An Embodied Biomimetic Model of Tactile Perception

Luke Burguete *† **Thom Griffith**‡ **Nathan Lepora** ‡

Abstract

Developing an artificial biomimetic model of tactile perception is a key goal for both practical applications in dexterous robotics and in our understanding of human touch. We present a novel, embodied model of tactile perception that mimics physical and computational features of the peripheral and central nervous systems. We deploy the model on a grating discrimination task, showing how integrating sensory evidence over time improves perceptual accuracy. We also provide a method for learning evidence thresholds which optimise the speed accuracy trade-off. Our model therefore provides a novel implementation of robotic tactile perception and sheds light on the computational features of this process in humans.

1 Introduction

Tactile perception is the process by which humans, animals, and, increasingly, artificial agents extract valuable information from their surroundings through the sense of touch. It plays a fundamental role in allowing us to navigate and interact with our environment, communicate, and make informed decisions. The immediacy of tactile perception, and our typical inattention to it, belies an intricate and remarkable ability that unites the peripheral and central nervous systems. During tactile perception, we continually collect, process, and combine ambiguous sensory information to seamlessly produce our somatosensory experience. Understanding the neural mechanisms of this process is crucial for our understanding of sensory processing and for designing intelligent systems that can effectively perceive and interact with the physical world. Here we address both of these aims by designing a model that is embodied and mimics structural and computational features of human tactile perception.

1.1 Transduction and decoding of tactile stimuli

Human tactile perception begins with sensory transduction of stimuli, whereby contact-induced deformation of the skin is turned into neural information. Artificial sensory transduction can be modelled by the TacTip, a soft biomimetic optical tactile sensor (figure 1a) [1]. Pins mounted on the inside of a soft latex tip mimic the morphology of human glabrous (hairless) skin and mechanically amplify skin deformation (figure 1b, c). White markers on the pin-tips are tracked by an internal camera. Marker displacement provides an analogy to the neural information encoded by slowly-adapting Merkel cell mechanoreceptors (SA-I) (figure 1d).

Stimulus discrimination is a key feature of sensory perception. Johnson and Phillips [2] use a two-interval, binary, forced-choice grating discrimination task to assess tactile discrimination in humans. In it, participants sequentially tap two square-wave grating stimuli (figure 3) and determine whether their orientation is the same or different. Perceptual acuity depends on the period of the grating used, with easier discrimination of larger gratings due to more reliable tactile information (figure 1).

Pestell et al. [3] proposed an embodied model which combines stimulus transduction and decoding to infer performance on the grating discrimination task described. Transduction is modelled by the

*Corresponding author (l.burguete@bristol.ac.uk)
†University of Bristol, School of Physiology, Pharmacology and Neuroscience
‡University of Bristol, School of Engineering Mathematics and Technology

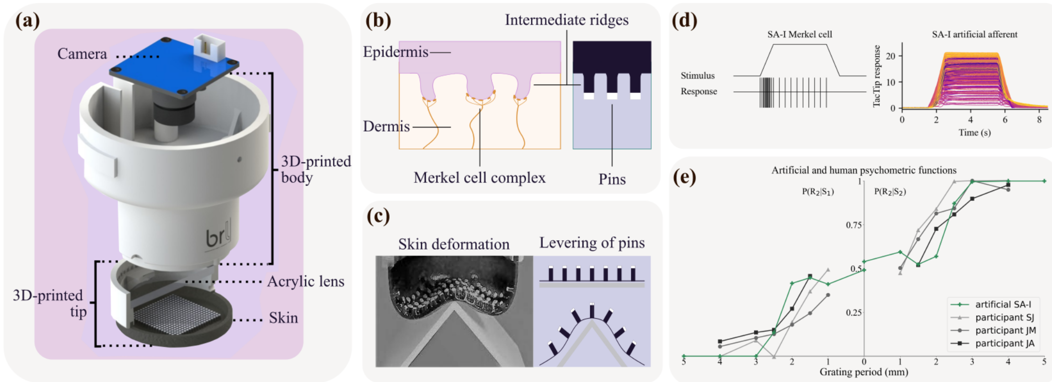

Figure 1: (a) Diagram of the TacTip. (b) Morphology of glabrous skin (left) and TacTip (right), pins correspond to intermediate epidermal ridges. (c) Contact-induced deformation of TacTip skin (left) causes levering of pins (right) which amplifies the mechanical signal. (d) Stimulus-response relationship for slow, sustained contact of Merkel SA-I mechanoreceptor cell (left) and artificial SA-I TacTip afferents (right) (reproduced from figure 4 in [3]). (e) Psychophysical responses to the grating discrimination task. Human subjects from [2] are plotted in grey and the embodied model of Pestell et al. [3] is shown in green (reproduced from their figure 8).

TacTip and decoding, which in humans is provided by the somatosensory cortex, is modelled by a convolutional neural network (CNN). When trained to predict grating orientation, this model can produce human-like responses to the grating discrimination task (figure 1e).

1.2 Integration of tactile evidence

The model of Pestell et al. [3] can be built upon by incorporating another key feature of perception: evidence integration over time. This integration process is essential whenever a sensor can only apprehend part of an object at a time or when perceptual evidence is noisy or ambiguous. This arguably describes all realistic situations requiring tactile perception.

Many kinds of evidence integration model, such as the sequential probability ratio test (SPRT) and the drift-diffusion model (DDM), are shown to be equivalent under reasonable assumptions [4; 5]. There is widespread evidence that humans and other mammals employ this class of algorithm in a range of perceptual and other decision-making tasks [6; 7; 8]. Furthermore, there is a strong link between the mammalian basal ganglia, a collection of subcortical brain structures involved in action selection and decision making, and the computational features of the SPRT [9; 10], including a precise mapping of the multiple-hypothesis SPRT algorithm onto the anatomy and function of the BG [11] (figure 2a) .

The SPRT decides between two hypotheses based on a single variable, the cumulative log-likelihood ratio, Z (see equation 1). This value encodes which hypothesis is more likely given the evidence. The key model parameter, the evidence threshold, θ_i, defines the point at which evidence accumulation stops and a decision is made (figure 2b). A larger threshold requires more evidence, which leads to a slower but more accurate decision. The evidence threshold therefore naturally defines a balance between speed and accuracy, which is an essential trade-off in any realistic decision process.

2 Methods

Grating discrimination The stimuli are square-wave gratings oriented at $\mp 45°$ relative to the midline of the sensor, labelled A and B respectively. At each timestep, stimuli are presented sequentially in state S_1 (AA) or S_2 (AB). After presentation, the model either re-samples the stimuli or gives response R_1 or R_2, indicating a perceptual decision that the state is S_1 or S_2, respectively.

Model description The full model schematic is shown in figure 3. We deploy this model on the grating discrimination task using the tactile images collected by Pestell et al. [3].

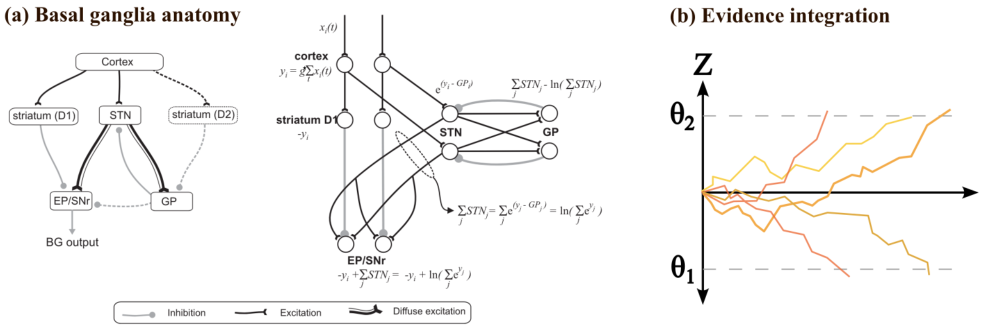

Figure 2: (a) Basal ganglia anatomy (left) and functional mapping to the SPRT (right) (reproduced from Bogacz and Gurney [11]). (b) Typical evidence trajectories to threshold for the SPRT.

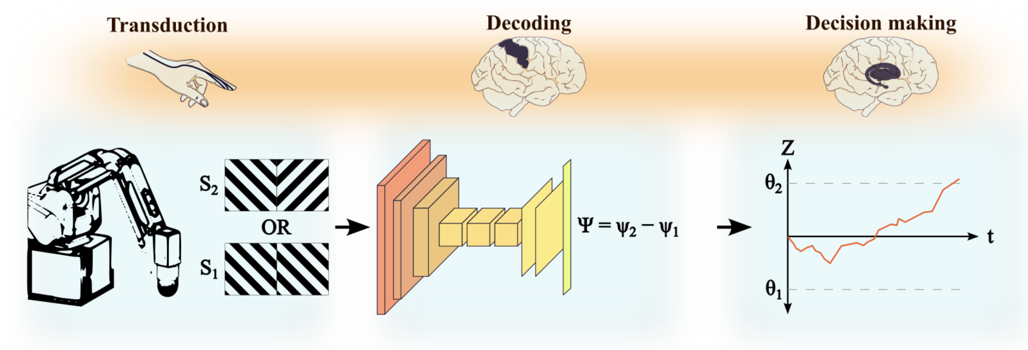

Figure 3: A schematic pipeline of tactile perception, elements of our model are shown below their corresponding regions in humans. Transduction (left, fingertip): tactile images are collected by a robotic arm- TacTip. Decoding (middle, somatosensory cortex): tactile images are decoded by a CNN which predicts the grating angle difference. Decision making (right, basal ganglia): predictions are integrated by the SPRT until evidence reaches the decision threshold (θ_i).

At each timestep t, a trained CNN (as described in [3]) predicts the angle of both grating presentations, $\psi_{1,t}$ and $\psi_{2,t}$. The evidence is defined as $\Psi_t = \psi_{2,t} - \psi_{1,t}$. The decision variable is defined as

$$Z_t = Z_{t-1} + \log \frac{p(\Psi_t|S_2)}{p(\Psi_t|S_1)} \tag{1}$$

where $p(\Psi_t|H)$ is the likelihood of the evidence given some hypothesis H. The threshold gap, δ, is defined such that the model terminates a decision process and gives response R_1 if $Z_t \leq -\delta/2$ and R_2 if $Z_t \geq \delta/2$. If $|Z_t| < \delta/2$, the model will continue to sample from the evidence unless $t = T_{max}$ when the model makes forced response R_1 if $Z_{T_{max}} < 0$ or R_2 if $Z_{T_{max}} > 0$.

The model receives reward

$$r = \xi W - T \tag{2}$$

where ξ is the cost ratio, T is decision time, and W is -1 for incorrect decisions and 0 otherwise. We use an implementation of REINFORCE with baseline [12] to optimise δ for a given value of ξ.

3 Results

The model performs the grating discrimination task with psychophysical responses that reflect human performance and become more accurate as the evidence threshold gap widens (figure 4a). We show the speed-accuracy relationship for different grating widths (figure 4b), demonstrating how increased decision time leads to higher accuracy. Grating widths $\geq$ 3mm are not shown because the perceptual model makes accurate decisions after a single sample, thus eliminating a speed-accuracy trade-off.

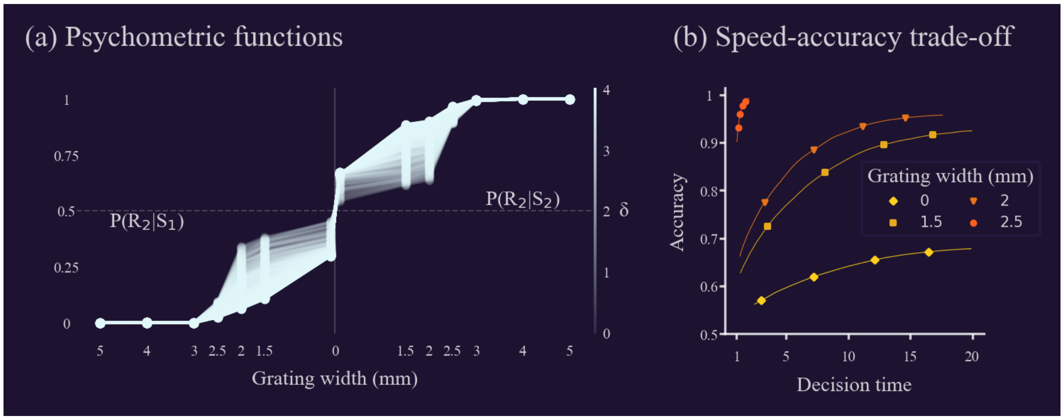

Figure 4: (a) Psychophysical responses to the grating discrimination task generated by our model. Accuracy improves as the evidence threshold gap increases (bolder lines). (b) Speed-accuracy curves for individual gratings (smoothed over 10 trials), showing the increase in accuracy with decision time.

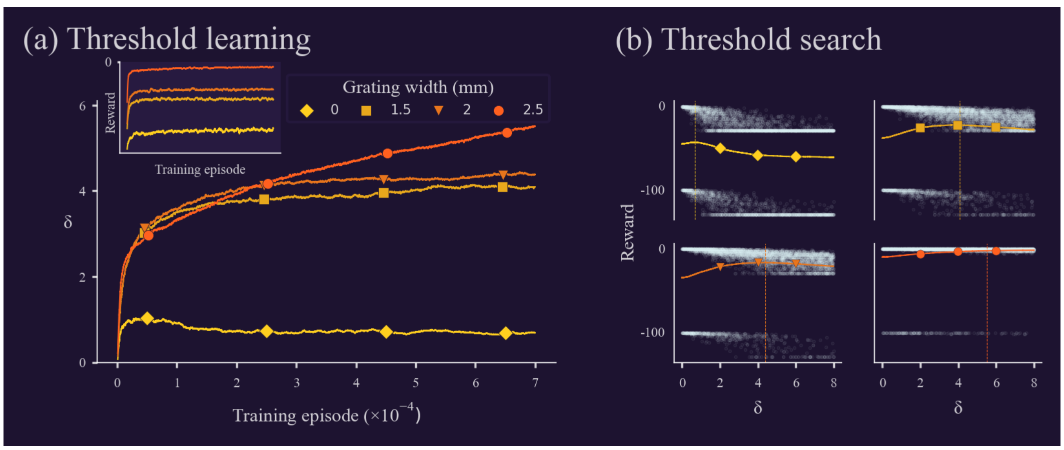

Figure 5: (a) Acquisition curves for threshold learning using REINFORCE (average 6-fold cross validation with 10 repeats per fold, smoothed over 100 episodes). The models converges (inset), and learns threshold values which approximate optimal thresholds found by exhaustive search. (b) Sampled rewards from 1.4×10^6 trials over the threshold range $\delta \in [0, 8]$, with the average reward curve plotted . Dashed lines show learned (mean) value from last 500 trials of REINFORCE.

The model learns a threshold value that optimises the reward rate for each grating (figure 5a). To validate this result, we performed an exhaustive search by sampling the reward obtained from 1.4×10^6 threshold values. The learned value approximates the optimal value of the average reward function (figure 5b). A very flat reward function (grating 2.5) means the threshold may continue to change after the reward value has converged (figure 5a, inset).

4 Conclusions

We have presented a novel, embodied model for robot tactile perception that is capable of accurately performing a well-studied human psychophysical grating discrimination task by integrating sensory evidence over time. Our model learns, by reinforcement, to balance speed with accuracy by optimising the evidence threshold parameter. This provides it with the capability to adapt to changes in both the environment and the internal goals of the agent. Furthermore, since the process of evidence integration here is agnostic to the source of evidence, our model has a wide range of potential applications in both tactile perception and other sensory modalities. Finally, the biomimetic features of our model provide an embodied testbed for hypotheses about human tactile perception, thus helping to clarify the computational features of this process.

References

[1] Nathan F. Lepora. Soft Biomimetic Optical Tactile Sensing With the TacTip: A Review. *IEEE Sensors Journal*, 21(19):21131–21143, October 2021. ISSN 1558-1748. doi: 10.1109/JSEN.2021.3100645. Conference Name: IEEE Sensors Journal.

[2] K. O. Johnson and J. R. Phillips. Tactile spatial resolution. I. Two-point discrimination, gap detection, grating resolution, and letter recognition. *Journal of Neurophysiology*, 46(6):1177–1192, December 1981. ISSN 0022-3077. doi: 10.1152/jn.1981.46.6.1177. Publisher: American Physiological Society.

[3] Nicholas Pestell, Thom Griffith, and Nathan F. Lepora. Artificial SA-I and RA-I afferents for tactile sensing of ridges and gratings. *Journal of The Royal Society Interface*, 19(189): 20210822, 2022. doi: 10.1098/rsif.2021.0822. Publisher: Royal Society.

[4] Rafal Bogacz, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D. Cohen. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113:700–765, 2006. ISSN 1939-1471. doi: 10.1037/0033-295X.113.4.700. Place: US Publisher: American Psychological Association.

[5] Sebastian Bitzer, Hame Park, Felix Blankenburg, and Stefan Kiebel. Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. *Frontiers in Human Neuroscience*, 8, 2014. ISSN 1662-5161.

[6] Joshua I. Gold and Michael N. Shadlen. Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5(1):10–16, January 2001. ISSN 13646613. doi: 10.1016/S1364-6613(00)01567-9.

[7] Joshua I. Gold and Michael N. Shadlen. The Neural Basis of Decision Making. *Annual Review of Neuroscience*, 30(1):535–574, 2007. doi: 10.1146/annurev.neuro.29.051605.113038. _eprint: https://doi.org/10.1146/annurev.neuro.29.051605.113038.

[8] Roger Ratcliff and Gail McKoon. The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural computation*, 20(4):873–922, April 2008. ISSN 0899-7667. doi: 10.1162/neco.2008.12-06-420.

[9] K. Gurney, T. J. Prescott, and P. Redgrave. A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biological Cybernetics*, 84(6):411–423, May 2001. ISSN 1432-0770. doi: 10.1007/PL00007985.

[10] Nathan F. Lepora and Kevin N. Gurney. The Basal Ganglia Optimize Decision Making over General Perceptual Hypotheses. *Neural Computation*, 24(11):2924–2945, November 2012. ISSN 0899-7667, 1530-888X. doi: 10.1162/NECO_a_00360.

[11] Rafal Bogacz and Kevin Gurney. The Basal Ganglia and Cortex Implement Optimal Decision Making Between Alternative Actions. *Neural Computation*, 19(2):442–477, February 2007. ISSN 0899-7667, 1530-888X. doi: 10.1162/neco.2007.19.2.442.

[12] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, May 1992. ISSN 1573-0565. doi: 10.1007/BF00992696.