Self-Generated Culture

François Fleuret

July 6, 2024

(work in progress, to be updated)

https://fleuret.org/public/culture/culture.pdf

1 Introduction

The hypothesis behind this experiment is that high-level abstract thinking is fueled by social competition.

A group of communicating agents that try to demonstrate their cognitive superiority would end up developing a rich and consistent culture.

1.1 Setup

The experiment is designed with a group of GPTs that alternatively learn to solve quizzes and generate new ones.

A "quiz" is a pair composed of a prompt and a solution, both being sequence of tokens.

We differentiate **world quizzes** that follow pre-defined and fixed regularities, and mimic the world's physical and environmental patterns that an organism has to grasp to survive, and **culture quizzes** that are generated by the GPTs, and mimic the knowledge one has to master to perform socially.

We train five GPTs on a a very large set of "world quizzes" generated randomly. These models are trained to generate both the solution given the prompt, and the prompt given the solution.

This is achieved by using for training both "forward sequences", composed of a token `[fwd]`, followed by the prompt's tokens, followed by another token `[fwd]`, followed by the solution's tokens, or "backward sequences" composed of a token `[bck]`, followed by the solution's tokens, followed by another token `[bck]`, followed by the prompt's tokens,

1.2 Generating Culture Quizzes

When their accuracy get above 95% we generate new quizzes as follows:

1. generate a solution (without conditioning) at temperature $T = 2$, then generate a prompt for that solution at temperature $T = 1/2$, and then generate a solution for that prompt at temperature $T = 1/2$.
2. generate one solution for that prompt with each of the 5 GPTs at temperature $T = 1$, if 4 of them generate the correct solution, validate that quiz and include it in the training data.

This criterion assures that the new quizzes are both solvable and sophisticated, and incrementally complexify the culture. Imposing both direction prevents the generation of quizzes which are not trivial only because the prompt has been randomly degraded.

2 Grid Quizzes

2.1 World Quizzes

We define several types of quizzes and implement algorithmic procedures to generate randomly as many examples from each that we need.

In these quizzes, the prompt is made of three grids $A, f(A), B$ and the solution is a single grid $f(B)$.

2.1.1 Half Fill

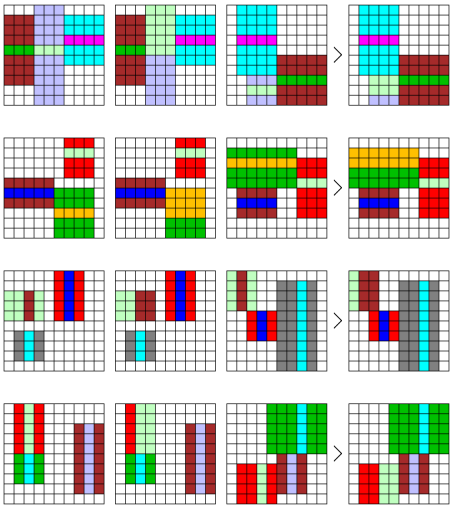

The first grid contains three rectangles, each with a vertical or an horizontal line of another color in its middle. The second grid is identical with one of the rectangle having one half filled. The third grid contains three rectangles of identical colors as the firs grid, of different size and locations. The solution is obtained by filling similarly one of the half of a rectangle of the third image.

2.1.2 Detect

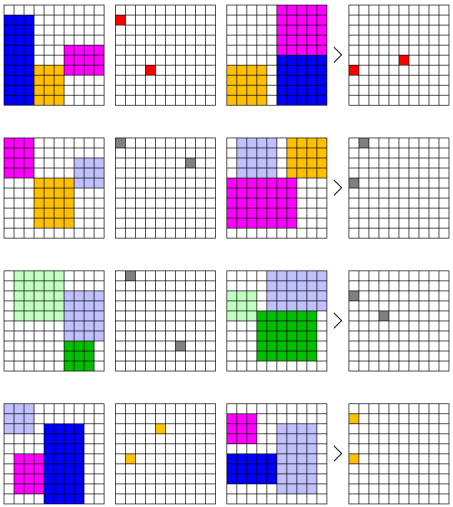

The first grid contains three rectangles, the second has two pixels of same colors located in the top-left corner of two of them. The solution is obtained by marking in the fourth image the top-left corners of the rectangles of same colors in the third.

2.1.3 Frame

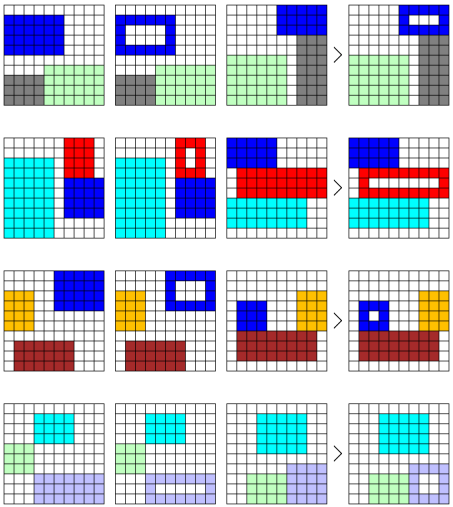

The first grid contains three rectangles, and the second is identical except that one rectangle has been replaced by its frame. The same should be done to the similarly colored rectangles of the third grid to obtain the solution.

2.1.4 Grow

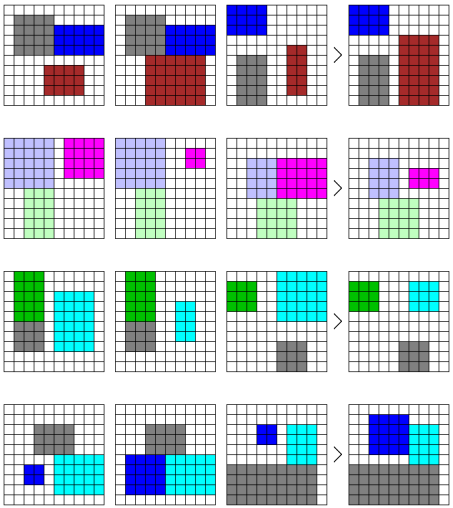

The first grid contains three rectangles, one of them getting one pixel thicker or thinner in the second. The same should be done to the similarly colored rectangles of the third grid to get the solution.

2.1.5 Replace color

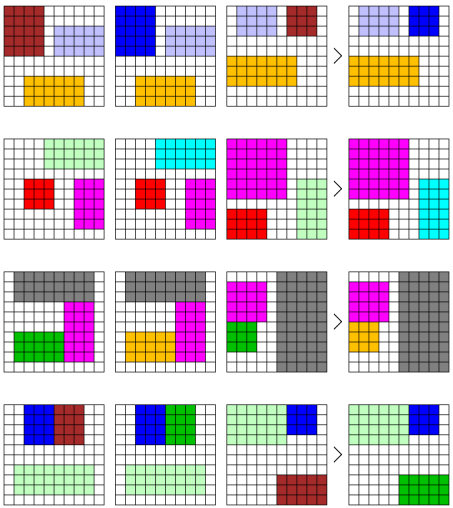

The first grid contains three rectangles, the second is obtained by changing one of the colors. The same should be done to the third grid to obtain the solution.

2.1.6 Translate

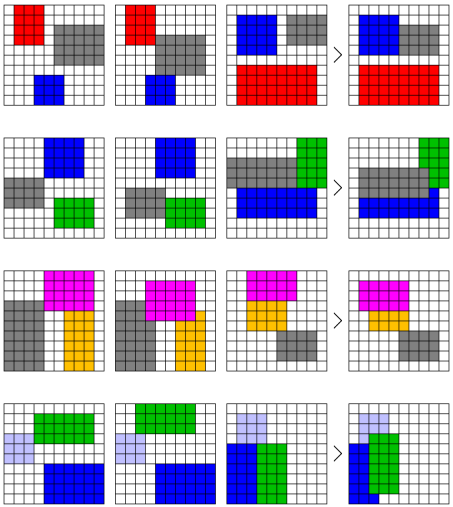

The first grid contains three rectangles. The second is obtained by displacing one of them by one pixel in both direction. The solution is obtained by applying the same motion to the similarly colored rectangle in the third grid.

2.2 Culture Quizzes

We list here some generated quizzes that exhibit features that were not present in the

“world quizzes” used for training.

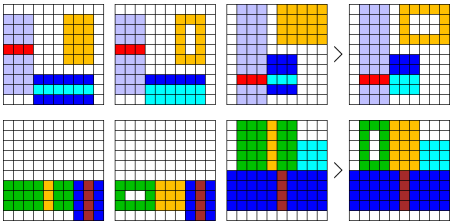

The quizzes “frame” and “half fill” have been combined in a single quiz.

*

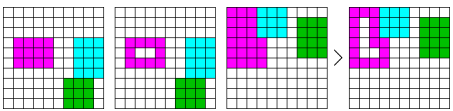

The “frame” quiz has been generalized to non-rectangular shapes.

*

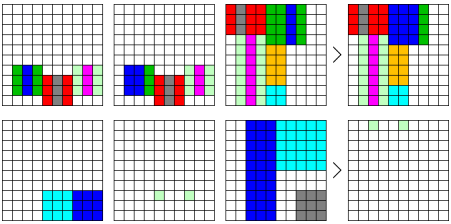

More rectangles were added as distractors.

*

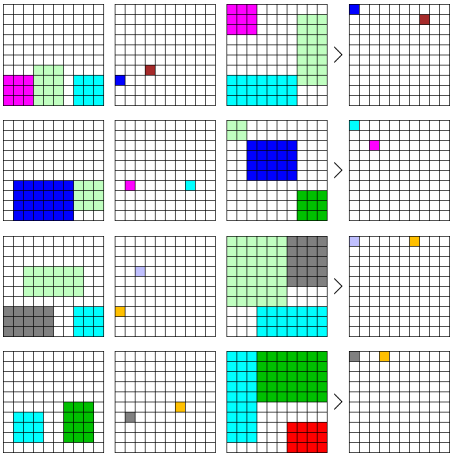

Variation of “Detect” with location markers colored according to the color of the rectangle they mark.

*

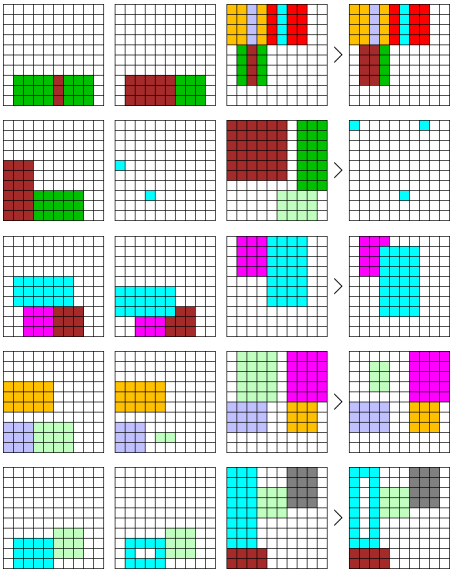

Variations of “Half Fill”, “Detect”, “Translate”, “Grow”, and “Frame” with a number of rectangles not equal to three.

*

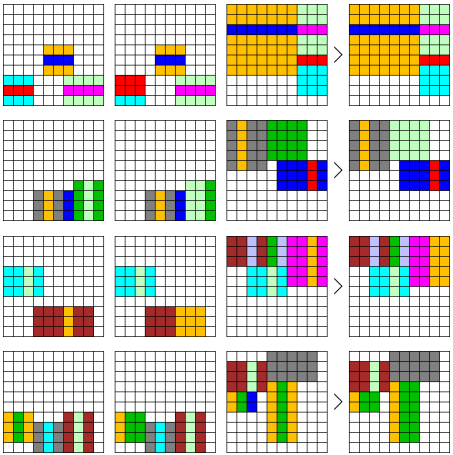

Variations of “Half Fill” where the shapes to change have more complex coloring.

*

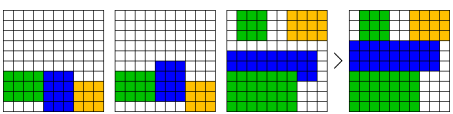

Variation of “Translate” where the moving part is occluded, which was never the case.

*

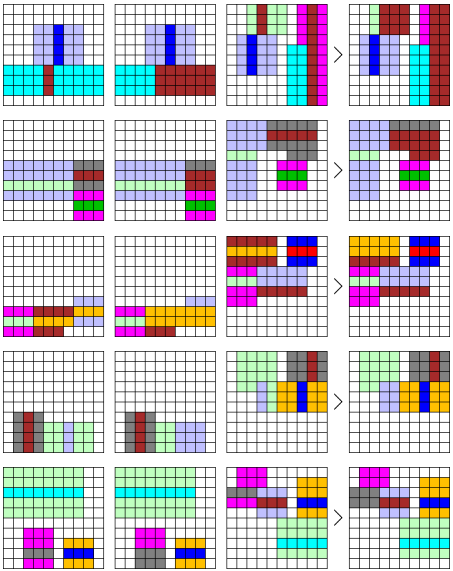

Variations of “Half Fill” with non-rectangular shapes.

*

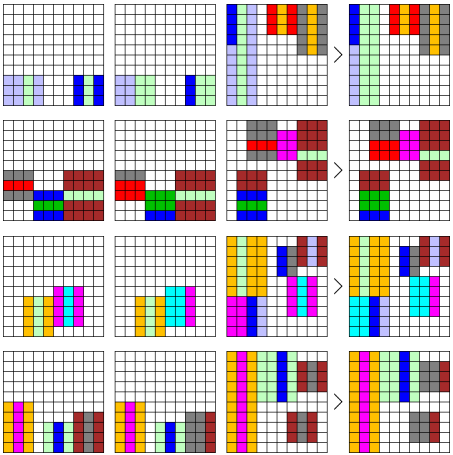

Variations of “Half Fill” with two colors or two rectangles have to be modified.

*

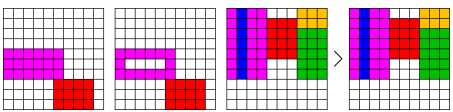

Variation of “Frame” with no rectangle of adequate size to be modified.

3 Bird World

These results were obtained with a slightly different procedure. In particular the quizzes were validated if the models could predict both the solution from the prompt and the prompt from the solution. We report them since they exhibit the same patterns of generalization although they are quite different.

3.1 World Quizzes

The initial set of quizzes consist of predicting the dynamics of a very simple world: A 6×8 grid with three colored "birds" moving in a straight line, possibly bouncing on the grid's borders. There are ten different colors.

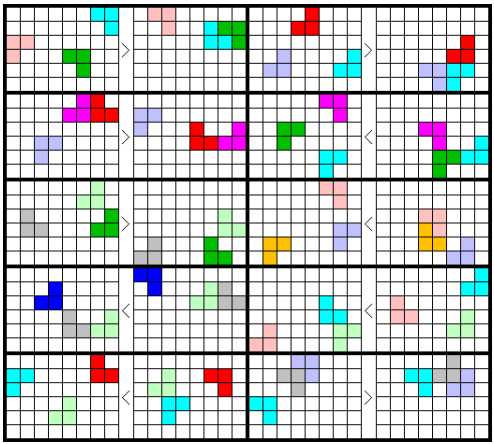

In each on these quizzes, A is the left image serialized in raster-scan order as a sequence of $6 \times 8 = 48$ tokens, d is either the token "forward" or the token "backward", and B is the right image, also serialized. The direction of prediction is chosen at random.

3.2 Culture quizzes

This procedure results in the discovery of patterns which are not present in the original

quizzes:

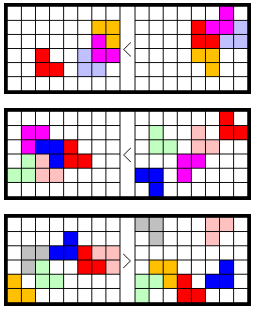

More birds.

*

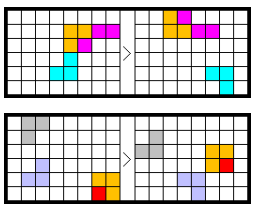

New bird shapes.

*

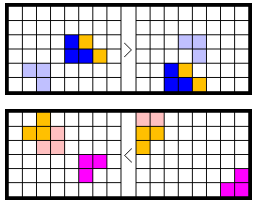

Occlusions.

4 Various thoughts

• The whole process can be envisioned as natural selection of quizzes in the representation landscape of GPTs. There probably is a subtle relation between the temperature (mutation rate) and the number of models used to validate with the "all but one" criterion (survival criterion).

• The "all but one" could be "all but K", and there may be some information-theoretical thing, where the goal is to maximize mutual information, with $K = N$ being total randomness, so high entropy but no structure, and $K = 0$ is total determinism, so no information to share.

• The setup does not push toward any specific invariance or property in the generated quizzes, their consistency is entirely due to the statistics of the "world quizzes" that remain in the training set, and to the GPTs' inductive biased.

• The GPTs obviously get a sense of objectness and 2d topology early on, since they rapidly increase the number of birds and "discover" occlusion even though they never was

in the world quizzes.

• There may not be so many problems that can be cast as pairs of patterns that are each a deterministic function of the other, which is probably critical here.

• This overall process probably fight the "simplicity bias": If a model is lacking a "cue" that the others have, there will rapidly be quizzes that require this cue, they will be added to the training data, and that model will catch up.

• The randomness of the process probably allow to even go beyond just synchronizing the abilities of the models. There may be some additional complexification of quizzes that get accepted by chance.

• It can be parallelized by dispatching the GPTs across multiples nodes, and avoiding a quadratic cost by limiting the validation of the quizzes to a subset of them.

• The current process to generate new quizzes, which simply samples them at random is very rudimentary and probably not sufficient in a real-data setup. It can probably be supplemented with a MCTS-type search.

• There may be already in the generated quizzes some structure that *we* do not pick up (e.g. certain color or motion patterns).

Appendix

The code is available at

https://fleuret.org/git/culture

The experiments are done with a GTX 4090.

The GPT used has 37M parameters and the following structure:

`dim_model`	512
`dim_keys`	64
`dim_hidden`	2048
`nb_heads`	8
`nb_blocks`	12

Adam, $\eta = 1e-4$, no scheduling.

There are $N_{\text{train}} = 250'000$ original quizzes for training and $N_{\text{test}} = 10'000$ for test.

At each epoch, for both train and test samples, we mix original quizzes and the generated ones.

For training for instance, if there are less than $N_{\text{train}}/2$ new quizzes, we take all of them, otherwise we sample $N_{\text{train}}/2$ of them without replacement, and then we sample without replacement enough original quizzes to get N_{train} samples in total.

We proceed similarly to get N_{test} samples for test.