SOLID robot perception

Nathan F. Lepora, Uriel Martinez-Hernandez, Tony J. Prescott

***Abstract*— In a series of papers, we have formalized an *active Bayesian perception* approach for robotics based on recent progress in understanding animal perception. A central aspect of this active perception approach is the inference of 'where' and 'what' objects are in the environment under uncertainty, termed Simultaneous Object Localization and IDentification (SOLID). Here we describe some of the details of this algorithm and approach, its extension to include reinforcement learning of the active control strategy and speed-accuracy balance, and how the methods connect with computational neuroscience.**

In a series of papers [1], [2], [3], [4], [5], [6], [7], [8], we have formalized a *Bayesian perception* approach for robotics based on recent progress in understanding animal perception. Our formalism extends naturally to active perception, by moving the sensor with a control strategy based on evidence received during decision making. Benefits of active Bayesian perception include: (i) robust perception in unstructured environments [5]; (ii) an order-of-magnitude improvement in acuity over passive methods [8]; (iii) a general framework for simultaneous object localization and identification (SOLID), or 'where' and 'what' [8]; and (iv) a formalism that naturally integrates with reinforcement learning so that both the active control strategy and the appropriate belief threshold can be tuned appropriately to the contextual situation [6]. Here we describe some of the details of this algorithm and approach, and how it connects with computational neuroscience.

The present approach of active Bayesian perception with reinforcement learning for Simultaneous Object Localization and IDentification is further motivated by recent progress in the neuroscience of human and animal perception over imperfect sensor information. Leading computational accounts involve the sequential accumulation of evidence to threshold, consistent with numerous psychological and electrophysiological experiments [9], [10]. Work in computational neuroscience also indicates these principles may relate to the macro-architecture of the brain, principally the basal ganglia and cortex [11], [12]. In particular, it has recently been suggested that the basal ganglia architecture appears configured for optimal decision making over multiple channels of sensory evidence under a wide range of assumptions about what the evidence represents. The correspondence between the basal ganglia anatomy and a network version of sequential analysis holds with generic probability distributions and evidence encoding log likelihoods, log likelihood ratios or the log odds, and is shown in Fig. 1. (The mapping may not be unique: others have been proposed in [11] and [13].)

This work was supported by EU Framework project EFAA (ICT-270490).
N. Lepora, U. Martinez-Hernandez and T. Prescott are with the Sheffield Center for Robotics, University of Sheffield, UK. Email: {n.lepora, uriel.martinez, t.j.prescott}@sheffield.ac.uk

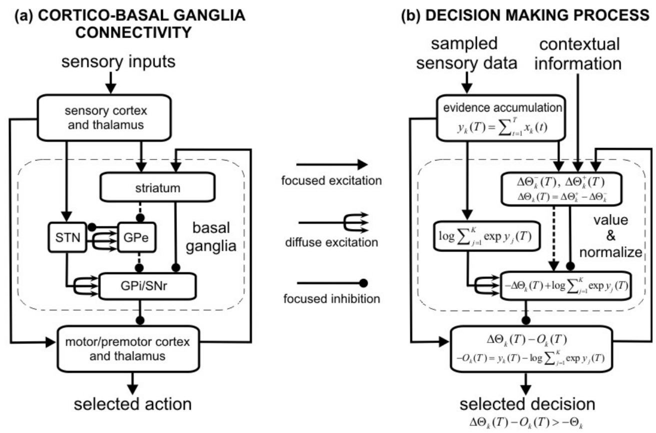

Fig. 1. Map between cortico-basal ganglia anatomy and sequential analysis. (a) A subset of the connections between the cortex and the basal ganglia nuclei. STN denotes the subthalamic nucleus, SNr the substantia nigra pars reticulata and GPe/GPi the globus pallidus external/internal segments. (b) The network version of Bayesian sequential analysis considered in [12], configured to map onto the cortico-basal ganglia anatomy. The striato-pallidal indirect pathway (dashed lines) is interpreted as conveying negative changes to decision thresholds $-\Delta\Theta_k^-$, which complement positive changes $\Delta\Theta_k^+$ along the direct striato-nigral pathway. Decisions are selected with upwards threshold-crossing in cortex, with the baseline and thresholds shifted consistently to ensure positive neuronal activity.

Sequential analysis methods for optimal decision making have also been applied recently to robot perception, focussing on robot touch [1]. A strength of the formalism is that it connects closely with leading work in neuroscience, allowing insights from animal perception to be transferred to robot perception. For example, these methods have enabled the first demonstration of hyperacuity in robot touch [4], giving perceptual acuity finer than the sensor resolution, as is common in animal perception. As discussed above, they also give robust perception in unstructured environments [5] in which there is uncertainty in both where and what objects are, as is also a central aspect of animal perception.

Our algorithm for active Bayesian perception with reinforcement learning is based on including a sensorimotor feedback loop in an optimal decision making method for passive perception derived from Bayesian sequential analysis [1]. Sequential analysis uses a free parameter, the decision threshold, to adjust the speed-accuracy tradeoff of the decisions [14]. Our control strategy for active perception also has another free parameter, the fixation point. We thus introduce reinforcement learning to set these two free parameters according to a reward function of the speed and accuracy of the decision outcome. Taken together, this

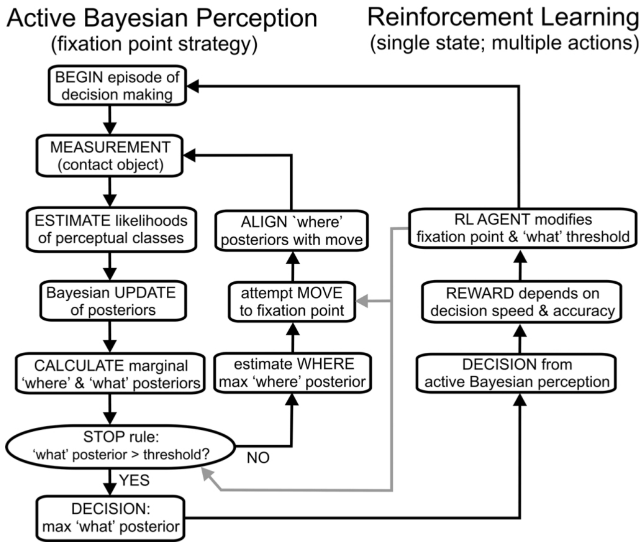

Fig. 2. Algorithm for active Bayesian perception with reinforcement learning. Active Bayesian perception (left) has a recursive Bayesian update to give the marginal 'where' and 'what' posteriors, allowing active control of the sensor position and decision termination at sufficient 'what' belief. Reinforcement learning (right) modifies the decision threshold and active control strategy based on reward information derived from the decisions.

results in the algorithm shown in Fig. 2. A key step in our combination of active perception and reinforcement learning is to interpret each active perception strategy (parameterized by the threshold and fixation point) as an action. In other work, we have viewed active Bayesian perception as passive Bayesian perception with a sensorimotor control loop [5], [8]. For passive Bayesian perception, the methods proposed here would optimize just the belief threshold. Thus, by including a sensorimotor loop, the active control strategy is then optimized with the same reinforcement mechanism.

In the future, we aim to draw further parallels between the cortico-basal ganglia network for perceptual decision making and algorithms for optimal robot perception, to use insights from animal perception to propose new methods for robot perception and interaction with their environments, while constraining the models of animal perception to be appropriate for robot applications.

References

[1] N.F. Lepora, C.W. Fox, M.H. Evans, M.E. Diamond, K. Gurney, and T.J. Prescott. Optimal decision-making in mammals: insights from a robot study of rodent texture discrimination. *Journal of The Royal Society Interface*, 9(72):1517–1528, 2012.

[2] N.F. Lepora, M. Evans, C.W. Fox, M.E. Diamond, K. Gurney, and T.J. Prescott. Naive bayes texture classification applied to whisker data from a moving robot. *Neural Networks (IJCNN), The 2010 International Joint Conference on*, pages 1–8, 2010.

[3] N.F. Lepora, J.C. Sullivan, B. Mitchinson, M. Pearson, K. Gurney, and T.J. Prescott. Brain-inspired bayesian perception for biomimetic robot touch. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 5111 –5116, 2012.

[4] N.F. Lepora, U. Martinez-Hernandez, H. Barron-Gonzalez, M. Evans, G. Metta, and T.J. Prescott. Embodied hyperacuity from bayesian perception: Shape and position discrimination with an icub fingertip sensor. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 4638–4643, 2012.

[5] N.F. Lepora, U. Martinez-Hernandez, and T.J. Prescott. Active touch for robust perception under position uncertainty. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 3005–3010, 2013.

[6] N.F. Lepora, U. Martinez-Hernandez, G. Pezzulo, and T.J. Prescott. Active bayesian perception and reinforcement learning. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, 2013.

[7] N.F. Lepora, U. Martinez-Hernandez, and T.J. Prescott. A SOLID case for active bayesian perception in robot touch. In *Biomimetic and Biohybrid Systems*, pages 154–166. Springer, 2013.

[8] N.F. Lepora, U. Martinez-Hernandez, and T.J. Prescott. Active bayesian perception for simultaneous object localization and identification. In *Robotics: Science and Systems*, 2013.

[9] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J.D. Cohen. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700, 2006.

[10] J.I. Gold and M.N. Shadlen. The neural basis of decision making. *Annual Reviews Neuroscience*, 30:535–574, 2007.

[11] R. Bogacz and K. Gurney. The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, 19(2):442–477, 2007.

[12] N.F. Lepora and K. Gurney. The basal ganglia optimize decision making over general perceptual hypotheses. *Neural Computation*, 24(11):2924–2945, 2012.

[13] R. Bogacz and T. Larsen. Integration of reinforcement learning and optimal decision-making theories of the basal ganglia. *Neural computation*, pages 1–35, 2011.

[14] A. Wald. *Sequential analysis*. John Wiley and Sons (NY), 1947.