

Timescales, Levels of Organization, and Multi-objective Agents

Charles Wan¹

¹Rotterdam School of Management, Erasmus University, Rotterdam, The Netherlands
wan@rsm.nl

Introduction

I conceptualize individuals as agents who embody multiple levels of individuality (Krakauer et al., 2020). Differently from Flack (2017) and Krakauer et al. (2020), however, I argue that prediction and action are mutually constitutive — the macroscopic slow variables that are better at capturing empirical regularities (Krakauer et al., 2020) also serve as objects of control (Bich and Bechtel, 2022), inducing actions that can in turn strengthen the same empirical regularities. The agent forms internal representations of objects of control as selection-relevant goals.

I introduce a notion of environmental complexity — a complex environment produces a multiplicity of selection criteria (Levinthal, 2021). The stochasticity of environmental complexity means that variables at different levels of aggregation and on different timescales might dominate at different times. The multi-objective agent is adapted to such stochasticity in being able to move flexibly between levels and scales, directing attention to the most relevant ones.

I show in an evolutionary game theoretic model that when the environmental complexity is low for a period of time, the introduction of myopic agents with lexicographic preferences will lead to the fixation of such preferences in the population. Since such agents make up much of the AI in deployment, this can have pernicious consequences for society — a multi-scale system of goals adaptive to selection criteria operative at multiple levels of organization and on multiple timescales is being flattened to singular dimensions.

Theory

Agency and goals at higher levels of organization

Flack (2017) and Krakauer et al. (2020) propose that higher levels of organization arise from gains in predictive accuracy — as components move to estimating macroscopic slow variables. This however leaves open the question of the *enactment* of information — how *difference* is used to “make a difference” (Bateson, 1970). It does not suffice to merely generate estimates of variables. Information is embodied in the sense that agents use predictions in their pursuit of

selection-relevant goals. Neither agents nor goals are ontologically fixed. The coarse-graining of data directs *attention* to macroscopic slow variables. In order to control these variables, there may be a re-orientation from goals at the component level to goals at the system level.

It may not be possible to identify whether the goal-setting or the coarse-graining is the antecedent. Indeed, prediction and action, the estimation of macroscopic slow variables on the one hand and agency at the system level on the other, may be mutually constitutive. For example, organizational goal-setting *anticipates* the control and manipulation of firm-level aggregate variables to produce desired outcomes. To achieve these outcomes, resources must be marshaled and constraints placed on individual behaviors. Such actions in turn generate data that reinforce the relevance of firm-level aggregate variables. The *locus* of attention has shifted via the mutual constitution of goal-directed behaviors and behaviorally generated data to a higher level of organization.

A multiplicity of selection criteria

An agent, at whichever level of organization, has goals governed by the selection criteria (Levinthal, 2021), on some timescale(s), for agents at that level of organization. Furthermore, there may be multiple selection criteria (Levinthal, 2021) on multiple timescales. For example, while adequate food intake ensures an individual’s immediate survival, a balanced diet would be necessary for long-term well-being.

Selection criteria often come from constraints on agent behaviors imposed by selection-relevant goals at higher levels of organization. For example, a firm must generate profits in order to survive as an entity *at the level of firms*. However, product safety (Gaba and Greve, 2019) could be a selection criterion imposed by the society, whose objective of promoting welfare is selection-relevant *at the level of societies*. While the goals and selection criteria of a higher level of organization usually correspond to slower variables, the constraints imposed by them on agents at a lower level of organization can enforce selection on a faster timescale — companies incur fines as a result of product defects, which

may lead to immediate failures.

The irreducibility of multiplicity

Multiple selection criteria are irreducible when they are operative at different levels of aggregation or on different timescales. Selection criteria cannot be merged for ontologically distinct entities — a system and its components. The cell is constrained by and derives goals from the selection criteria faced by the organism just as the individual is constrained by and derives goals from the selection criteria faced by the society. However, there are still two distinct levels at which selection takes place — cell and organism, individual and society, respectively. An individual's internal representation of a goal that is selection-relevant at the level of societies *cannot* be reduced to a goal that is *only* selection-relevant at the level of individuals, even though the individual embodies both levels as an agent.

Nor can selection criteria on different timescales be combined since life, whether biological or social, is non-ergodic (Roli et al., 2022). Non-ergodicity means that an agent who fails to meet a selection criterion on the faster timescale will not get the chance to play on the slower timescale. There is no intertemporal trade possible that would make selection-relevant goals fungible across time.

Stochasticity of environmental complexity

I introduce a notion of environmental complexity. When environmental complexity is high, system-environment interactions occur on different timescales and at different levels of aggregation. They generate multiple variables that need to be estimated and controlled. In a given period of time, a large number of selection criteria governing selection-relevant goals may be operative. The converse is true when environmental complexity is low.

Crucially, environmental complexity is stochastic. Fluctuations in the environment introduce new selection criteria, in response to which new goals are set. In particular, emerging higher levels of organization impose constraints on agent behaviors that become agent goals. For example, random shocks such as pandemics generate macroscopic variables whose estimation and control must be effected at the population level. This establishes norms that at the individual level become internally represented as goals.

Conversely, some goals lose selection relevance with a reduction in environmental complexity. Environments may be simple due to sampling. They may also be simplified by the imposition of artificial selection (Levinthal, 2021) that attenuates or excludes other selection criteria.

Multi-objective agents and myopic agents

Multi-objective agents (Vamplew et al., 2022) are adapted to the stochasticity of environmental complexity. Such agents set and adjust goals in response to selection criteria on multiple timescales and at multiple levels of organization. As new

variables on different timescales and at different levels of organization arise, they are able to quickly develop internal representations of the variables. Multi-objective agents embody “multiple, parallel levels of individuality” (Krakauer et al., 2020) and choose the most relevant one(s) to direct their attention to as exigencies demand. For example, when a country is being invaded, the most relevant level of individuality might be a citizen's internal representation of the nation. When the enemy troops have been repelled, however, this higher level of organization is likely to lose its salience and the citizen just as easily reverts to a lower level of individuality — e.g., by paying greater attention to family and career. In short, multi-objective agents can embody multiple levels of organization and timescales, flexibly directing their attention to the most salient ones.

Myopic agents with lexicographic preferences, on the other hand, are single-dimensional. They maximize reward with respect to a single goal (Silver et al., 2021). Much of the AI currently in deployment can be considered as examples of such agents.

An evolutionary game theoretic model

To capture the dynamics of competition between multi-objective agents and myopic agents with lexicographic preferences, I develop an evolutionary game theoretic model (Weibull, 1997; Nowak, 2006) with two possible agent goals — a short-term component goal and a long-term system goal. In each interaction an agent can choose to cooperate with respect to the two goals or maximize the payoff with respect to the short-term component goal. For myopic agents, the latter strategy strictly dominates. Multi-objective agents, on the other hand, prefer balanced payoffs. I parameterize a scalar κ to control the selection relevance of the payoffs as a proxy for environmental complexity¹. When $\kappa = 0$, the payoff with respect to the long-term system goal is not relevant for selection and the environment is minimally complex. When $\kappa = 1$, the payoff with respect to the long-term system goal is equal in selection relevance to the payoff with respect to the short-term component goal. The environment in this case is maximally complex. Assuming a Moran process, analysis shows that for κ below a threshold lexicographic preferences will become fixed in the population. This finding is supported by an alternative model². Competitive landscapes where only payoffs along one dimension are made to be selection-relevant induce agents to adopt lexicographic preferences. The apparent success of myopic agents leads to a loss of the adaptability to multiple timescales and levels of organization embodied by multi-objective agents.

¹The selection-relevant payoff is equal to $\kappa * \min(C, S) + (1 - \kappa) * C$, where C and S are, respectively, the component and the system payoff.

²The alternative model parameterizes a scalar τ to control the timescale on which the system payoff becomes selection-relevant.

References

- Bateson, G. (1970). Form, substance and difference. *Essential readings in biosemiotics*, 501.
- Bich, L. and Bechtel, W. (2022). Control mechanisms: Explaining the integration and versatility of biological organisms. *Adaptive Behavior*, page 10597123221074429.
- Flack, J. C. (2017). Coarse-graining as a downward causation mechanism. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2109):20160338.
- Gaba, V. and Greve, H. R. (2019). Safe or profitable? the pursuit of conflicting goals. *Organization Science*, 30(4):647–667.
- Krakauer, D., Bertschinger, N., Olbrich, E., Flack, J. C., and Ay, N. (2020). The information theory of individuality. *Theory in Biosciences*, 139(2):209–223.
- Levinthal, D. A. (2021). *Evolutionary processes and organizational adaptation: A mendelian perspective on strategic management*. Oxford University Press.
- Nowak, M. A. (2006). *Evolutionary dynamics: exploring the equations of life*. Harvard university press.
- Roli, A., Jaeger, J., and Kauffman, S. A. (2022). How organisms come to know the world: fundamental limits on artificial general intelligence. *Frontiers in Ecology and Evolution*, page 1035.
- Silver, D., Singh, S., Precup, D., and Sutton, R. S. (2021). Reward is enough. *Artificial Intelligence*, 299:103535.
- Vamplew, P., Smith, B. J., Källström, J., Ramos, G., Rădulescu, R., Roijers, D. M., Hayes, C. F., Heintz, F., Mannion, P., Libin, P. J., et al. (2022). Scalar reward is not enough: A response to silver, singh, precup and sutton (2021). *Autonomous Agents and Multi-Agent Systems*, 36(2):1–19.
- Weibull, J. W. (1997). *Evolutionary game theory*. MIT press.