Entropy Mastermind: Learning from Humans about Intelligent Systems

Lara Bertram[1,2], Eric Schulz[3], Matthias Hofer[4], and Jonathan D. Nelson[1,2]

[1] University of Surrey, Guildford, UK
[2] Max Planck Institute for Human Development, Berlin, Germany
[3] Harvard University, Cambridge, US
[4] Massachusetts Institute of Technology, Cambridge, US

Abstract. Despite the rapidly increasing computational power of today's computers, a key challenge in robotics and artificial intelligence is successful and smooth interaction with the environment – something that comes naturally for most humans. Which strategies enable humans to learn about and adapt to an environment with such high computational complexity? What can machines learn from human strategies in order to exploit their computational powers efficiently and be a natural, understandable and manageable part of our everyday lives? We review how a proposed research program, centered around a novel version of the code-breaking game Mastermind, sets out to address these questions.

Keywords: Entropy · Mastermind · Uncertainty · Information Search · Cognition and Emotion.

1 Intelligent Masterminds

1.1 The Fruit Salad Mastermind Game

Imagine a magic fruit bowl filled with six different kinds of fruit: apples, oranges, pears, grapes, blueberries and pineapples. Whenever you take a fruit out of the fruit bowl and eat it, it is refilled with the same kind of fruit. If I now randomly sampled three fruits from the fruit bowl and put them on a skewer one by one, the result would be a fruit skewer with one out of 216 possible fruit combinations. In order to find out the fruit combination on my fruit skewer, you can test out different combinations and I will give you feedback about your guess.

1.2 Entropy and Fruit Salad Mastermind

This is the kind of situation people encounter in our novel experimental paradigm, Fruit Salad (or Entropy) Mastermind (see Fig. 1). Entropy Mastermind is an app-based, customizable version of the classic Mastermind code-breaking game, in which a hidden fruit code is drawn from a fruit bowl. The player has to guess this code by repeatedly making queries and getting feedback.

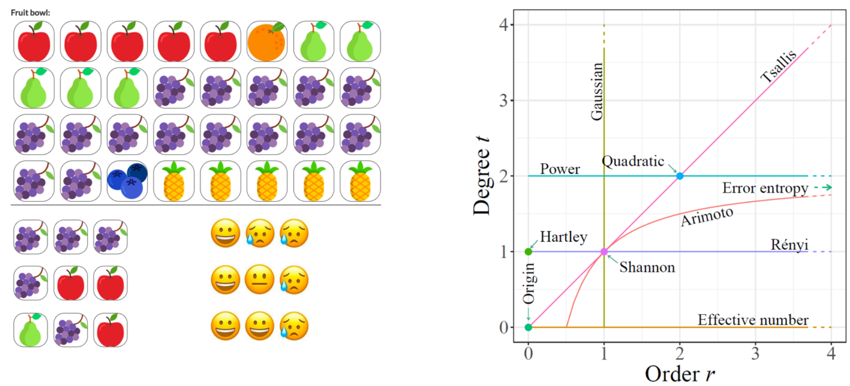

Fig. 1. Entropy Mastermind App and the Sharma-Mittal Space
Left: Example fruit bowl that generated the secret code (high entropy condition). Players receive feedback in the form of smileys. You can try the game at: http://jonathandnelson.com/curious/masterminding.html
Right: The Sharma-Mittal family of entropy measures is represented in a Cartesian quadrant with values of the order parameter r (how much minor hypotheses are disregarded) and the degree parameter t (how prominent the goal of getting as close as possible to the state of certainty is). Each point in the quadrant corresponds to a specific entropy measure, each line corresponds to a distinct one-parameter generalized entropy function. Several special cases are highlighted.

Importantly, the level of entropy in the fruit bowl differs between rounds of the game. For example, the fruit bowl could contain a lot of apples but only a few oranges, pears, grapes, blueberries and pineapples. Compared to a fruit bowl with equal distribution of fruits, this fruit bowl would be relatively low in entropy. The level of entropy in turn affects the probability of randomly sampling fruits: In the low entropy example, the probability of drawing an apple would be higher than the probability of drawing an orange. In the high entropy example, the probabilities would be the same. Mathematically, the uncertainty in a discrete random variable $K = k_1, k_2, ...k_n$, in our case the fruit skewer, can be measured by its entropy. The generalized Sharma-Mittal space of entropy measures [1] defines entropy as:

$$\text{entropy}(K) = \frac{1}{t-1}\left[1 - \left(\sum_{i=1}^{n} P(k_i)^r\right)^{\frac{t-1}{r-1}}\right], \tag{1}$$

where r is the *order* and t the *degree* of the entropy measure. The different entropy measures quantify the average surprise that would be experienced if the value of a random variable K, in our case the composition of the fruit skewer, was learned. We model people's information search behavior within this framework to find the information theoretical metrics which best describe how people search for information in differently entropic environments [2].

How do people perform on tasks like Fruit Salad Mastermind compared to computers? Human performance scales well from small (e.g., 216 possible codes) to large (e.g., 43 million possible codes) versions of the game. Our current information theoretic algorithms, however, do not scale well, and quickly become computationally intractable. What heuristic "tricks" does human cognition use that enable it to scale so well? And can those tricks be harnessed to improve machine learning systems?

2 The Mastermind Research Agenda

Our first goal in this research is to better understand human cognition and intelligence, by investigating and modeling cognitive variables such as numeracy, the psychological feature space, belief updating, and memory processes. People's emotional and motivational states affect cognition, learning and behavior [3, 4] and thus we are also researching how these variables shape information search in Mastermind. Early results suggest that self-concept and emotional dominance better predict how many queries people will need than numeracy alone.

Our second research goal is to learn from human intelligence to improve artificially intelligent systems. One example is to enhance adaptive game-based tutoring systems [5], by making them more sensitive to self-evaluative cognition and the emotional state of players.

In the field of human-like computing, insights from psychological studies on human cognition, motivation and emotion can help improve interactive and adaptive machines [6]. By using Fruit Salad Mastermind in combination with psychological experimental methods, we aim to contribute to this line of research.

References

1. Crupi, V., Nelson, J. D., Meder, B., Cevolani, G., Tentorie, K. (2018). Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cognitive Science, **42***, 1410–1456.
2. Schulz, E., Bertram, L., Hofer, M., Nelson, J. D. (2019). Exploring the space of human exploration using Entropy Mastermind. *BioRxiv, 540666*. https://doi.org/10.1101/540666
3. Volz, K. G., Hertwig, R. (2016). Emotions and decisions: Beyond conceptual vagueness and the rationality muddle. *Perspectives on Psychological Science **11***, 101–116.
4. Lerner, J. S., Keltner, D. (2000). Beyond valence: Toward a model of emotion-specific influences on judgement and choice. *Cognition and Emotion, **14***, 473–493.
5. DMello, S., Jackson, T., Craig, S., Morgan, B., Chipman, P., White, H., Person, N., Kort, B., El-Kaliouby, R., Picard, R. W., Graesser, A. (2008). AutoTutor detects and responds to learners affective and cognitive states. *Workshop on Emotional and Cognitive Issues at the International Conference on Intelligent Tutoring Systems*, 306–308. https://doi.org/10.1109/TE.2005.856149
6. Picard, R. W. (1997). *Affective computing.* Cambridge, Massachusetts: MIT Press. ISBN 9780262161701.