

Testing Computational Models of Goal Pursuit

Florian Mohnert, Mateo Tošić, Falk Lieder

Rationality Enhancement Group, Max Planck Institute for Intelligent Systems
Max-Planck-Ring 4, Tübingen, Baden-Württemberg, Germany

Abstract

Goals are essential to human cognition and behavior. But how do we pursue them? To address this question, we model how capacity limits on planning and attention shape the computational mechanisms of human goal pursuit. We test the predictions of a simple model based on previous theories in a behavioral experiment. The results show that to fully capture how people pursue their goals it is critical to account for people's limited attention in addition to their limited planning. Our findings elucidate the cognitive constraints that shape human goal pursuit and point to an improved model of human goal pursuit that can reliably predict which goals a person will achieve and which goals they will struggle to pursue effectively.

Keywords: goal-pursuit; decision-making; computational modelling; bounded rationality

Introduction

Human behaviour and cognition are fundamentally goal directed (Carver & Scheier, 2001). It has been proposed that goals serve to simplify complex decision problems so that people can solve them more effectively despite their limited computational resources (Lieder & Griffiths, 2019). Optimal goal pursuit entails many computationally intractable problems (Bourgin, Lieder, Reichman, Talmon, & Griffiths, 2017). The brain has to somehow approximate these solutions within the constraints of its bounded computational resources.

Previous work on human problem solving (Newell, Simon, et al., 1972) suggested that rather than making a complete plan for how to achieve their goal people often just look a single step and choose their action based on a heuristic estimate of resulting reduction in their distance to the goal. Furthermore, recent work on human decision making has highlighted that people's decisions are highly constrained by their limited attentional resources (Gabaix, 2014). Despite these and other insights (O'Doherty, Cockburn, & Pauli, 2017), there are still no definite computational models of how people pursue their goals in complex dynamic environments.

Here, we leverage the principle of resource-rationality (Lieder & Griffiths, 2019) to derive a model of how cognitive constraints shape the computational mechanisms of human goal pursuit. We test the resulting models against people's performance in a newly developed paradigm for studying how people pursue goals in complex dynamic environments. We find that our model correctly predicts which goals people will achieve easily and which one's they will fail to achieve. This is a significant step towards a grounding recommendations and tools for helping people to set better goals into computational models of human goal pursuit.

A simulated-microworld paradigm for studying goal setting and goal pursuit

Researchers interested in how people solve complex problems introduced Simulated Microworlds (SMWs) (Brehmer & Dörner, 1993) as a more realistic alternatives to the puzzles that were predominant in the problem-solving literature at the time. Simulated microworlds are complex dynamic systems that model situations from real-life. In a SMW human participants might, for example, manage a fictional airline company over the course of 5 years and at each timestep t (e.g. days, months etc.) humans intervene with the system by manipulating a set of exogenous variables such as the current salary of the service staff or buying a certain amount of fuel and then observe the corresponding changes in the environment state or endogenous variables such as revenue or customer satisfaction of the company. The SMW conceived for our experiments are based on simple systems of linear equations like in (Funke, 1993). Formally, a simulated microworld consists of a set of D_e exogenous variables e_1, \dots, e_{D_e} and a set of D_s endogenous variables s_1, \dots, s_{D_s} whose dynamics over time are determined by a system of linear equations, that is

$$\mathbf{s}_{t+1} = f(\mathbf{s}_t, \mathbf{e}_t) = \mathbf{A} \cdot \mathbf{s}_t + \mathbf{B} \cdot \mathbf{e}_t, \quad (1)$$

where the vectors \mathbf{s}_t and \mathbf{e}_t contain the current values of the endogenous and exogenous variables respectively. The matrix \mathbf{A} is in $\mathbb{R}^{D_s \times D_s}$ and matrix \mathbf{B} in $\mathbb{R}^{D_s \times D_e}$. Matrix \mathbf{A} determines both the *eigendynamics* of the system (i.e. how variables affects themselves from t to $t + 1$) and also *side effect* (i.e. how variables influence other variables from t to $t+1$). Matrix \mathbf{B} on the other hand determines the effect of the human intervention of the exogenous variables on the endogenous variables. The task posed to participants is to reach a certain goal state of the endogenous variables. We define a goal g by two vectors namely a location \mathbf{g}_l and a scale \mathbf{g}_s . The location components are the desired states of the system and the scale components define a diagonal covariance matrix S a standard deviation around the desired states, defining how specific the goal is (i.e. how close one needs to get in order to reach it). An appropriate distance measure to determine how close a given endogenous state \mathbf{s}_t is to a goal g is the standardized Euclidean distance, $d(\mathbf{s}_t, g) = \sqrt{(\mathbf{s}_t - \mathbf{g}_l)^T S^{-1} (\mathbf{s}_t - \mathbf{g}_l)}$. This distance measure can be used to define a threshold ω for goal achievement, that is a goal g is reached at timestep t if $d(\mathbf{s}_t, g) < \omega$.

While simulated microworlds have previously been used to study problem solving (Brehmer & Dörner, 1993), managerial or other trainings, we present a way of using them to study goal setting and goal pursuit.



Models of goal-pursuit

Previous work problem solving emphasised that people's ability to achieve their goals is bounded by people's limited planning horizon (Simon & Newell, 1971). Recent work on economic decision-making emphasised the role of limited attention (Gabaix, 2014). Here we consider how both of these constraints might jointly shape how people pursue their goals.

Limited planning leads to hill-climbing

Unlike strategies that plan multiple steps ahead, hill-climbing myopically chooses one step at a time based on the immediate reduction in the (perceived) distance to the goal. The strategy of hill-climbing has been previously proposed as a process model of how humans solve problems (Simon & Newell, 1971). People's reliance on this strategy might reflect that planning is costly and can only be performed for a limited number of steps.

In a SMW hill-climbing amounts to setting the exogenous variables such that the distance between the endogenous variables \mathbf{s}_t and the goal \mathbf{g} is minimized. More formally, a hill-climbing agent arrives at it's new exogenous inputs \mathbf{e}_{t+1} by taking a step in the direction of the negative gradient of the distance w.r.t. the exogenous inputs, that is

$$\mathbf{e}_t = -\lambda \cdot \nabla_{\mathbf{e}_t} d(\mathbf{s}_{t+1}(\mathbf{e}_t), \mathbf{g})|_{\mathbf{e}_t=\mathbf{0}}, \quad (2)$$

where λ denotes the learning rate. Note that the gradient is computed at the point $\mathbf{e}_t = \mathbf{0}$, which means that the agent always computes it's next exogenous input to the system assuming that the current exogenous input is at 0 for all variables. We additionally constrain the agent to a *budget* β . If the sum of the absolute value of the exogenous input is greater than β the exogenous inputs are clipped to a maximum of β .

Adding limited attention: the sparse-max hill-climbing model

Certainly, hill-climbing captures some but not all aspects of human goal-pursuit. In particular, the hill-climbing strategy ignores the increasing cognitive costs posed by limited attention and memory when an agent focuses on larger sets of endogenous variables at once. The *sparse-max* operator (Gabaix, 2014) is a psychologically plausible version of the max operator that takes the trade-off between maximized utility and cost of attention into account. A sparse solution to the minimization problem posed to the hill-climbing agent is to only minimize the distance of the goal and a subset of the endogenous variables. The sparse-max is a two step process, 1) the agent chooses an attention vector $\mathbf{m} \in \{0, 1\}^{D_s}$ that indicates which subset of the endogenous variables the agent should attend to, 2) the agent performs hill-climbing by minimizing the distance between the goal and the chosen subset of endogenous variables. The choice of \mathbf{m} is performed as follows,

$$\operatorname{argmin}_{\mathbf{m}} d(\mathbf{s}_{t+1}^{\mathbf{m}}, \mathbf{g}) + c \cdot \sum_{i=1}^{D_s} \mathbf{m}_i \quad (3)$$

where $\mathbf{s}_{t+1}^{\mathbf{m}}$ denote the endogenous state at $t + 1$ only taking the active dimensions in \mathbf{m} into account and c denotes the cognitive costs of attending to a variable. Once \mathbf{m} is chosen the hill climbing step can be computed as before but only reducing the distance for the variables active in \mathbf{m} .

Null-model

As a third model of goal pursuit we propose a model that chooses its exogenous input at each timestep as follows:

1. n endogenous variables are selected at random.
2. For each of the chosen endogenous variables s the agent randomly selects one of the exogenous input variable e that has an effect on s .
3. Each selected exogenous variable e is given a random input (within the allowed budget) that adjusts the corresponding endogenous variable s in the direction of its target value.

This model has one free parameter, namely n .

Performance metric To measure an agent's performance in a SMW, we compute it's *average closeness* as follows,

$$c_{avg} = \frac{\sum_{t=1}^T (d(\mathbf{s}_1, \mathbf{g}) - d(\mathbf{s}_t, \mathbf{g}))}{T} \quad (4)$$

Additionally, we record whether an agent model or participant has reached it's goal during the episode.

Model selection and parameter estimation

Assuming that the error ϵ around the system state that the participant intended to reach in each step is normally distributed ($\Delta \sim \mathcal{N}(0, \sigma_{err})$), the likelihood of a participants data \mathcal{D} under a model m with parameters θ is

$$p(\mathcal{D}|m, \theta) = \prod_{t=1}^T \frac{1}{\sigma_{err} \sqrt{2\pi}} \exp\left(\frac{-\Delta_t^2}{2\sigma_{err}^2}\right), \quad (5)$$

where $\Delta_t = \|\mathbf{s}_{t+1} - f(\mathbf{s}_t, \text{model}(\mathbf{s}_t; \theta))\|_2$ is the Euclidean distance between the next state the participant reached and the state that the model's input $\text{model}(\mathbf{s}_t; \theta)$ would have moved the system to. To estimate the parameters θ and σ_{err} we maximize the log likelihood of the data using Bayesian optimization.

How do people pursue goals in SMWs?

We developed an interactive SMW game to collect data on how people pursue goals and to evaluate our models of human goal pursuit. Using this paradigm we designed an experiment to test whether the standard hill-climbing model can predict which goals people will achieve and which they will struggle to pursue effectively.

Stimuli and Procedure

For this experiment we designed a SMW with 4 exogenous and 5 endogenous variables and embedded it in an interactive

experiment ¹. Participants are told that they are situated on a far-away planet and that their job is that of an Alien-Farmer. The task consists of controlling five “farming measures” (the endogenous variables) using a set of four “resources” (the exogenous variables). As we aim to study goal pursuit all the dynamics of the system are displayed from the beginning, such that participants do not have to explore first in order to be able to plan effectively. Participants are told that across a *season* consisting of 20 rounds they have to work towards a goal consisting of *target values* and *ranges* for each farming measure respectively. In the rightmost column the current goal is shown to participants and in the leftmost column the resources can be adjusted via slide-bars. The budget β is set to 25 and the current budget that is still available is displayed on top of the leftmost column. After a participant decided on an exogenous input, the next round can be reached via a click on the *next round* button.

We constructed 30 situations consisting of initial values for the farming measures and a goal. These 30 situations were designed such that the hill-climbing model predicts that participants should always achieve the goal in the 10 “Easy” situations, fail to achieve the goal in the “Medium” situations but come closer to it, and end up farther away from the goal than they started out in the 10 “Difficult” situations. The situations within each level of difficulty were generated by randomly generating situations and recording the performance of the hill-climbing agent. The first 10 situations were chosen such that the hill-climbing agent reaches the goal at least once (by goal achievement threshold ω) and achieves an average closeness of between 0 and 50. In situations with medium difficulty the goal was not reached by the agent but the agent achieved an average closeness of between -20 and 50. Finally, hard goals were chosen such that the agent did not reach the goal and its average closeness was in between -100 and -200. The experiment was conducted as follows, first participants were given instructions on how to perform the task, followed by a training situation with a single easy situation. If the participants managed to reach the goal at least once in this training situation they were randomly allocated to one out of the 30 test situations. Participants could earn a performance dependent bonus in the test situation.

Participants

We recruited 215 participants via the online platform *Positly* ². In total 43 of participants were allocated to the easy situations, 84 to the medium difficulty situations and 88 to the hard situations.

Results

Consistent with the predictions of the hill-climbing model, most people were able to achieve goals in the easy category but only very few goals are reached in the medium and hard categories (see Figure 2). Furthermore, as shown in Figure 1, both




Figure 1: Average closeness of humans and hill-climbing agent with a learning rate of 1.




Figure 2: Goal achievement humans and hill-climbing agent with a learning rate of 1.

hill-climbing and human average closeness are significantly lower in the medium category compared to the easy category (human: $t = 7.59, p < .001$; model: $t = 4.65437, p < .001$). However, while the hill climbing model performs significantly worse ($t = 37.76, p < .001$) in the hard category compared to the medium category, humans do not show such a large drop in performance ($t = 0.23, p = .815$). This highlights that while the hill-climbing model can capture some important aspects of human goal-pursuit, it remains incomplete.

As shown in Table 1, the choices of the sparse-max model and people showed robust performance in the difficult situations in which the hill-climbing agent moved away from the goal. This might be because – unlike the hill-climbing model – both people and the sparse-max model have limited attentional resources that they preferentially allocate to the most helpful information. This might have allowed both people and the sparse-max model to ignore the misleading lures that led the hill-climbing model astray in the difficult situations. Our model selection results confirmed that the sparse-max model captured people’s goal pursuit in difficult situations more accurately than the basic hill-climbing model (see Table 1). These findings suggest that taking into account the rational use of limited attention is critical for understanding human goal pursuit.

Which model explains people’s goal pursuit strategies best? After fitting each model to each participant, we per-

¹We host a trial version of our goal pursuit task at <http://goal-setting-experiments.herokuapp.com/demo>

²<https://www.positly.com/>

	Easy	Medium	Hard
hill-climbing	-8.49 vs. 24.09 (N = 39)	-47.33 vs. 22.15 (N=83)	-285.09 vs. -119.87 (N =19)
sparse-max	24.9 vs. -9.96 (N=4)	-39.7 vs. 36.51 (N=1)	16.20 vs. 15.34 (N=69)

Table 1: Average closeness of humans vs. models in the 3 difficulty categories. The results are grouped depending on the model that best explains the respective participant's data.

formed model selection according to Akaike's information criterion (AIC). We found that 141 participant's data was best explained by the hill-climbing model, 74 participant's data was best explained by sparse-max hill-climbing and 0 participants data was best explained by the Null-model.

How well do the selected models fit the participants' data? To evaluate how well the fitted models explain participant's trajectories in the SMW we devised a metric that is similar to the proportion of explained variance. At a given timestep t we measured how much of the change in position is explained by the model as follows,

$$F(t) = \frac{\|s_{t+1} - s_t\|_2 - \|f(s_t, \text{model}(s_t; \theta)) - s_{t+1}\|_2}{\|s_{t+1} - s_t\|_2} \quad (6)$$

For each participant we computed \bar{F} which is the average over all $F(t)$. We found that for 125 of the 215 participants, we find that \bar{F} was larger than 0. For these participants, the best fitting model explained about $18\% \pm 1\%$ of the variance in their state sequences on average. This suggests that our model can successfully explain a notable portion of participant's trajectories. However, it also shows that for a large subgroup of participants (90/215) our models were unable to explain their goal pursuit strategies. Table 1 compares the performance of participants to the performance of the model that best explained their behavior.

Qualitative predictions Previously, when comparing participants' performance to the hill-climbing model that was used to generate the situation, we found that participants in the hard category did not perform as poorly as the model did 1. In the third column of Table 1 this gap is explained by the fact that for 69 out of 88 participants in the hard category the sparse-max model was the best fitting model whereas for another 19 the hill-climbing model was the best model. The sparse-max model performs as well as people with an average closeness of 15.34 compared to the 16.2 of the people.

One specific prediction the sparse-max model makes is that given a nonzero cost parameter the best input might not take all endogenous variables into account. For the participants best explained by the sparse-max model we compared the average number of exogenous inputs that was set to a non-zero number to the quantity that the sparse-max model predicts. In Figure 3 we show that the sparse-max model indeed captures the fact that most participant's used only a small subset of the exogenous inputs at each timestep. The hill-climbing model, by contrast, would have manipulated an average of 3.92 of the 4 inputs on each step of those participants' problems.




Figure 3: The average number of modified exogenous inputs of the participants for which the sparse-max model had the best fit.

Discussion

Based on Simulated Microworlds we designed a new interactive experimental model how humans pursue goals in complex control tasks. Additionally, we introduced two models of human goal pursuit: the first model performs hill-climbing, and the second model additionally accounts for people's limited attention. We found that the hill-climbing model correctly predicted which goals people achieve and which they fail to achieve. However, people's performance was robust to goals that would have led the standard hill-climbing model astray. This discrepancy could be reconciled by taking into account that peoples' attentional resources are limited. The limited number of inputs variables that participants' modified on each trial provided additional support for people's limited attentional resources. Future work will leverage these insights to investigate which sub-goals are most effective at helping people achieve challenging goals in complex environments.

References

- Bourgin, D., Lieder, F., Reichman, D., Talmon, N., & Griffiths, T. (2017). The structure of goal systems predicts human performance. In *Cogsci 2017*.
- Brehmer, B., & Dörner, D. (1993). Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior*, 9(2-3), 171–184.
- Carver, C. S., & Scheier, M. F. (2001). *On the self-regulation of behavior*. Cambridge University Press.
- Funke, J. (1993). Microworlds based on linear equation systems: A new approach to complex problem solving and experimental results. In *Advances in psychology* (Vol. 101, pp. 313–330). Elsevier.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4), 1661–1710.
- Lieder, F., & Griffiths, T. L. (2019). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 1–85.
- Newell, A., Simon, H. A., et al. (1972). *Human problem solving* (Vol. 104) (No. 9). Prentice-Hall Englewood Cliffs, NJ.
- O'Doherty, J. P., Cockburn, J., & Pauli, W. M. (2017). Learning, reward, and decision making. *Annual review of psychology*, 68, 73–100.
- Simon, H. A., & Newell, A. (1971). Human problem solving: The state of the theory in 1970. *American Psychologist*, 26(2), 145.