Human Geoscientist Objective Functions for Robot-Aided Field Data Collection Decisions

Shipeng Liu[1], Cristina G. Wilson[2,3], Bhaskar Krishnamachari[1] and Feifei Qian[*,1]

***Abstract*—Geoscientists can use field robots to explore ongoing global changes like desertification, the process whereby fertile lands change to desert. High spatiotemporal resolution data collected by robots can inform understanding of the complex causes of desertification, critical to preventing additional land loss and protecting biodiversity. A key challenge in developing more intelligent robots, that move beyond mobile data collection devices and begin to aid human experts with sampling decisions, is the lack of understanding of how scientists make such decisions and adapt strategies when presented with new information.**

In this study we examined the dynamic data collection decisions of 108 expert geoscientists using a simulated field scenario. Human data collection behaviors suggested two distinct objectives: an objective to improve information coverage, and an objective to verify beliefs about the hypothesis. We developed a highly-simplified quantitative decision model that allows the robot to predict potential human data collection locations based on the two observed human data collection objectives. Predictions from the simple model successfully captured sampling location choices for 70% of human expert scientists, and revealed a transition in objectives as the level of information increased. The findings will enable decision support algorithms that allow robotic teammates to infer experts' desired data collection strategy based on abstract scientific objectives, in the long-term supporting the development of cognitively-compatible robotic field assistants.

I. Using robots to study desertification and global climate interactions

Arid, semi-arid, and dry sub-humid areas – covering 46.2% of the global land area and home to approximately 3 billion people – are vulnerable to land degradation in the form of declines in soil quality, vegetation, water, or wildlife, collectively referred to as desertification [1], [2]. Anthropogenic climate change, in interaction with human activities like unsustainable land management practices, has exacerbated desertification in some dryland areas and risks from desertification are projected to increase with continued global warming [1]. Scientists are deploying robotic technology in desertification research to understand complex and rapidly changing sediment dynamics *in-situ*. For example, in prior work, we deployed a legged robot, RHex [3], to assist human geoscientists in field data collection across deserts in NM and CA [4], to reveal how environmental properties such as soil strength and vegetation density influence sediment transport and desertification processes. The robot provided human geoscientists with high spatiotemporal resolution data on leg-soil interactions *in-situ*, bringing the precision of laboratory experimentation to the field [4]–[6]. Access to this data allowed scientists to test hypotheses about soil strength in the field, update their beliefs, and dynamically adapt their data collection strategies to enable important scientific discoveries [7].

[1]Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA.

[2]Department of Psychology, Temple University, Philadelphia, PA, USA

[3]Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA

[*]Corresponding author. feifeiqi@usc.edu

II. Moving robot field assistants towards more intelligent teammates

Most state-of-the-art robots assisting in data collection – including the robot RHex used in our research [4], [7] – are used as mobile sensor suites, taking low-level command from humans to execute the navigation, sensing and sampling, while human experts bear the full burden of integrating and interpreting data for future data collection decision making. When this burden exceeds the processing capacity of the human mind, experts are likely to rely on mental shortcuts and rules-of-thumb (heuristics) [8], [9] and have trouble flexibly adapting thinking and behavior in response to new information [10]; leaving scientists vulnerable to decision biases that can result in missed scientific discoveries [11].

For robots to move beyond mobile sensor suites and begin to assist human scientists with high-level decision-making and sampling strategy adaptations, a better understanding of how scientists make data collection decisions is essential. By understanding 1) what the most important objectives that influence experts' data collection strategy are, and 2) how experts combine different objectives to select and adapt strategies in response to new data, a robot could begin to suggest sampling locations based on inferred combinations of objectives and then learn from human responses. In this manner, robots can become more intelligent teammates that can can take on increased responsibility in collaborative exploration, thereby freeing up the expert to engage in the type of abstract hypothetical thinking that the human mind excels at [12].

III. Assessing expert scientist data collection behavior in simulation

Here, we use a simulated data collection scenario to determine how expert geoscientists make spatiotemporal data collection decisions, and how search is adapted in response to real-time data (Fig. 1). The simulated data collection scenario is inspired from the real-world field data collection process in [5]. In the task, participants are provided with a hypothesis about the relationship between sediment strength and moisture along a sand dune (Fig. 1, H_{A1}), and asked to create a data collection strategy to test the hypothesis – a simulated

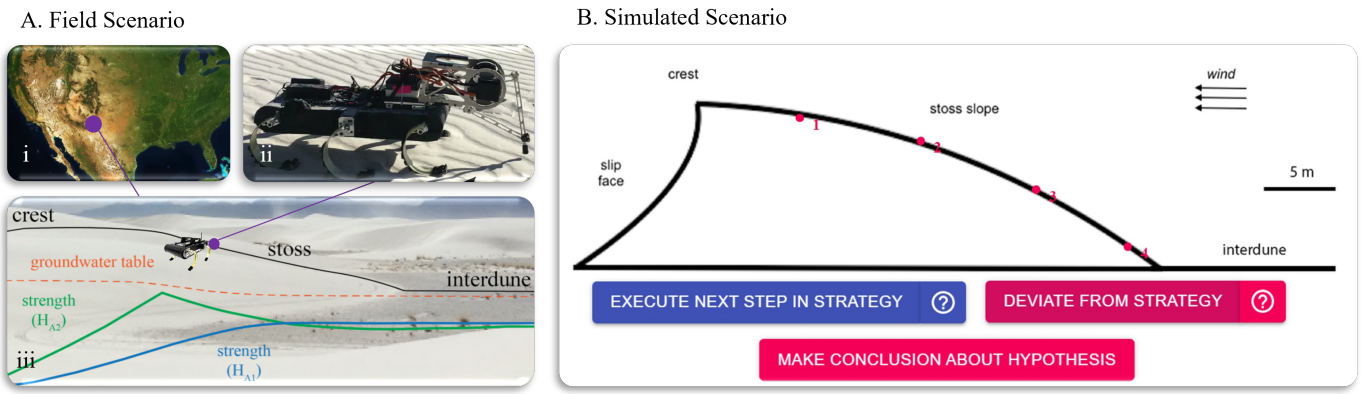

Fig. 1: Robot-aided *in-situ* sampling [5] inspired decision-making scenario. (A) Field sampling site at White Sands, NM, a dune field in the southwest of United States (i), where the RHex [3] robot (ii) assisted human scientists collect soil property measurements [4], [5] along a sand dune. Black line highlights the transect of a dune where we observe the largest gradient in soil properties (iii). At the crest of the dune, where the soil was driest because of its distance from the groundwater table (orange line), soil strength was expected to be low. As moisture increased on the stoss face moving towards the interdune, strength was expected to also increase before leveling of at the point of moisture saturation. This pattern of expected results, H_{A1}, is displayed in blue and is provided to participants in the simulated scenario. Through field work, geoscientists discovered that soil strength actually increases rapidly to its maximum as soil becomes slightly wet, and then decreases slightly as soil moisture becomes more saturated nearing the interdune area just before leveling off [5]. This alternative pattern of results, H_{A2}, is displayed in green. Participants in the simulated scenario were randomly assigned to sample from data sets supporting H_{A1} or H_{A2}. (B) The interactive data collection page of the web-based decision-making scenario (user/password: rhex), which was inspired from the robot-assisted field data collection scenario. Expert geoscientist participants select data collection locations on dune cross-section and measurements are provided in real-time.

robot executes the strategy and provides real-time data that participants can use to adapt their strategy. Previous work with an older version of the task showed expert geoscientists rely on simple data collection heuristics of equally spacing sampling locations and taking a consistent "magic" number of measurements at each location [7]. Using an updated version of the task, the current study sought to determine the objectives underlying these heuristics, and how changes in hypothesis beliefs (measured by subjective confidence) alter the weighting of objectives and corresponding data collection decisions.

IV. Experts' initial sampling strategy is exploration oriented and heuristic driven

We replicated equal spacing and magic number heuristics amongst a new group of 108 expert geoscientists: 94% of participants chose uniform location intervals (Fig. 2 A), and 85% of participants chose a constant number of samples at each location (Fig. 2 B). Furthermore, we found that the majority (60%) of the participants chose a constant number of samples, $n \in [3, 6]$, and uniform intervals, $\Delta l \in [2, 4]$. We propose two hypotheses to explain the observed heuristics: (i) a diminished information reward hypothesis to explain the magic number heuristic, and (ii) an information inference hypothesis to explain the equal spacing heuristic.

(i) Diminished information hypothesis. We hypothesize that with an increasing number of samples from the same location, the amount of new information decreases (Fig. 2 C, slope of the blue curve), while the cost to obtain each sample stays constant (Fig. 2 C, slope of the red curve). As a result, the net information reward (*i.e.,* new information minus sampling cost) is the highest for the first few samples, and decreases with continued sampling at the same location (Fig. 2 C, yellow curve). We posit that the observed magic number between 3-6 (Fig. 2 A) allows experts to efficiently increase information coverage at each sampling location.

(ii) Information inference hypothesis. We hypothesize that experts could obtain information about a location by either direct sampling at the location, or inferring information indirectly from a nearby sampled locations. As a result, the net information reward (Fig. 2 D, yellow) exhibits a non-monotonic trend with location interval: when sampling too densely (Fig. 2 D, small interval), the amount of information (Fig. 2 D, blue) is sufficiently large, but the total sampling cost (Fig. 2 D, red) is also high, resulting in a reduced net reward (Fig. 2 D, yellow); on the other hand, when sampling too sparsely (Fig. 2 D, large interval), the indirect information that can be inferred from each sampled location is too small, resulting in a much smaller total information (Fig. 2 D, blue), and consequently a smaller net reward (Fig. 2 D, yellow) as well. Therefore, the observed heuristics of sampling every 2-4 locations (Fig. 2 B) allows experts to efficiently collect both direct and indirect information while balancing sampling cost.

Based on the two hypotheses, the information reward, I, for sampling n measurements at location l, can be modelled as:

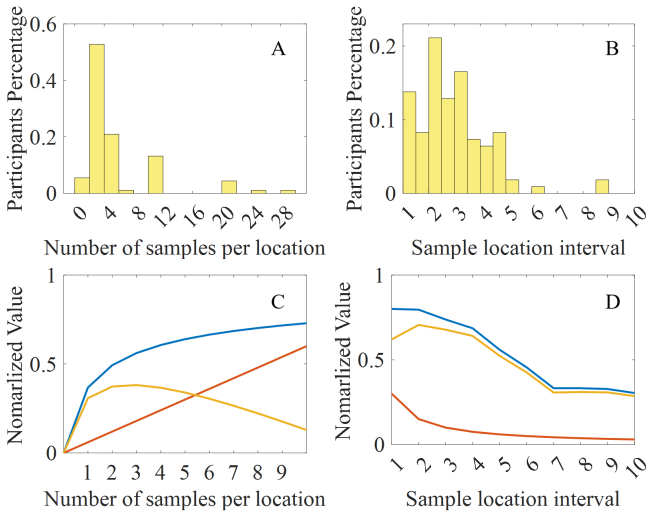

Fig. 2: Observed sampling strategy heuristics from human responses (A, B), and hypothesized information reward models to explain the observed heuristics (C, D). (A) Magic number (*i.e.,* the constant number of samples per location) distribution from participants who chose to take a constant number of measurements at each location in their initial strategy. (B) Average location interval distribution from participants who selected evenly-spaced sampling locations in their initial strategy. (C) Model-predicted information reward for choosing different number of samples per location. Blue represents total information. Red represents the total sampling cost. Yellow represents the net reward, computed as the information (blue) minus the sampling cost (red). (D) Model-predicted information reward for choosing different sampling location intervals. Color scheme is the same as (C).

$$I(l,n) = \sum_{l_s \in L_s} I_m \cdot e^{\frac{-\delta}{\sqrt{n}}} \cdot e^{\frac{-(l-l_s)^2}{2\beta^2}}, \tag{1}$$

where I_m represents the complete amount of information one could obtain from location l, l_s represents a sampled location, δ represents measurement noise, and β represents the information inference decay factor. Here $e^{\frac{-\delta}{\sqrt{n}}}$ represents the diminished information reward with increased number of samples at the same location, whereas $e^{\frac{-(l-l_s)^2}{2\beta^2}}$ represents the reduced inferred information reward with increased distance between location l with the sampled location, l_s.

Model predictive success. To test the information reward model, we compute the distribution of I among all possible sampling locations, and compare the locations with large information reward with human expert' actual sampling location choices. Fig. 4 shows that with the highly-simplified information reward function (Eqn. 1), the robot-predicted high-information-reward locations (Fig. 4, blue) matches closely with experts-selected locations at low information level (*i.e.,* total information $<$ 60%). However, at higher information level, the information based reward fails to capture experts' sampling strategies. This suggests experts dynamically update their sampling objective and priority in response to incoming data, with the objective to effectively increase the amount of information coverage only initially driving data collection behavior.

V. Experts' Sampling Strategy Adaptation is Hypothesis Verification Oriented and Discrepancy Driven

To investigate what other objectives drive experts' data collection behavior, particularly at high information coverage, we examined their reported hypothesis beliefs (measured by subjective confidence). Once experts collected a sufficient amount of initial information, they began to report specific beliefs towards the given hypothesis (Fig. 3 A), and their reported beliefs exhibited a strong relationship to the amount of discrepancy between their collected measurements and the given hypothesis: when discrepancy between the data and hypothesis was high, experts reported low confidence in hypothesis, and when discrepancy was low they reported higher confidence (Fig. 3 B). Based on this observation, we hypothesize that once experts have formed an initial belief towards the given hypothesis – either supporting or not supporting – their data collection decisions are driven by an objective to further verify this belief.

The potential discrepancy between measurement and hypothesis one may observe if they were to sample at location, l, can be modelled as:

$$D(l) = \sum_m \sum_y P(m|l)P(y|m)\,|H_y - y|\,, \tag{2}$$

where H_y represents the expected shear strength based on the given hypothesis. $P(y|m)$ represents the probability of getting a specific shear strength measurement value, y, given moisture m. $P(m|l)$ represents the probability of getting a moisture measurement within the range, m, at location l.

Model predictive success. To test our hypothesis that sampling strategy at high information level is discrepancy driven, we compute the distribution of D among all possible sampling locations. For experts holding the belief that their measurements do not support the given hypothesis, the robot predicts sampling locations with maximal potential discrepancy:

$$R_{invalidate} = \underset{l}{argmax}(D(l)), \tag{3}$$

For experts holding the belief that their measurements do support the given hypothesis, the robot predicts sampling locations with minimal potential discrepancy:

$$R_{validate} = \underset{l}{argmin}(D(l)) \tag{4}$$

Fig. 4 shows that predictions using the discrepancy reward functions (Eqn. 3, 4) successfully captured sampling location choices for more than 70% of human experts. In addition, the high prediction rate of information reward function (Fig.

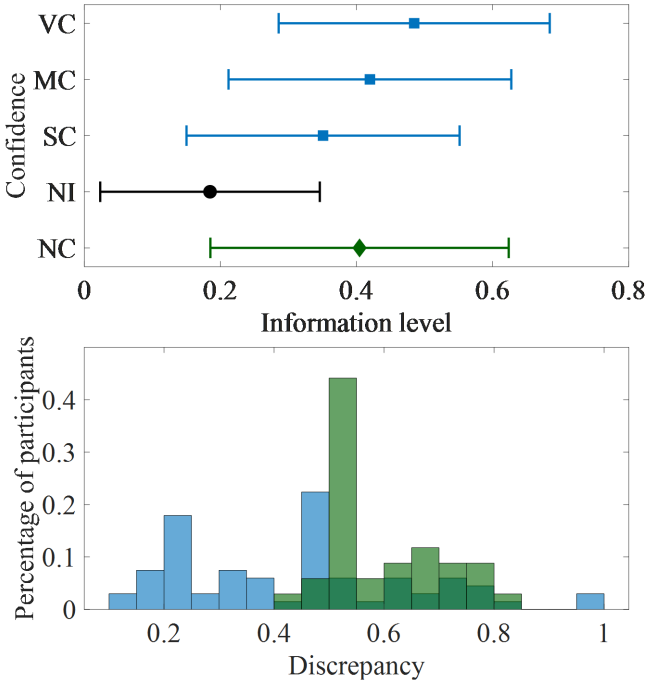

Fig. 3: (A) Experts' confidence in specific beliefs – either supporting (blue) or not supporting (green) – towards the given hypothesis increases with the increase of information level. VC, MC, SC, NI, NC stands for Very confident, Moderately confident, Slightly confident, No information, and Not confident. (B) Expert-reported beliefs towards the given hypothesis are related to the discrepancy between their measurements and the given hypothesis. Blue represents supporting and green represents not supporting.

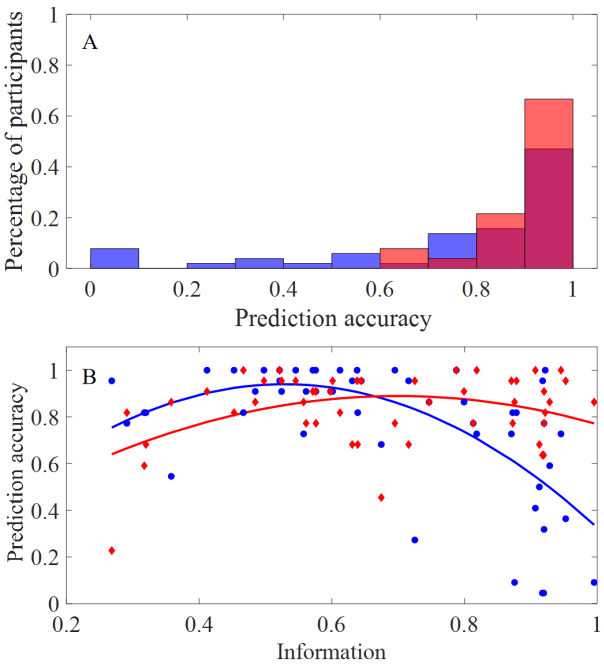

Fig. 4: Accuracy of robot-predicted sampling locations based on the information reward (blue) and discrepancy reward (red). (A) Histogram of the prediction accuracy. (B) Prediction accuracy at different information coverage levels. Information reward (blue) captures human sampling behaviors at lower information level, whereas discrepancy reward (red) captures sampling behavior at higher information level.

4B, blue) at lower information coverage (*i.e.,* total information $< 60\%$), and the high prediction rate of discrepancy reward function (Fig. 4, red) at high information coverage (*i.e.,* total information $> 60\%$), reveals a transition from an information coverage oriented objective to a hypothesis verification oriented objective as the level of information increased.

The strikingly high prediction rate of the extremely simple model reflects the success of our approach informed by the study of expert human behavior. Going forward, our findings will help enable the development of more cognitively-compatible robots teammates that can infer human experts' abstract objectives to better support exploration and understanding of complex earth environments vulnerable to desertification processes compounded by the effects of anthropogenic global climate change.

References

[1] H.-O. P. Mbow, A. Reisinger, J. Canadell, and P. O'Brien, "Special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems (sr2)," *Ginevra, IPCC*, vol. 650, 2017.

[2] S. Veron, J. Paruelo, and M. Oesterheld, "Assessing desertification," *Journal of Arid Environments*, vol. 66, no. 4, pp. 751–763, 2006.

[3] U. Saranli, M. Buehler, and D. E. Koditschek, "Rhex: A simple and highly mobile hexapod robot," *The International Journal of Robotics Research*, vol. 20, no. 7, pp. 616–631, 2001.

[4] F. Qian, D. Jerolmack, N. Lancaster, G. Nikolich, P. Reverdy, S. Roberts, T. Shipley, R. S. Van Pelt, T. M. Zobeck, and D. E. Koditschek, "Ground robotic measurement of aeolian processes," *Aeolian research*, vol. 27, pp. 1–11, 2017.

[5] F. Qian, D. Lee, G. Nikolich, D. Koditschek, and D. Jerolmack, "Rapid in situ characterization of soil erodibility with a field deployable robot," *Journal of Geophysical Research: Earth Surface*, vol. 124, no. 5, pp. 1261–1280, 2019.

[6] G. Picardi, M. Chellapurath, S. Iacoponi, S. Stefanni, C. Laschi, and M. Calisti, "Bioinspired underwater legged robot for seabed exploration with low environmental disturbance," *Science Robotics*, vol. 5, no. 42, p. eaaz1012, 2020. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.aaz1012

[7] C. G. Wilson, F. Qian, D. J. Jerolmack, S. Roberts, J. Ham, D. Koditschek, and T. F. Shipley, "Spatially and temporally distributed data foraging decisions in disciplinary field science," *Cognitive Research: Principles and Implications*, vol. 6, no. 1, pp. 1–16, 2021.

[8] W. D. Neys, "Dual processing in reasoning: Two systems but one reasoner," *Psychological science*, vol. 17, no. 5, pp. 428–433, 2006.

[9] J. S. B. Evans and K. E. Stanovich, "Dual-process theories of higher cognition: Advancing the debate," *Perspectives on psychological science*, vol. 8, no. 3, pp. 223–241, 2013.

[10] T. S. Braver, "The variable nature of cognitive control: a dual mechanisms framework," *Trends in cognitive sciences*, vol. 16, no. 2, pp. 106–113, 2012.

[11] C. G. Wilson, C. E. Bond, and T. F. Shipley, "How can geologic decision making under uncertainty be improved?" *Solid earth*, pp. 1–34, 2019.

[12] J. S. B. Evans, "In two minds: dual-process accounts of reasoning," *Trends in cognitive sciences*, vol. 7, no. 10, pp. 454–459, 2003.