ϵ – a number polynomial in $1/\epsilon$ and $1/\delta$. Such a problem is said to be 'PAC-learnable', where PAC stands for 'probably approximately correct' [16]. However, here a sensor does not detect infallible Booleans of event/non-event, but some set of numbers $\{(s_1, \ldots, s_n) : s_i \in \Re\}$, with only the vague guarantee that during an interesting event, at least one of the s_i will be 'large'. (In the case of visual data, this 'sensor' may actually be a vision-processing module, such as a motion detector.) In addition, the environment and the sensor's own noise occasionally produce high values for the s_i that are at least as high as those produced by self-generated action. To learn the timing of its self-feedback, it seems as if the robot must first distinguish between its own actions and sensory data generated by noise and the environment. The reader may well wonder, then, how is this is to become a method of self-identification!

The answer lies in the environment model. If the model assumes that the likelihood of an event occurring within a time interval of length t is the same probability p regardless of when the measurement is taken, and that these events are stochastically independent of each other, then taking the limit as $t \to 0$ results in the Poisson distribution:

$$p(k; \lambda t) = e^{\lambda t} \frac{(\lambda t)^k}{k!} \quad (1)$$

where $p(k; \lambda t)$ is the probability of k events occurring within time t, given an event density of λ. In particular, $p(0; \lambda t) = e^{-\lambda t}$, and therefore the probability of at least one event happening within a time interval t is $1 - e^{-\lambda t}$. Because of its relatively easy to achieve prerequisites, the Poisson distribution has been observed in many natural domains, including radioactive disintegrations, flying bomb hits on London during World War II, chromosome interchanges in cells, wrong numbers, and distributions of bacteria on petri plates [17].

The assumption of stochastic independence must deal with the fact that natural events are likely to occur over time, and thus will produce values for the s_i that are potentially correlated with each other. The model presented here circumvents this problem by only dealing with the *onsets* of events, and assumes that events can be tracked across their durations.

An event takes place at time t if at least one value of s_i has risen above the threshold T, and that value was not generated by an already tracked event. The event density $\lambda(T)$ is therefore a function of this threshold. The results that follow require that $\lambda(T)$ is nonincreasing, with higher sensor values no denser than lower values.

The algorithms seeks a threshold T that sets the probability of an environmental event occurring between the self-initiated action and the true sensory feedback to be less than some small number ϵ. (This ϵ is not to be confused with the ϵ mentioned in the discussion of the PAC-learning model.) Using the Poisson model of the environment and an estimate t_0 of the true time delay gives the equation:

$$1 - e^{-\lambda(T)t_0} < \epsilon \quad (2)$$

Rearranging the terms to find $\lambda(T)$ gives:

$$\lambda(T) < \frac{-\ln(1-\epsilon)}{t_0} \quad (3)$$

In any sample of N sensor readings from S taken over t seconds, $Nt\lambda(T)$ readings are expected to have an event above the threshold T. This suggests the following linear time algorithm for finding an appropriate threshold, called *Poisson threshold learning*:

$(S_0, S_1, \ldots, S_N) = \mathrm{SAMPLE}(N, S, t)$
$M = (max(S_0), max(S_1), \ldots, max(S_N))$
$T = \mathrm{SELECT}(N - \lfloor \frac{-N \ln(1-\epsilon)}{t_0 t} \rfloor, M)$

Here, $\mathrm{SAMPLE}(N, S, t)$ returns N samples of the value sets returned by the sensor S over t seconds. Since a sensor set exceeds the threshold if and only if its maximum value exceeds the threshold, only the maximum value for each sensor set must be retained. This can be done online during sampling, to reduce memory requirements. On this reduced data set, the algorithm calls $\mathrm{SELECT}(i, M)$: a linear time algorithm described in [18] for finding the ith smallest element of a set without sorting. Here, it guarantees that only $N\lambda(T)$ values exceed the threshold. Because every operation presented here is linear time, the algorithm runs in $O(N + n)$ time, where N is the number of samples and n is the number of values in each sample.

Once the environment has been sampled and the threshold calculated, the robot can begin to act. For N trial movements, the system finds the length of time between sending a command to the actuator and sensing a value that exceeds the sensory threshold. Once this data has been collected, the system discards the $\lfloor 1 - e^{-\lambda(T)t_{1_{min}}} \rfloor N$ smallest values from t_1, since these are the expected number of false positive samples. The maximum and minimum of the remaining values become $[t_{1_{min}}, t_{1_{max}}]$.

Waiting between steps ensures that the robot does not attempt to act if the environment has suddenly changed; presumably, an environment that once had a reasonable λ will eventually quiet down again. Finally, the algorithm discards the $\lfloor 1 - e^{-\lambda(T)t_{1_{min}}} \rfloor N$ environmentally generated events that are expected to occur between sending a command and $t_{1_{min}}$, to obtain a closer estimate of the true window.

Detection of a social response is similar, except that the timing is measured from the offset of a self-generated event. Learning the social window $[t_{2_{min}}, t_{2_{max}}]$ therefore requires both identifying and tracking a self-generated event with some success. This implies that at least some self-learning must occur before social learning. In a real environment, social learning should only take place in the presence of some additional cue that signals the presence of a social partner, such as a forward-looking face; however, no such condition was required in the experiment presented below. The inclusion of a secondary sensor reading for verification would improve performance.

3. THE ROBOT PLATFORM

Nico is a humanoid robot built to serve as a testbed for models of infant learning (Fig. 1). Nico currently has a single arm with six degrees of freedom, while its head-neck assembly has an additional seven degrees of freedom.

Four miniature CCD cameras provide streaming video in wide and narrow fields of view for each of Nico's eyes; for simplicity, the following describes the path of a single image as it passes through the image processing pipeline. A frame grabber operating at 29-30 Hz receives a 320 × 240 image from one of the cameras. Once lens distortion is removed using a fixed lookup table, the image is passed to a motion detection module, which finds the absolute difference between the grayscales of the current image and the previous one. This difference image is smoothed and thresholded with an adjustable threshold T_m and passed to a pre-attentive vision module. There, large regions of thresholded motion are joined using a region-growing algorithm, which creates bounding box objects for each region. These events are then passed on to a tracking module before finally reaching the learning algorithm, at a frame rate of roughly 25Hz. Fig. 2

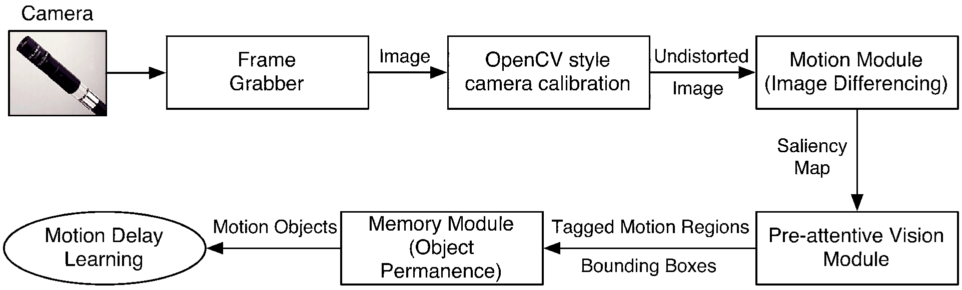

Figure 2: The visual preprocessing pipeline.

Figure 1: The robot, Nico. Nico is an upper-torso robot with a 6-DOF arm (left) and a 7-DOF head. Its four CCD cameras are located on either side of the bottom of Nico's 'forehead' cylinder (top, left).

provides an overview of the vision pipeline.

In the audio domain, Nico uses a two-channel microphone connected to a Sound Blaster Live card for audio recording. Data was sampled at 44100Hz on each channel, and was sampled again to provide the raw energy of each 100ms segment as input to the learning module. Because of the instability of the Sound Blaster Live drivers on QNX, sound production was performed by a Creative CT5880 card. The resulting sound was issued through normal PC speakers.

Each module ran on one of 16 processors running the QNX Neutrino real-time operating system. The nodes communicated over a 100Mbit switch, with each module providing its output to the other modules via shared memory in a concurrency-safe fashion.

4. SELF-RECOGNITION IN THE VISUAL DOMAIN

4.1 Methodology

Since the data from this portion of the paper were collected before the development of the Poisson threshold learning algorithm, T was set manually to provide a low $\lambda(T)$. An estimate of $\epsilon = 0.05$ was used to discard values after learning, but that was not based on a principled estimate of the underlying density $\lambda(T)$.

A series of 40 random poses was generated for Nico's arm. Nico would then wait to ensure that the background rate of new bounding boxes was acceptably low; in practice, this meant waiting until there was no detected motion for several seconds. Nico then iterated through the poses, calculating the delay in its actions t_1 as described above by measuring the time from sending the motor command to receiving the first untracked motion bounding box. After each action, Nico returned briefly to monitoring the background rate of bounding boxes to ensure that it had not significantly changed. Finally, at the end of the trials, two outlying values were dropped (since $\epsilon N = 2$) and the remaining minimum and maximum were designated $[t_{1_{min}}, t_{1_{max}}]$. The resulting window could then be used to filter motion bounding boxes for self-generated motion (Fig. 3).

4.2 Results and analysis

A subsequent 60 movement test run using the learned window resulted in only 3 movements that did not receive the self-generated label – exactly the 95% accuracy expected

Figure 3: An image from one of Nico's cameras. A bounding box around Nico's arm indicates that its motion was identified as self-generated using the t_1 time window.

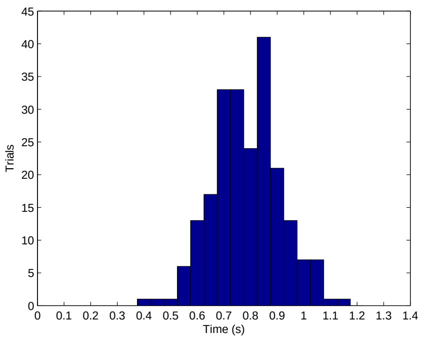

Figure 4: Distribution of times elapsed between sending a random motor command to the robot's arm and the perception of the resultant motion after the visual image has passed through the whole processing pipeline. The robot iterated through ten randomly selected motions 22 times.

from the Poisson threshold learning parameters. The learning thus took much less time than the theoretical upper bound described in the introduction, despite the relative variability of the response times (Fig. 4).

Under normally placid laboratory conditions, it was difficult to assess how well this system would perform in a busier natural environment Therefore, to test the limits of the system, an independent subject was asked to mislead the robot by attempting to move when it did. If any part of the subject's thresholded motion fell within the time window, it counted as a false positive, even if there was already some bounding box with the self-label. Despite this demanding adversarial setup, the widely variable t_1 times, and the relative simplicity of the method, the robot falsely labeled the subject's movements only about half the time (33/60 movements).

Since the data shows that the self-feedback times form a quite well-behaved normal distribution (Fig. 4), it is presumably possible to implement a Bayesian decision procedure for whether an event was more likely to come from the learned normal distribution for the self or the estimated Poisson

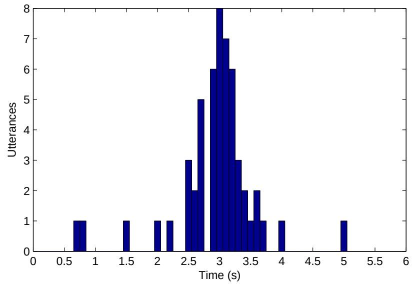

Figure 5: The distribution of social response times for humans interacting with the robot. The times are measured from when the robot sent the command to vocalize; the vocalization ended between 2 and 2.5 seconds later.

distribution of environmentally generated events. Such an implementation would be potentially useful and straightforward to implement, but it is beyond the scope of the current paper.

5. SELF AND SOCIAL RECOGNITION IN THE AUDITORY DOMAIN

5.1 Methodology

Nico's lack of text-to-speech software required some creativity in its production of social-response eliciting actions. To this end, the robot played a selection of sound files of the character R2D2 from the film *Star Wars*: a character known for its emotive – if inscrutable – beeps, chirps, and whistles. A background threshold was set using the threshold algorithm described earlier, with 10 seconds of samples recorded at 44.1 kHz. The parameter ϵ, characterizing the false positive rate, was set to 0.001, and the conservative guess t_0 was set to 2 seconds.

After setting the threshold, the robot went through an exploratory session alone in which it learned the delay t_1 and the duration u_i for each sound. Unlike the motor-motion loop, this loop was highly reliable, with all measurements for the same sound falling within 100ms of each other.

Two subjects were then instructed to respond to these sounds as if the robot were engaging them in intelligible conversation. The robot waited for a delay $t_{1_{max}}$ plus the learned length of each utterance, then measured the nearest gap in above-threshold values to find the length of delay between its self-generated sound and the subject's response. The exchanges continued for 30 utterances per subject.

5.2 Results and analysis

The resulting distribution of social response times is shown in Figure 5. Seven responses were not detected, and thus are not shown. In fact, the threshold was automatically set so high (because of some speech in the background) that the robot's own vocalizations did not pass the learned background threshold. The times were thus measured from the time that it sent the command to vocalize.

Even when the robot was in a noisy environment, and not all robot-directed speech surpassed the automatically set threshold, it still found a normally distributed response time for the humans. This means that it should be possible to distinguish between socially contingent responses and other noises from the environment that pass threshold, such as speech not directed at the robot, using the timing information from the response.

That the social responses should obey such a narrow distribution may be somewhat surprising, given the large amount of freedom subjects had in their responses. One may suspect that the subjects were simply responding to the robot non-socially, simply to get the experiment over with. This was not the case; the subjects faithfully responded as if Nico were attempting to communicate, though in some cases they protested to the robot that they couldn't understand what it was trying to say.

Tracking of the auditory events was performed only by following the raw sound intensity, but this approach had trouble distinguishing the pause after the robot had finished its utterance from the pauses within its own utterance. In fact, simply waiting for the first gap in above-threshold sound originally produced what appeared to be a negative exponential distribution on social response times, until it was found that some of the detected "responses" were actually the last few beeps of Nico's utterances. When the robot's behavior was changed so that it would always wait for the learned duration of an utterance before waiting for a response, the more accurate distribution on social response times was found. Because the auditory production pipeline was much more reliable than the motor pipeline, this approach performed adequately; however, in less predictable systems, the spectrogram of the auditory signal should probably be used to generate more reliable tracking. Without adequate tracking of the robot's own feedback, accurate social detection is unlikely.

6. APPLICATIONS

The method presented here for self-recognition makes very few assumptions about the behavior, appearance, and kinematics of the robot. Instead of requiring a specific behavior for finding the self, as in [6], these self-recognition and social modules can easily piggyback on top of behaviors performed for some other purpose, or for no purpose except random exploration. The lack of assumptions about the robot's appearance is also useful in case that appearance changes – for example, if the robot were to wear a glove, or if it were to become damaged in a way that changed its kinematics. A low-level detector for contingent motion may help the robot quickly adapt to new tools and new situations.

In particular, the ability to identify when a visual event is self-generated is key to the test of intelligence known as the 'mirror test' [19]. Briefly, the ability to recognize oneself in the mirror has long been used (possibly without much good justification) as proof of 'self-awareness' for two-year-old infants [20], chimpanzees [21], and dolphins [22]. For whatever reason, the ability to recognize the self in the mirror is not common in the animal kingdom, and seems to only be present in species that are 'intelligent' by other standards as well. Though setting the 'mirror test' as a benchmark for robot intelligence would inevitably lead to some solutions that utterly miss the point, as has happened with Turing's imitation game, it may be useful to consider why this is such a rarified ability in the animal kingdom. It may be that the ability to recognize new physical extensions or transformations of the self is a prerequisite for other important aspects of embodied intelligence. A robot that is able to identify its own motion in a mirror (Figure 6), while still far from 'passing' the mirror test by any definition, is a tentative step toward this kind of adaptive intelligence.

Figure 6: The robot identifies the self-generated motion of its reflection. The shot is from one of Nico's wide-angle cameras; Nico's reflection (center) is moving its hand as Nico moves a real arm in the foreground (bottom right). Because the motion occurs within the the learned time interval $[t_{1_{min}}, t_{1_{max}}]$, the motion is tagged as potentially self-generated (highlights).

In the social domain, the most obvious application of the social time window is to detect potential social partners that are not within the field of view. A potentially more interesting application is that this Boolean feature corresponds well to the semantic content of the word 'you.' In related experiments, Nico has learned that the word 'I' refers to the speaker and 'you' refers to the listener [5]. To move to first-person usage, the robot will need properties grounded in its own experience to which those words can be anchored. Ongoing work seeks to ground these terms using the properties described in this paper.

7. CONCLUSIONS

The work presented here has made three novel contributions.

The first contribution is a novel framework for learning contingent causal chains in noisy environments and arbitrary domains. Given a relatively innocuous set of assumptions about the environment and the nature of the sensory data, it is possible to find a threshold that reduces the error in the possible time windows $[t_{1_{min}}, t_{1_{max}}]$ to an arbitrarily low value in linear time. Using principles of PAC-learning [4], this number of samples is expected to have an upper bound that is polynomial regardless of the distribution.

Second, this framework was applied to the problem of detecting the robot's self-generated feedback in the environment. The number of samples required to achieve an error rate ϵ on the true positives was shown to be much smaller in practice than the upper bound suggested by the VC dimension of the problem [4]. The distribution of the self-generated motion is quite clearly Gaussian, allowing the substitution of a probabilistic model for the Boolean attribute model given here. Nevertheless, real robotic systems should probably not use the learned time window as a sole determiner of self-identification, but as a filter to eliminate those actions that are not self-generated, or as a cue to investigate responses further.

Third, the framework was applied to the problem of detecting contingent social responses. The data suggests that the human social responses are well-modeled by a normal distribution, and thus should be learnable for the purposes of self-other discrimination and turn-taking.

But much work still lies ahead in this area. The utility of employing Bayesian methods in the time window domain has yet to be determined. Poisson threshold learning has yet to see very many field tests. A proof of how many samples are needed to bound the error of the $\lambda(T)$ estimation would be useful, as would a more formal proof that the problem of learning real-valued intervals in a noisy environment is learnable in polynomial time. How reasonable is the assumption of uncorrelated event onset times in the real world? How should the environment be sampled to best preserve this assumption, while still taking linear time to calculate the threshold? And is temporal contingency a useful property to associate with the words 'I' and 'You'? These questions are the subject of ongoing research.

Acknowledgments

Support for this work was provided by a National Science Foundation CAREER award (#0238334). Some parts of the architecture used in this work was constructed under NSF grants #0205542 (ITR: A Framework for Rapid Development of Reliable Robotics Software) and #0209122 (ITR: Dance, a Programming Language for the Control of Humanoid Robots) and from the DARPA CALO/SRI project.

The authors would also like to thank Liz Darbie, Frederick Shic, Elizabeth Kim, Marek Doniec, and Viksit Gaur for their volunteering as test subjects, and the attendees of the AAAI 2005 Spring Symposium on Developmental Robotics for their helpful comments. Thanks also to Philip Michel, who collaborated on on the authors' early self-recognition work [15], and to Ganghua Sun and Andrew Lovett, who coded several of the modules mentioned in this paper.

8. REFERENCES

[1] S. C. Johnson, "Whose gaze will infants follow? Features that elicit gaze-following in 12-month-olds." *Developmental Science*, vol. 1, pp. 233–238, 1998.

[2] J. Nadel, I. Carchon, C. Kervella, D. Marcelli, and D. Réserbat-Plantey, "Expectancies for social contingency in 2-month-olds." *Developmental Science*, vol. 2, no. 2, pp. 164–173, 1999.

[3] B. J. Scholl and P. D.Tremoulet, "Perceptual causality and animacy," *Trends in Cognitive Sciences*, vol. 4, no. 8, pp. 299–309, 2000.

[4] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth, "Learnability and the vapnik-chervonenkis dimension," *Journal of the Association for Computing Machinery*, vol. 36, no. 4, pp. 929–965, 1989.

[5] K. Gold and B. Scassellati, "Using context and sensory data to learn first and second person pronouns," in *Human-Robot Interaction*, 2006, to appear.

[6] P. Fitzpatrick, "From first contact to close encounters: a developmentally deep perceptual system for a humanoid robot," Ph.D. dissertation, MIT, 2003.

[7] P. R. Cohen, C. Sutton, and B. Burns, "Learning effects of robot actions using temporal associations," in *Proceedings of the 2nd International Conference on Development and Learning*, 2002.

[8] J. F. Allen, "An interval based representation of temporal knowledge," in *International Joint Conference on Artificial Intelligence*. Morgan Kaufmann Publishers, Inc., 1981, pp. 221–226.

[9] Y. Yoshikawa, K. Hosoda, and M. Asada, "Cross-anchoring for binding tactile and visual sensations via unique association through self-perception," in *Proceedings of the Fourth International Conference on Learning and Development*, 2004.

[10] Y. Yoshikawa, Y. Tsuji, K. Hosoda, and M. Asada, "Is it my body? Body extraction from uninterpreted sensory data based on the invariance of multiple sensory attributes," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems*, 2004.

[11] J. Movellan, "Infomax control as a model of real time behavior: Theory and application to the detection of social contingency," UCSD MPLab, Tech. Rep., 2005.

[12] A. Turing, "Computing machinery and intelligence," *Mind*, vol. 49, pp. 433–460, 1950.

[13] J. Weng, J. McClelland, A. Pentland, O. Sporns, I. Stockman, M. Sur, and E. Thelen, "Autonomous mental development by robots and animals," *Science*, vol. 291, no. 5504, pp. 599–600, 2000.

[14] J. Zlatev and C. Balkenius, "Introduction: Why 'epigenetic robotics'?" in *First International Workshop on Epigenetic Robotics*, vol. 85, 2001, pp. 1–4.

[15] P. Michel, K. Gold, and B. Scassellati, "Motion-based self-recognition," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sendai, Japan, 2004.

[16] T. M. Mitchell, *Machine Learning*. McGraw-Hill, 1997.

[17] W. Feller, *An Introduction to Probability Theory and its Applications, Volume I*, 3rd ed. Wiley, 1968.

[18] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. MIT Press, 2001.

[19] G. Gallup, J. Anderson, and D. Shillito, "The mirror test," in *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*, M. Bekoff, C. Allen, and G. Burghardt, Eds. MIT Press, 2002.

[20] B. Amsterdam, "Mirror self-image reactions before age two." *Developmental Psychobiology*, vol. 5, no. 4, pp. 297–305, 1972.

[21] G. Gallup, "Chimpanzees: self-recognition," *Science*, vol. 167, no. 914, pp. 86–87, 1970.

[22] D. Reiss and L. Marino, "Mirror self-recognition in the bottlenose dolphin: a case of cognitive convergence," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 98, no. 10, pp. 5937–5942, 2001.