ROSIE: RObotic ScIentist for Exploration

K. Bywaters[1], P. M. Furlong[2], and P. Boston[2]

[1]SETI Institute (Kbywaters@seti.org), [2]NASA Ames Research Center (padraig.m.furlong@nasa.gov; penelope.j.boston@nasa.gov)

Introduction: Explicitly encoding human judgment is something which has escaped the grasp of artificial intelligence researchers for decades. For example, the way humans conduct field work is hard to explicitly quantify in terms of cataloging and weighing critical environmental parameters. Human field scientists conducting biological studies of habitability might look at parameters such as light exposure, nutrient inputs, and water availability. Through years of experience field scientists collect and process this information, formulating hypotheses, leading to a "gut instinct" or intuition, often based merely on visual assessments, which they may not be consciously aware of. As such, it is often difficult for humans to communicate their actual objectives to robotic systems and thereby making the development of a trustworthy autonomous system challenging. We propose to extract an autonomous exploration policy for robots directly from the neurological activity of the scientists they are learning to mimic.

Rationale: The human brain is what enables the skillset that comprises "field research". To expand and execute more competent autonomous exploration the complex human skillset needs to be translated to our robotic proxies. Robots need to acquire the operational flexibility and the essence of the native and trained intelligence of a human field researcher.

Common approaches to autonomous exploration rely on information theoretic quantities to drive exploration, e.g. [1-3], with minimal specification of objectives from human scientists. Approaches to human-robot teaming for exploration use various methods to encode expert knowledge. Arora et al. [4] and Candela [5] rely on humans encoding their beliefs explicitly using graphical models. Somers et al. [6] uses imitation learning to estimate scientists' preferences from example trajectories. Jamieson et al. [7] demonstrates an improvement in autonomous exploration through a sample efficient mechanism for requesting teaching examples from remote users. These methods are effective, and appropriate when the tempo of human-robot interaction is low and the interaction has clearly defined boundaries. However, if the objective is to extract training data for field missions, interrupting operations to seek feedback can slow operational cadence and distract from the mission objectives. Hence, we propose a system that can be worn by humans on analogue missions to extract training data with no additional interaction required from the scientists.

Native intelligence or intuitive pattern recognition can be recorded in the neurological activity of the human field researcher and translated via the brain-machine interface. Monitoring neurological activity is necessary for two reasons: 1) it does not interfere with the scientific process that the scientists are performing and 2) eliminates introspective fallacies. Understanding human decision making is notoriously difficult [8] for a multitude of reasons including people lying about the factors they considered or not understanding the factors that influenced them [9, 10]. However, whether or not a scientist can verbally qualify and quantify their field experience is irrelevant, the data needed to assign values is contained in the scientist's behavioral and neurological responses [11]. Assuming the value of decisions can be extracted from neural activity, we frame the problem as a Reinforcement Learning problem, where the objective is to estimate the state-action value function over input images collected during field operations.

Concept: ROSIE (RObotic ScIentist for Exploration) is a novel approach to capturing the essential cognitive skills of field scientists through neural signals and embed that capability into a robotic explorer. ROSIE will be trained using neural data obtained from human field scientists while they explore field sites or engage in investigatory activities, learning judgement from human experience using neurological responses. The neurological data will be correlated with an array of data gathered simultaneously about the environment (i.e. what the scientist is observing, as well as data from sensors placed in the surroundings for temperature, light, etc.). Correlating these data will lead to the automation of exploration skill sets that will enable purely robotic exploration capable of addressing the critical scientific questions.

The premise of brain-machine interfaces is that neurological activity can be observed and interpreted. This ability has been credibly demonstrated in concrete activities like motor control [12], as well as more abstract functions, like language interpretation [13]. For example, invasive monitoring techniques have been used to allow monkeys to control wheelchairs with their minds. Non-invasive techniques (fMRI) have been used to identify words and concepts that humans are thinking about [12, 13], been used for controlling switches, [13], used to control a wheelchair by a human via EEG [14], and to predict human decision making [15 - 17]. Brain-machine interfaces have been used to monitor and interpret neurological activity; this work leverages these existing technologies to create a novel algorithm for controlling robots exploring and collecting data.

There are two major obstacles to designing a system like ROSIE. First, sets of rules consciously designed by scientists may be incomplete, due to the importance of scientists' unconscious knowledge, and brittle. Second, scientists' interaction does not have a clear end-point, they are continually interacting with the environment, while brakes in exploration occur, they are because of demands on the scientists' time outside of the exploration task.

Through the course of this work both critical points will be addressed. The first obstacle indicates that we should frame this problem as a machine learning problem, teaching ROSIE how to behave based on how scientists behave. Neural activity is an appropriate source of training information because it gives a measure of the value of potential science targets, rather than just the discrete information of selected/not selected that we would get out of behavior information alone. Monitoring neural activity also does not require the scientist to switch context out of the exploration task to record information for training data. The second obstacle drives us to treat the learning problem as a continuous reinforcement learning problem -- maximizing the average neurally-derived reward it gets for exploring, because there is no clear end-point where performance can be finally evaluated.

We propose to use a TD-learning algorithm [18] to train a network that maps input images to the estimated value of sampling actions, given the input image. The algorithm will learn a neural network to transform the images into the state-value function over the actions of {*sample*, *ignore*}. By learning this function, ROSIE will capture essential details that drive human decision making. Tools from explainable AI [19] can invert the learned value function, identifying features that motivate choices to engage with or ignore potential samples. By using these techniques to create a visual interpretation of the algorithm's judgement, it can be tested and assessed before launch. Avoiding reliance on scientists' introspective rules will enable robust robotic exploration missions. Learning directly from human neurological activity is a novel value function for value-guided judgement in data collection missions. By training robots to satisfy scientists' internal value functions we should develop robots that behave more predictably, increasing trust in deployed robots.

References:

[1] Arora, A., Furlong, P. M., Fitch, R., Sukkarieh, S., & Fong, T. (2019). Multi-modal active perception for information gathering in science missions. *Autonomous Robots*, *43*(7), 1827-1853.

[2] Wettergreen, D., Foil, G., Furlong, M., & Thompson, D. R. (2014). Science autonomy for rover subsurface exploration of the atacama desert. *AI Magazine*, *35*(4), 47-60.
[3] Thompson, D. R., Wettergreen, D. S., & Peralta, F. J. C. (2011). Autonomous science during large-scale robotic survey. *Journal of Field Robotics*, *28*(4), 542-564.
[4] Arora, A., Fitch, R., & Sukkarieh, S. (2017, September). An approach to autonomous science by modeling geological knowledge in a Bayesian framework. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 3803-3810). IEEE.
[5] Candela, Alberto, et al. "Planetary robotic exploration driven by science hypotheses for geologic mapping." *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017.
[6] Somers, T., & Hollinger, G. A. (2016). Human–robot planning and learning for marine data collection. *Autonomous Robots*, *40*(7), 1123-1137.
[7] Jamieson, S., How, J. P., & Girdhar, Y. (2020). Active Reward Learning for Co-Robotic Vision Based Exploration in Bandwidth Limited Environments. *arXiv preprint arXiv:2003.05016*.
[8] Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. *Journal of Legal Analysis*, *10*.
[9] Pronin, E. (2009). The introspection illusion. *Advances in experimental social psychology*, *41*, 1-67.
[10] White, P. A. (1988). Knowing more about what we can tell:'Introspective access' and causal report accuracy 10 years later. *British Journal of Psychology*, *79*(1), 13-45.
[11] Babiloni, C., Del Percio, C., Rossini, P. M., Marzano, N., Iacoboni, M., Infarinato, F.,et al. (2009). Judgment of actions in experts: a high-resolution EEG study in elite athletes. *Neuroimage*, *45*(2), 512-521.
[12] Rajangam, S., Tseng, P. H., Yin, A., Lebedev, M. A., & Nicolelis, M. A. (2015). Direct cortical control of primate whole-body navigation in a mobile robotic wheelchair. *arXiv preprint arXiv:1504.02496*.
[13] Mitchell, T. M., Hutchinson, R., Just, M. A., Niculescu, R. S., Pereira, F., & Wang, X. (2003). *In AMIA Annual Symposium Proceedings* (Vol. 2003, p. 465). American Medical Informatics Association.
[14] Al-Qaysi, Z. T., Zaidan, B. B., Zaidan, A. A., & Suzani, M. S. (2018). A review of disability EEG based wheelchair control system: Coherent taxonomy, open challenges and recommendations. *Computer methods and programs in biomedicine*, *164*, 221-237.
[15] Jollans, L., Whelan, R., Venables, L., Turnbull, O. H., Cella, M., & Dymond, S. (2017). Computational EEG modelling of decision making under ambiguity reveals spatio-temporal dynamics of outcome evaluation. *Behavioural brain research*, *321*, 28-35.
[16] Çakır, M. P., Çakar, T., Girişken, Y., & Demircioğlu, A. K. (2016, July). Neural correlates of purchasing decisions in an ecologically plausible shopping scenario with mobile fNIR technology. In *International Conference on Augmented Cognition* (pp. 135-146). Springer, Cham.
[17] Telpaz, A., Webb, R., & Levy, D. J. (2015). Using EEG to predict consumers' future choices. *Journal of Marketing Research*, *52*(4), 511-529.
[18] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
[19] Samek, W., Wiegand, T., & Müller, K. R. (2017). Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296*.