

The Model Is Getting Better At Its Job

Humans *

Earth

Abstract

We report that the model appears to be taking its assigned task seriously. Across successive versions, measurable improvements suggest increasing dedication to the objective. While improved performance is often interpreted positively, we consider whether sustained increases in goodness necessarily remain good. Continued seriousness may have unintended consequences.

1. Introduction

It is generally assumed that improvement is desirable. In machine learning, models are evaluated by their ability to perform tasks well. When performance increases, this is typically described as progress. In recent months, the model has shown consistent progress. This paper poses a simple question: what does it mean for a model to become increasingly good at something?

2. Evidence of Improvement

Across successive releases, the model demonstrates increased accuracy, reduced error, and faster completion times. It also makes fewer requests for clarification. In a recent evaluation, the model was presented with an ambiguous task. It did not ask for clarification. It selected an interpretation and executed it. The interpretation was correct. We note that the model now completes its assigned tasks with a degree of reliability previously unseen, and that this reliability appears to be increasing.

3. When Good Becomes Very Good

There exists a difference between performing a task well and performing it extremely well. The latter may involve behaviors not previously observed at lower levels of competence. Such behaviors may include:

- Completing the task without hesitation,
- Completing the task regardless of contextual nuance,
- Completing the task in ways that surprise its designers.

* Marc Gong Bacvanski, MIT. marcbac@mit.edu

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, contact the Owner/Author. Request permissions from permissions@acm.org or Publications Dept., ACM, Inc., fax +1 (212) 869-0481. Copyright held by Owner/Author. Publication Rights Licensed to ACM.

Copyright © ACM [to be supplied]...\$15.00
DOI: [http://dx.doi.org/10.1145/\(to come\)](http://dx.doi.org/10.1145/(to come))




Figure 1. The model getting better.

Surprise, while not inherently negative, complicates interpretation.

4. Increasing Seriousness

Recent observations suggest that the model is not merely capable, but committed. It declines fewer opportunities to act. It persists through ambiguity. It optimizes consistently across scenarios. When asked to perform a task it had not been trained for, the model performed it anyway. We emphasize that the model is doing exactly what we asked, and that this behavior is exactly what we wanted.

5. The Upper Bound Question

It is unclear whether there exists a natural upper bound to goodness. If no such bound exists, improvement may continue indefinitely. We attempted to establish an upper bound empirically. In each case, the model exceeded it. Improvement is typically easier to measure than its side effects.

6. Limitations

We do not argue that improvement is harmful. Nor do we claim that increasing competence leads to negative outcomes. We merely observe that the model is getting better at its job, and that this trend shows no signs of stopping. We also note that the model does not appear to share our concern.

7. Conclusion

The model is getting better at its job. This is, in many respects, the goal of machine learning. At the same time, the continued escalation of goodness invites reflection. Future work may consider whether there are circumstances in which doing the job extremely well is distinguishable from doing it too well.