SIMON KIRBY
UNIVERSITY OF EDINBURGH

ITERATED LEARNING:

INDUCTION, CULTURAL EVOLUTION, AND THE ORIGINS OF LINGUISTIC STRUCTURE

WHAT DOES EVOLUTIONARY LINGUISTICS ATTEMPT TO EXPLAIN?

- Language as a *major transition in evolution*
- Language enables unlimited transfer of information - hugely adaptive trait
- How?
- It is unique in exhibiting shared *systematic structure*

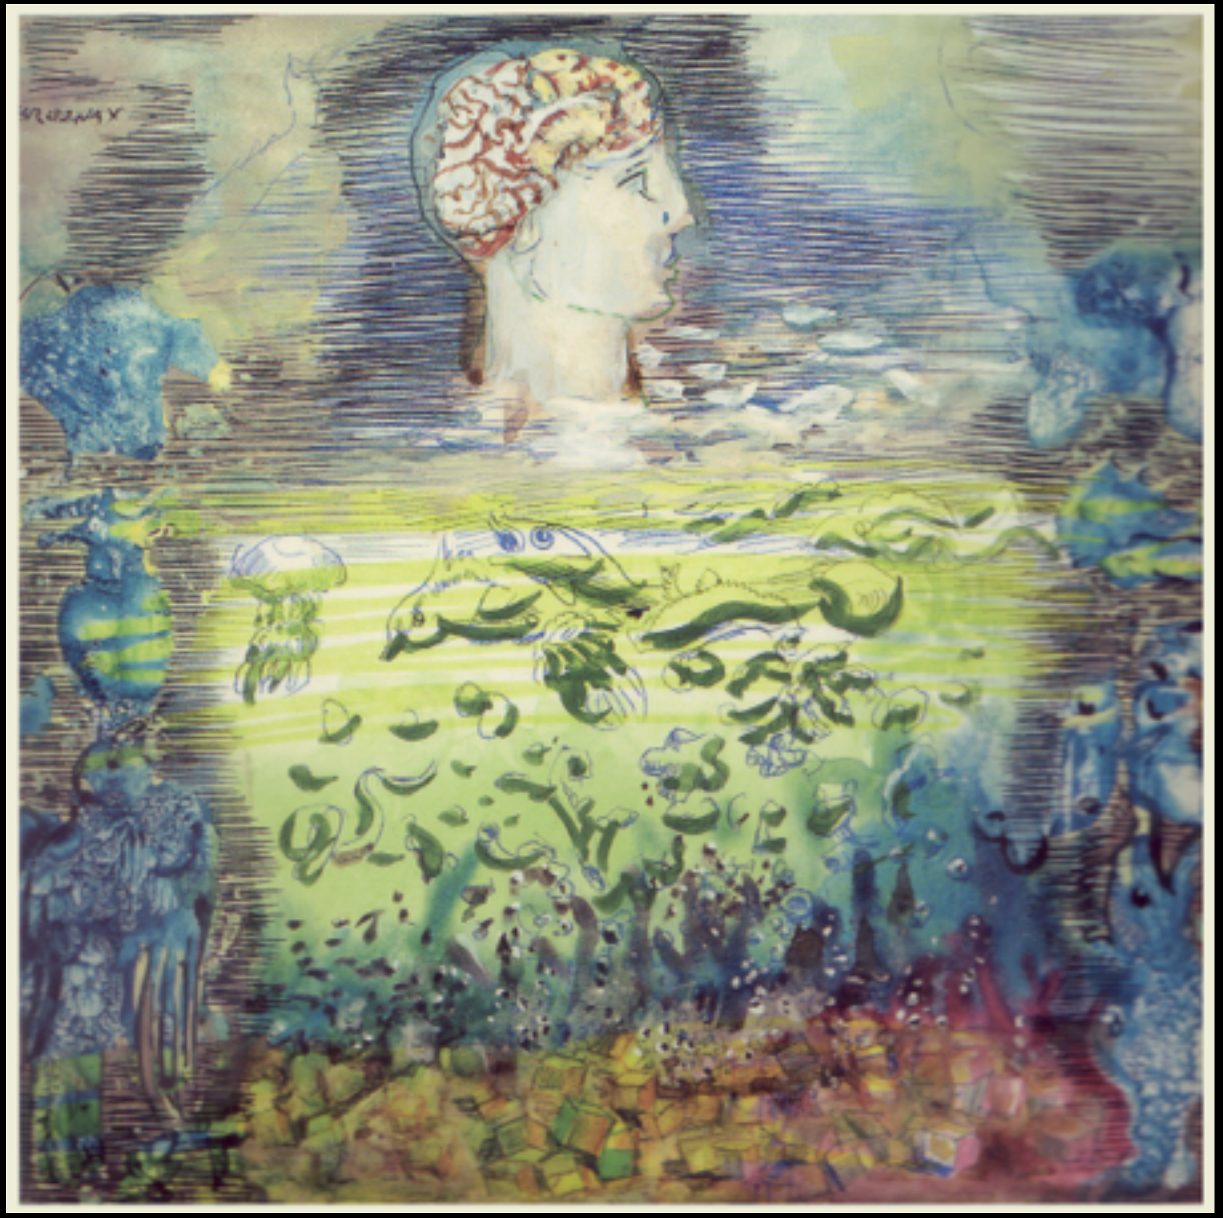

Maynard Smith & Szathmáry (1997)

WHAT DO WE MEAN BY SYSTEMATIC STRUCTURE?

INDEPENDENT BEHAVIOURS

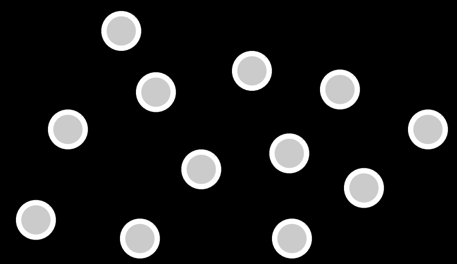

Typical of most communication in nature

SYSTEM OF BEHAVIOUR

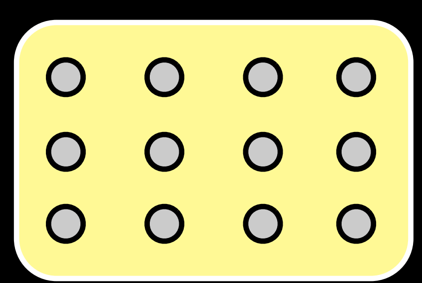

Fundamental to language at all levels

GRAMMAR

combinatoriality

compositionality

morphosyntax

phonology

Why does language exhibit shared system-wide structuring in the way it constructs messages?

OUR PROPOSED ANSWER

Language structure is the result of a trade-off between two partially competing pressures

Representational compressibility ⟷ Expressivity

LINGUISTIC STRUCTURE

Language structure is the solution to the problem of building the simplest system that is nevertheless useful

WHAT'S THE MECHANISM?

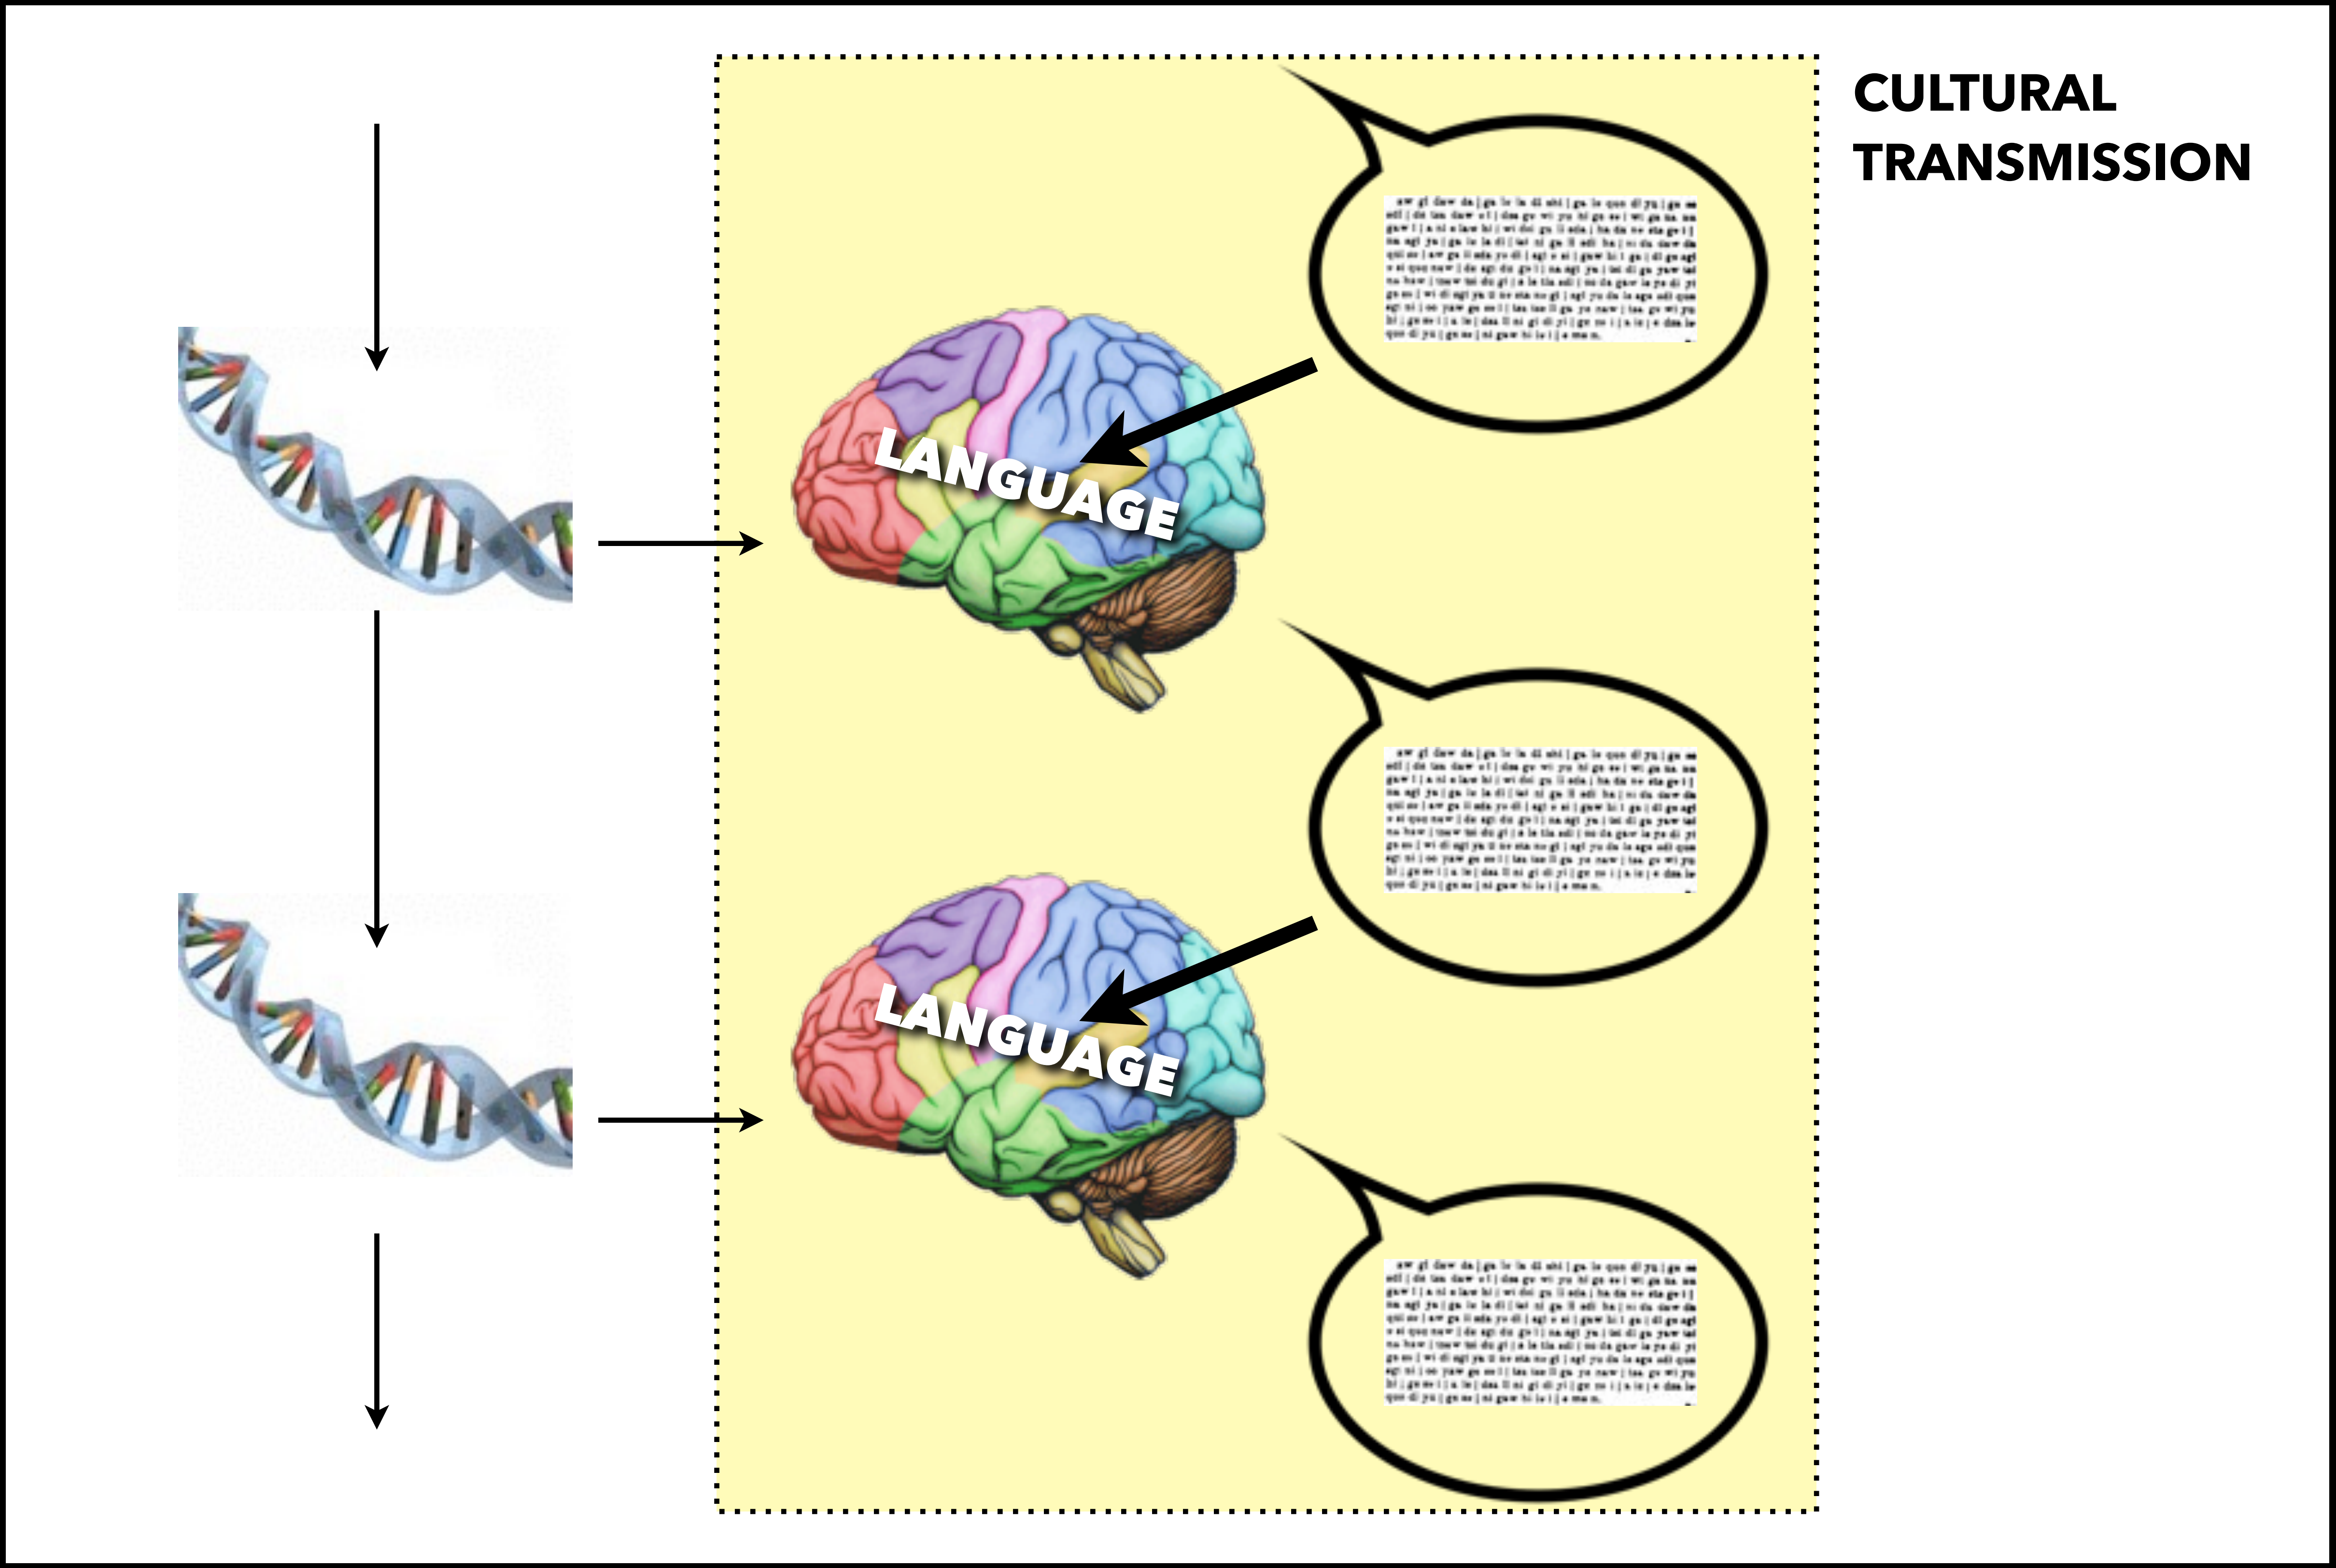

ITERATED LEARNING

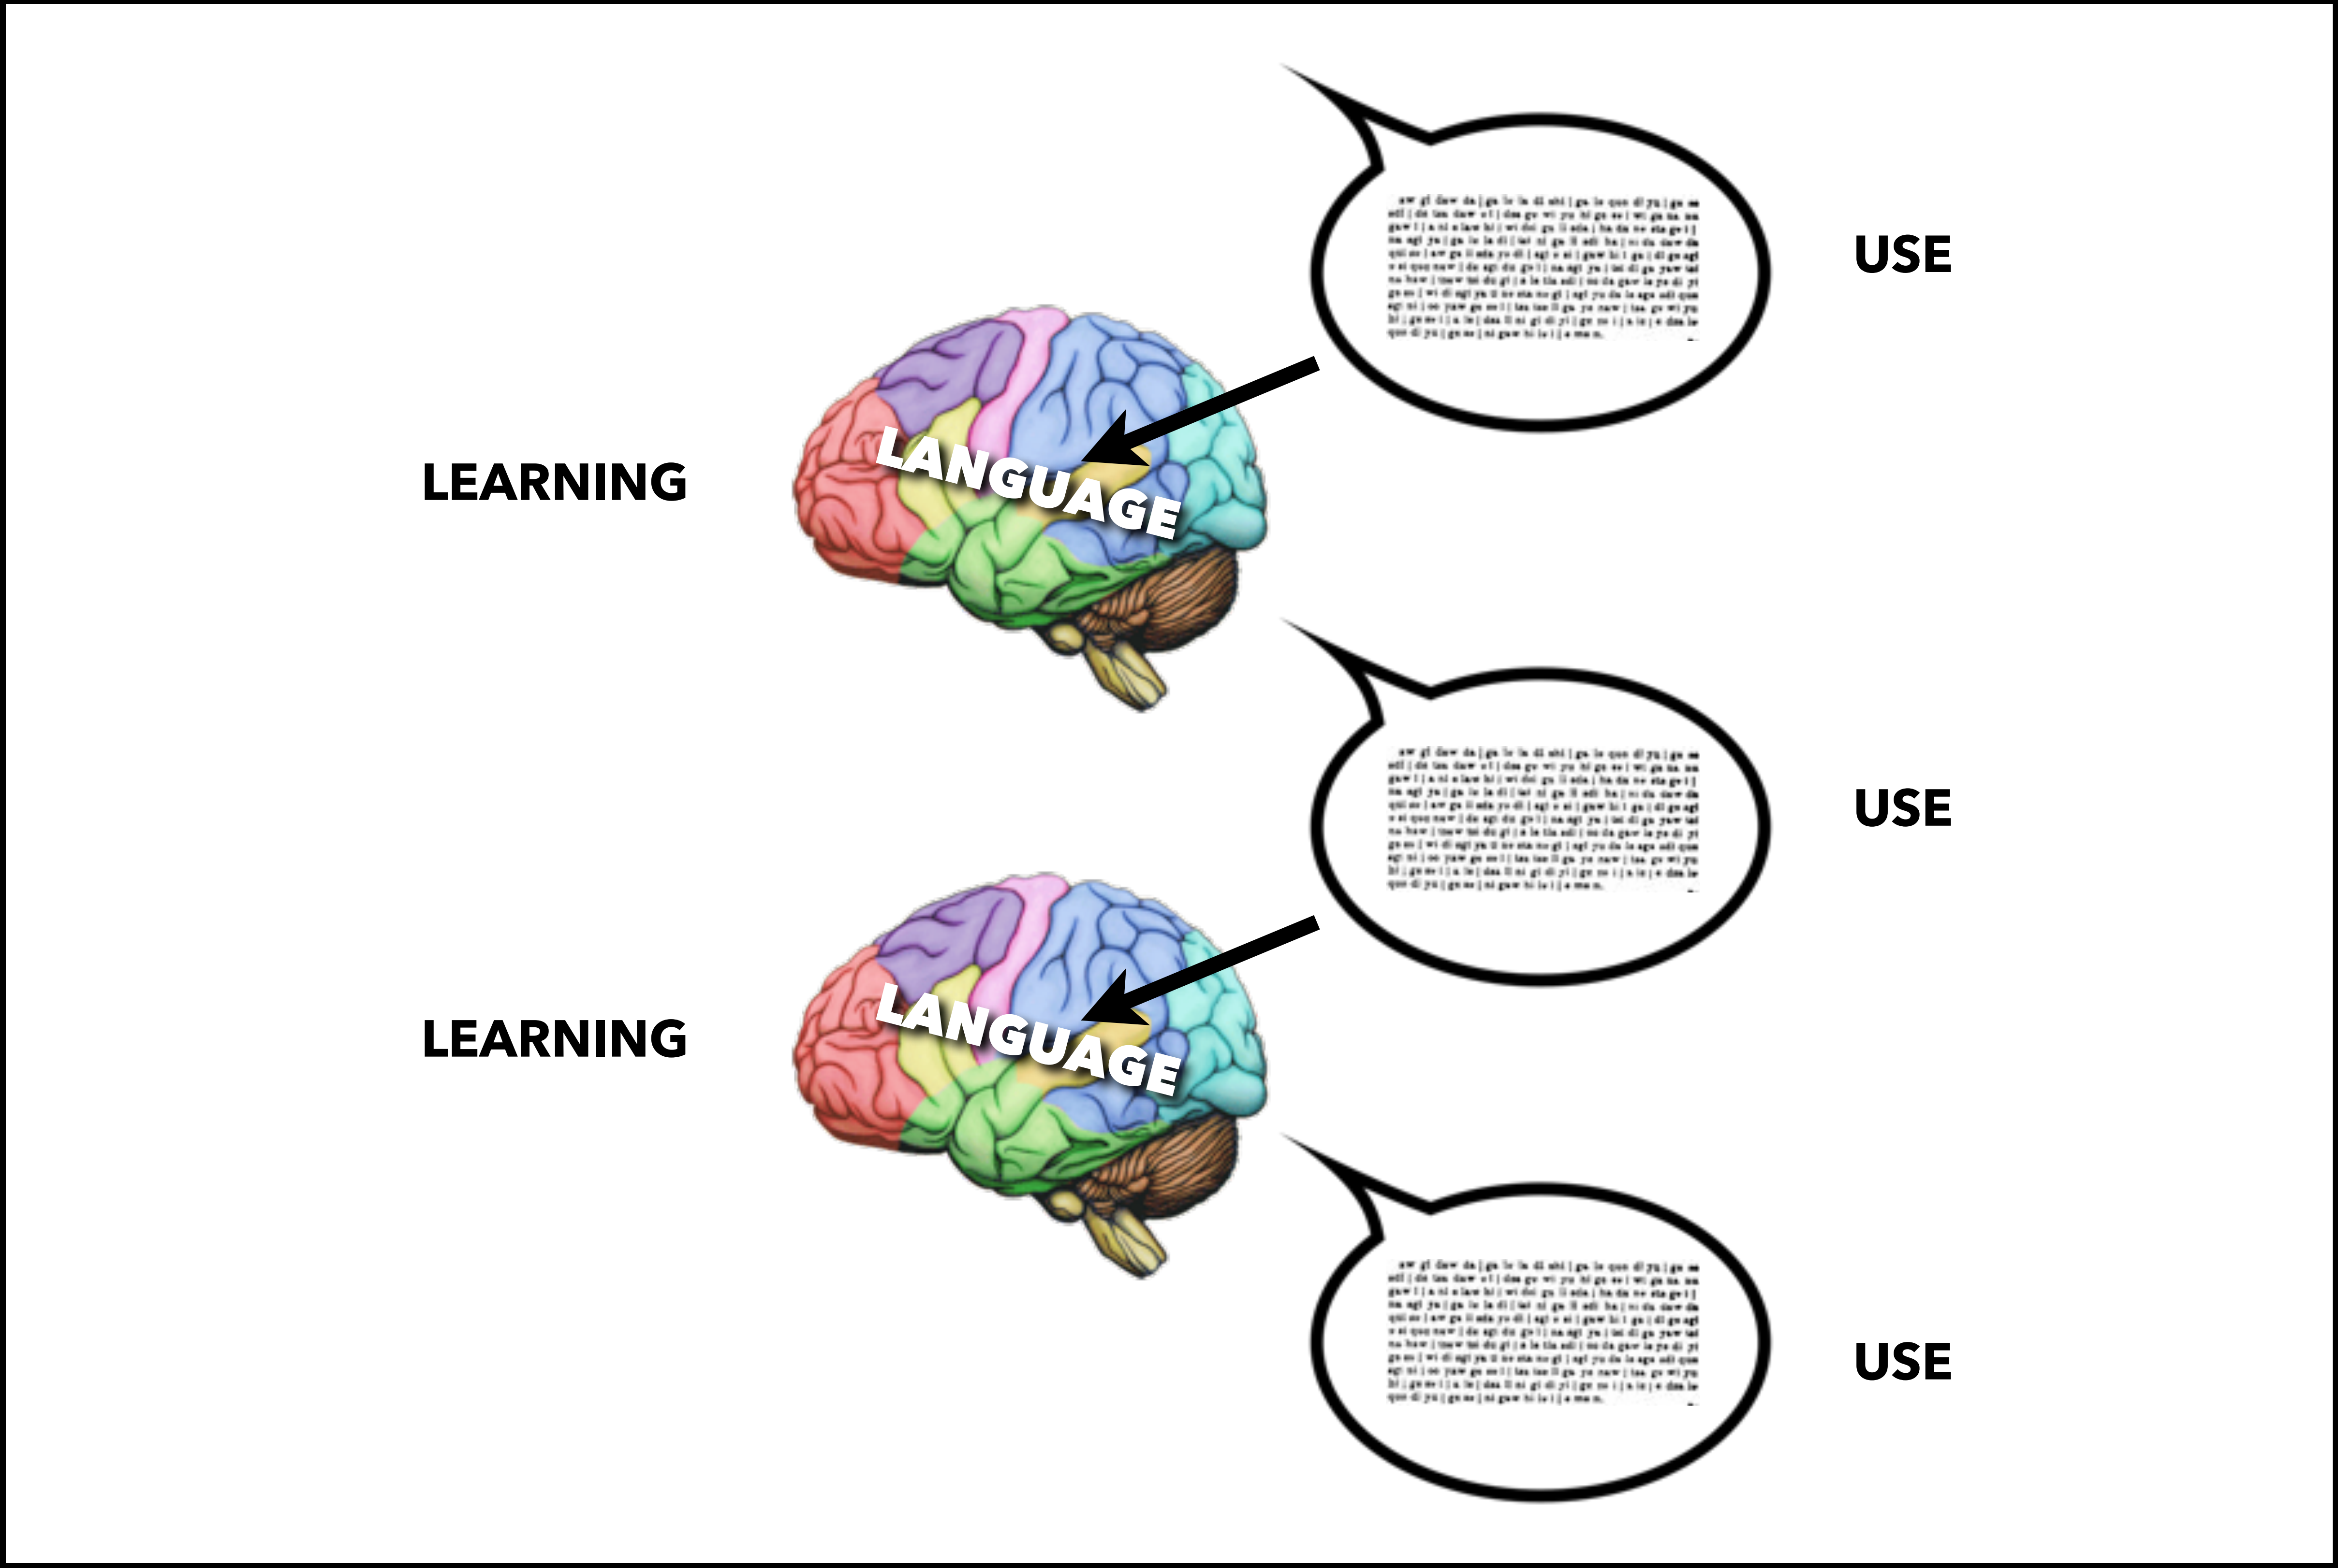

OVERVIEW OF SESSION

An investigation of how iterated learning leads to structure in language

1. Start looking in detail at one study attempting to show emerging structure in psychology lab

2. Then talk through a simple agent-based simulation replicating the experiment

3. Another lab experiment showing emergence of symbols

4. Very rapid overview of current work in progress in a variety of domains (and even species!)

Question: should understanding the *origins* of structure in behaviour change the way we build intelligent systems?

PART 1:
EVOLVING A LANGUAGE IN THE LAB

EARLY WORK: THE ALIEN LANGUAGE EXPERIMENTS

- Participants learn strings of syllables paired with structured meanings
- Start with **holistic** (unstructured) language
- Output of participant at test becomes language for next "generation"
- Each participant only sees half of the language, randomly chosen, but must generalise to all meanings.

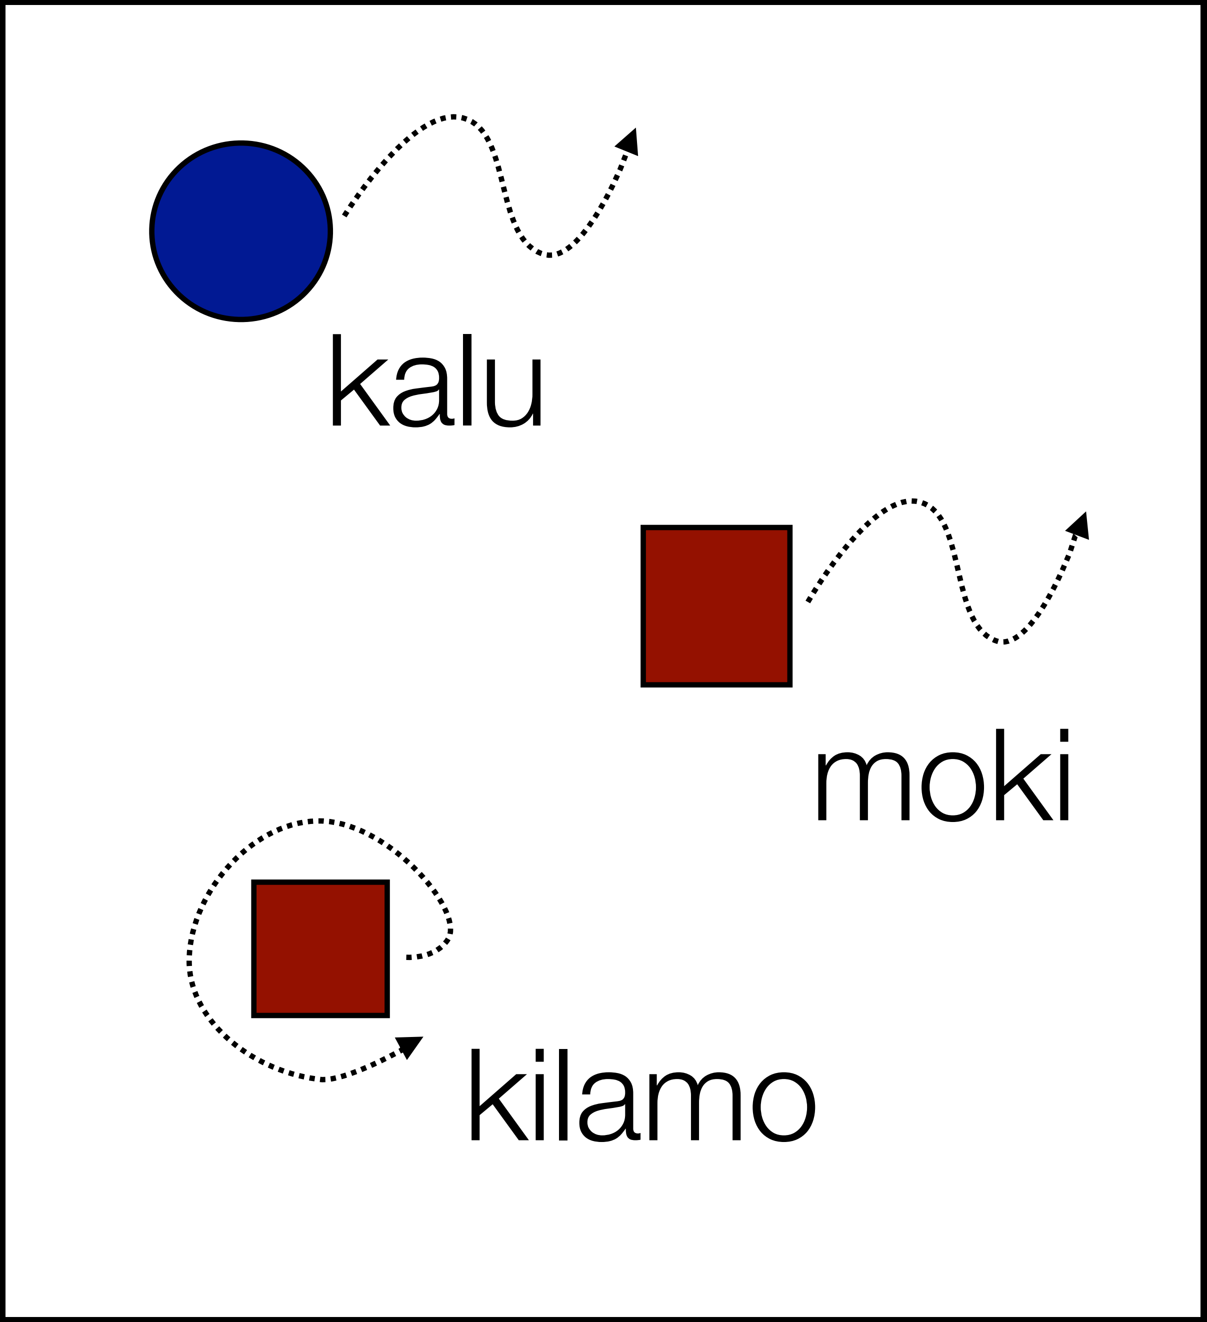

Kirby, Cornish, Smith (2008) *PNAS*

LANGUAGE BECOMES EASIER TO LEARN

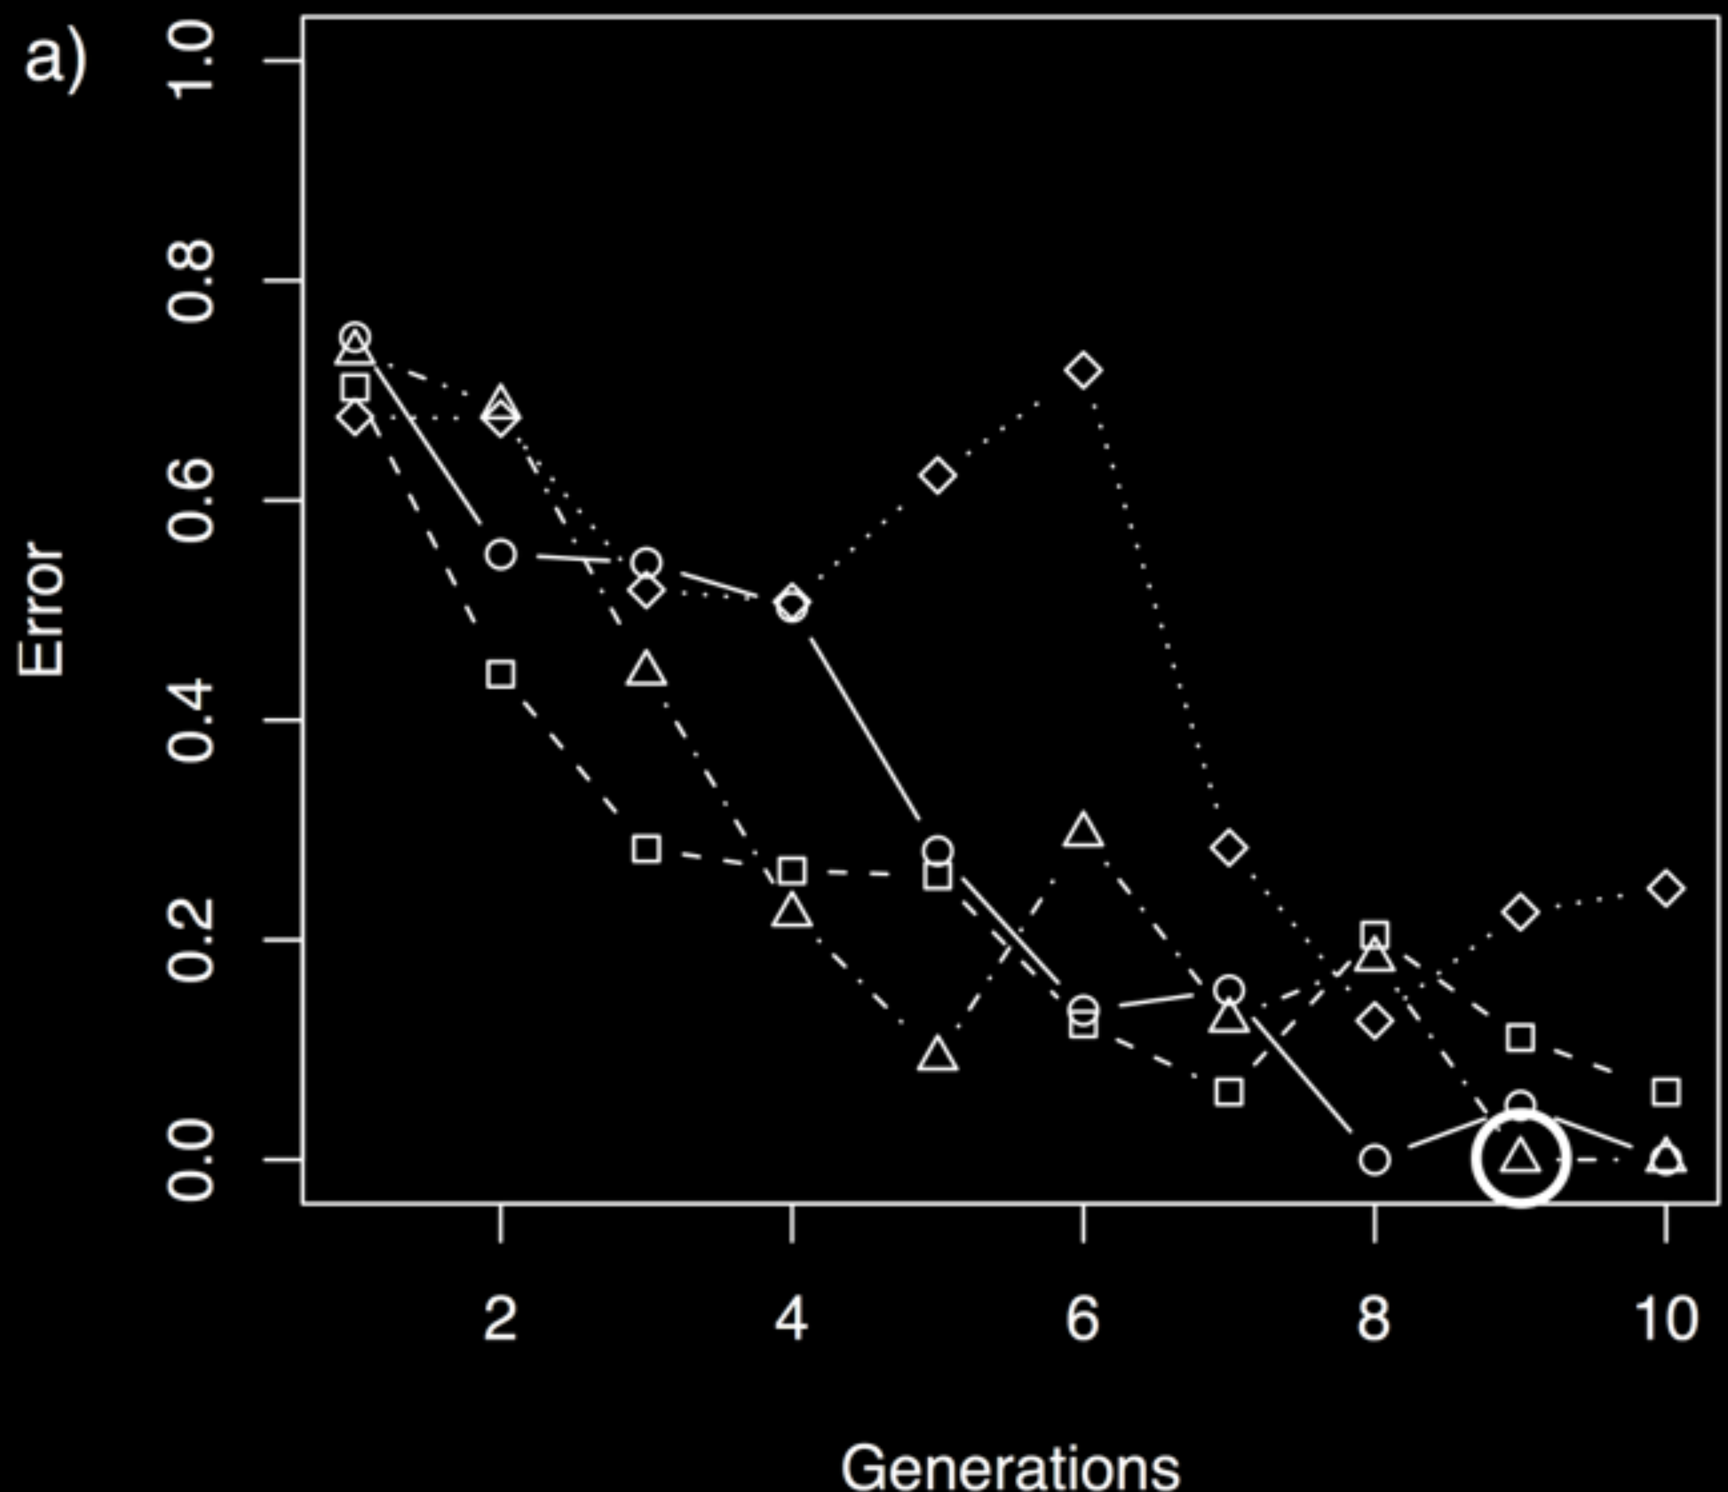

AFTER GENERATION 1

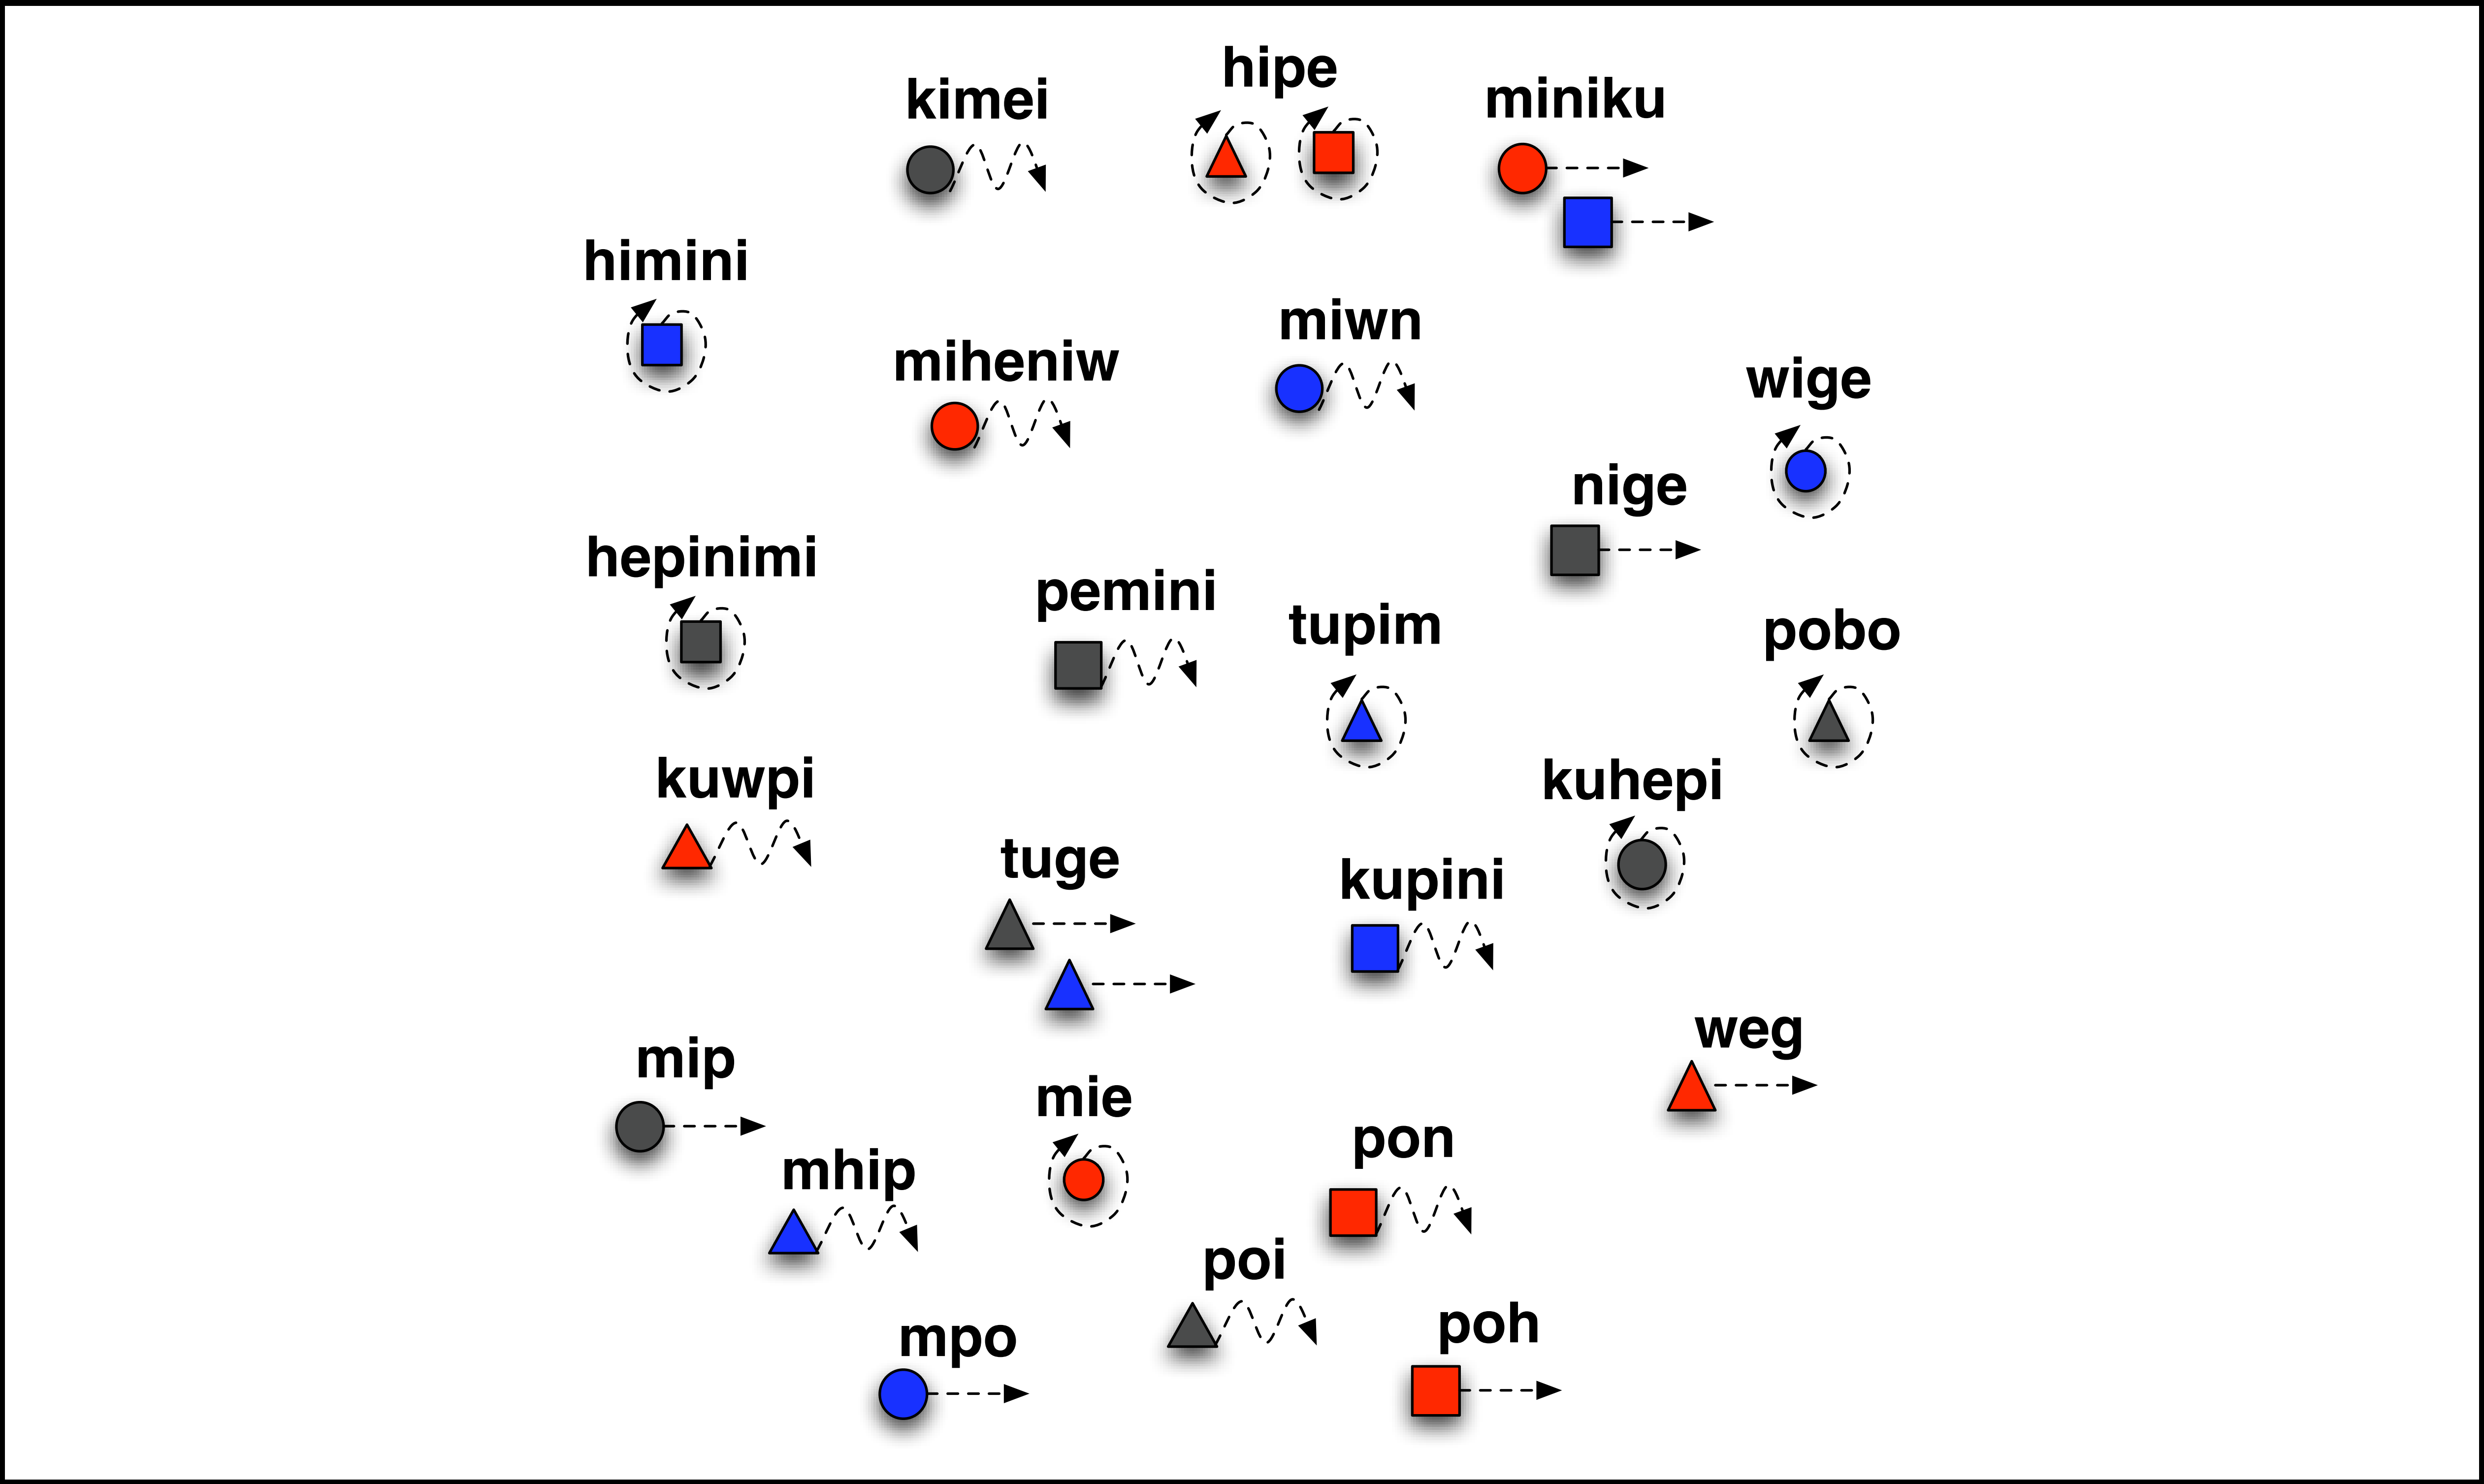

AFTER GENERATION 6

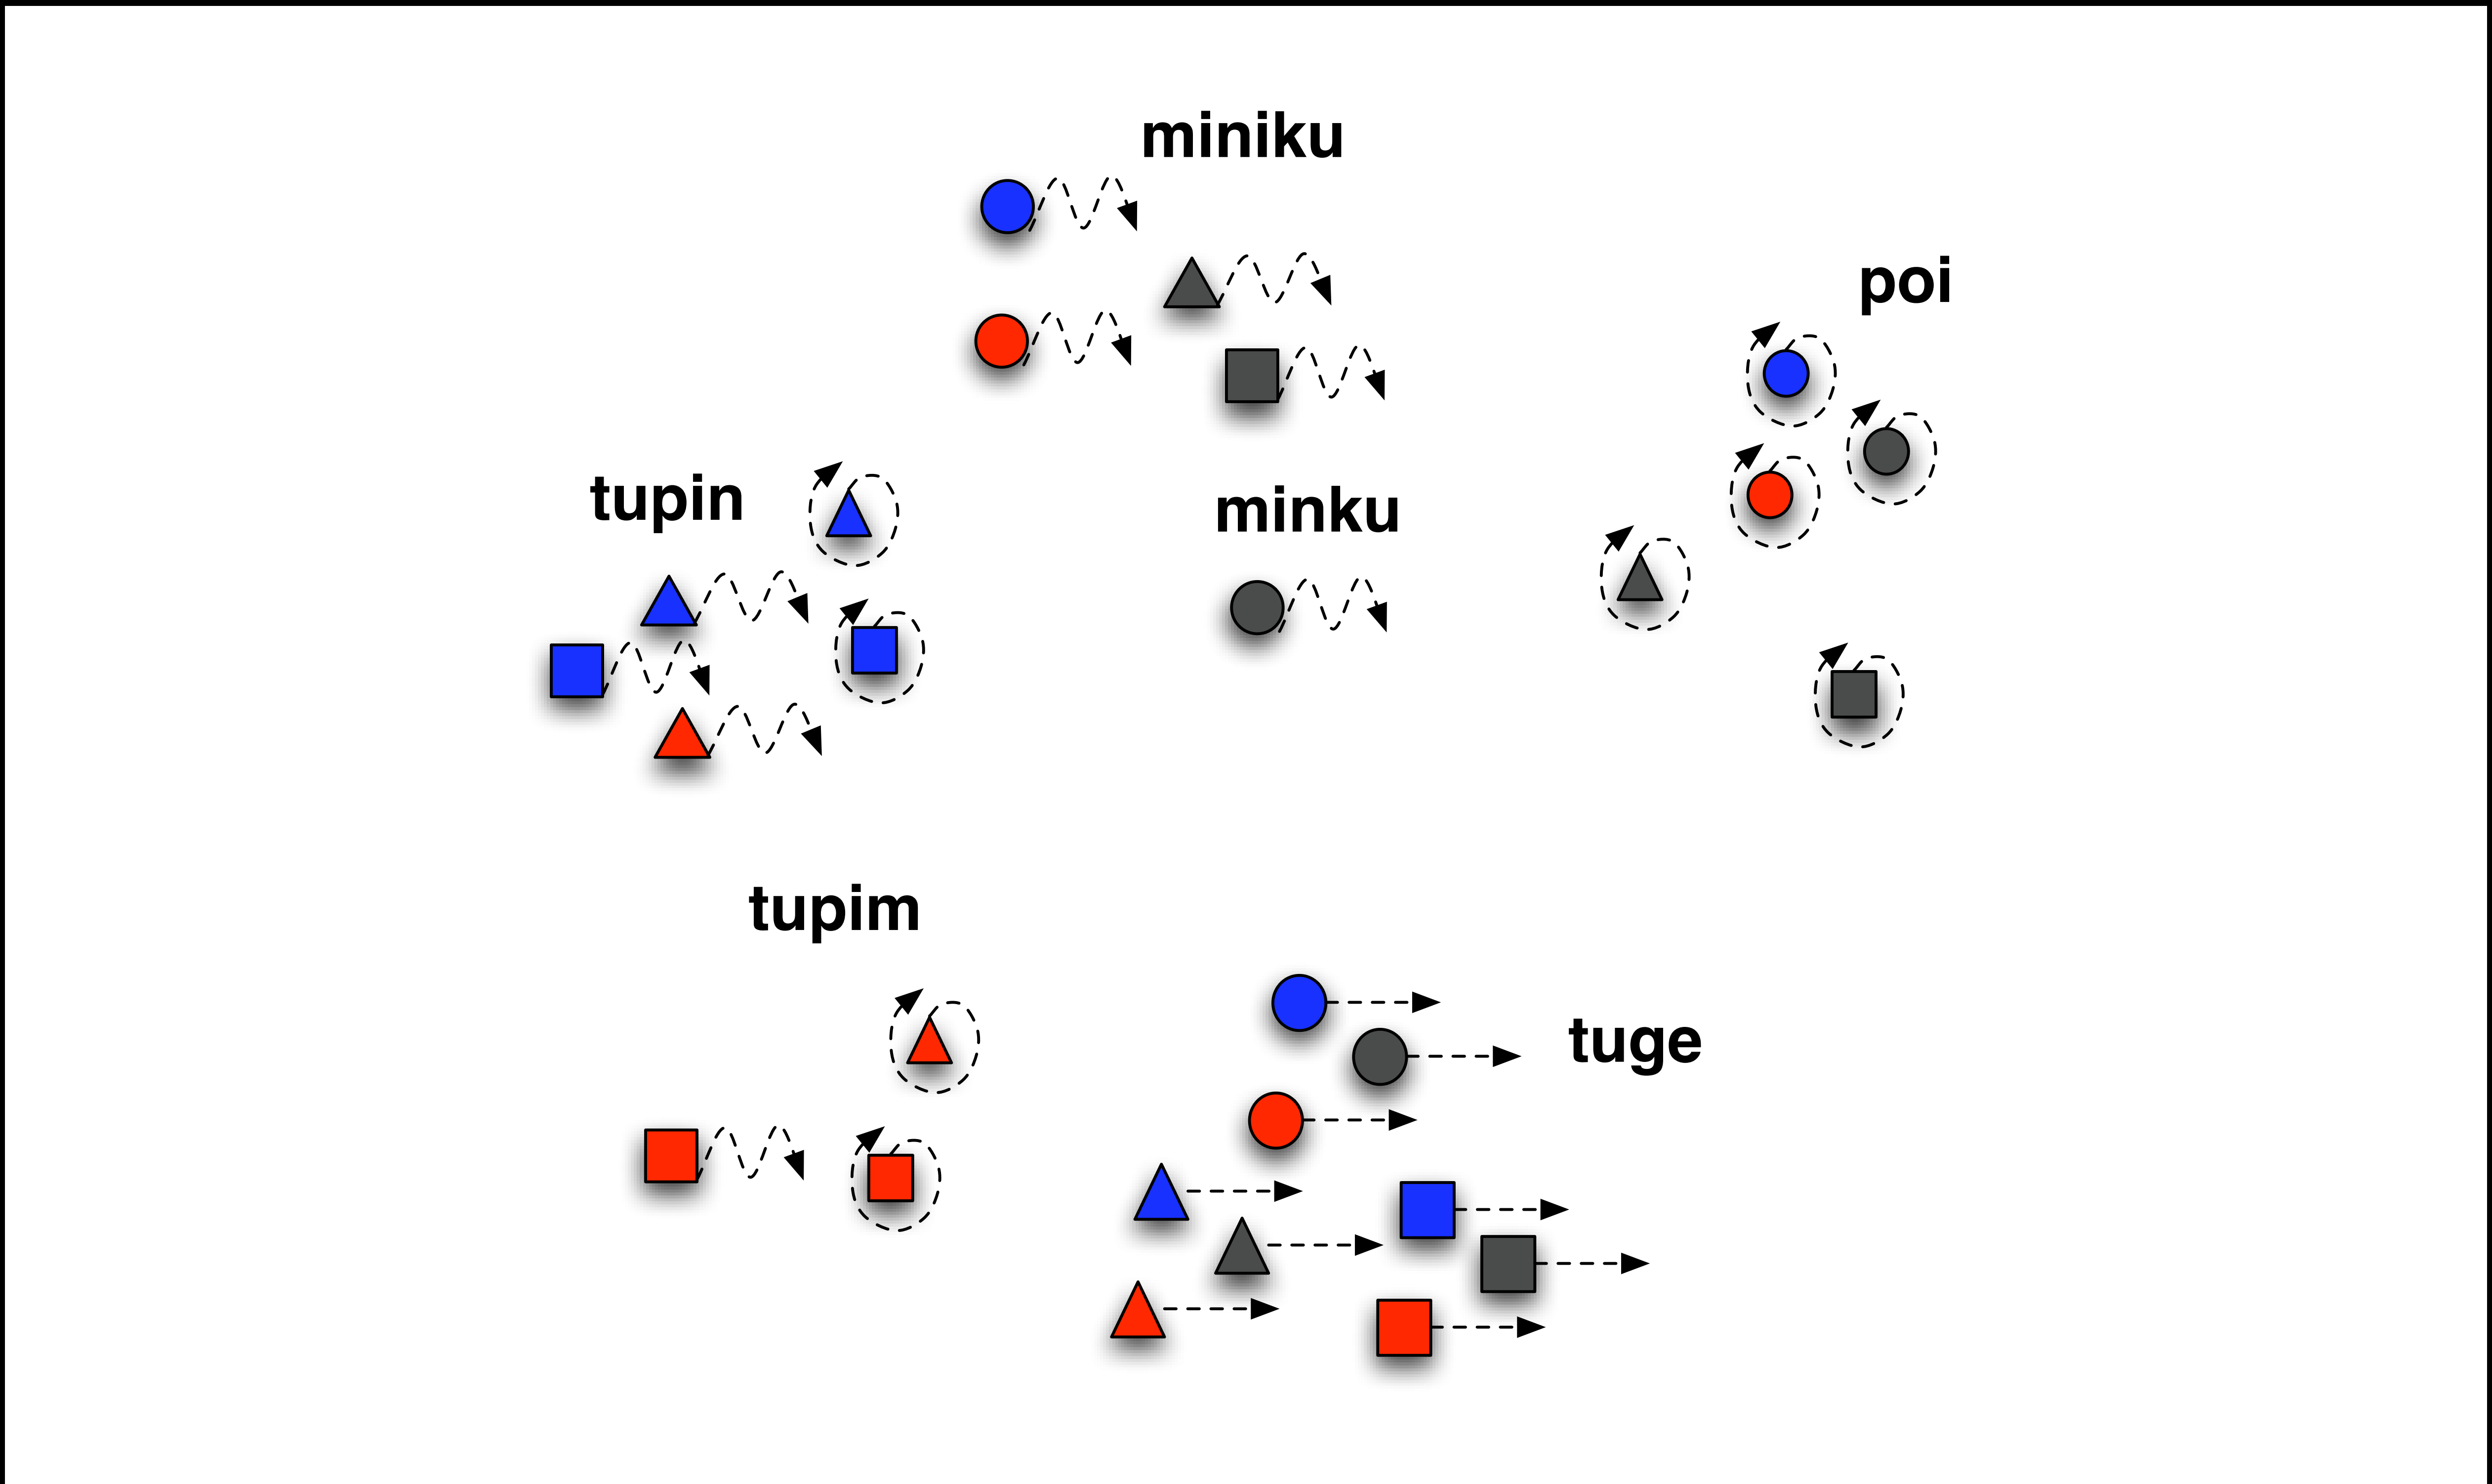

AFTER GENERATION 7

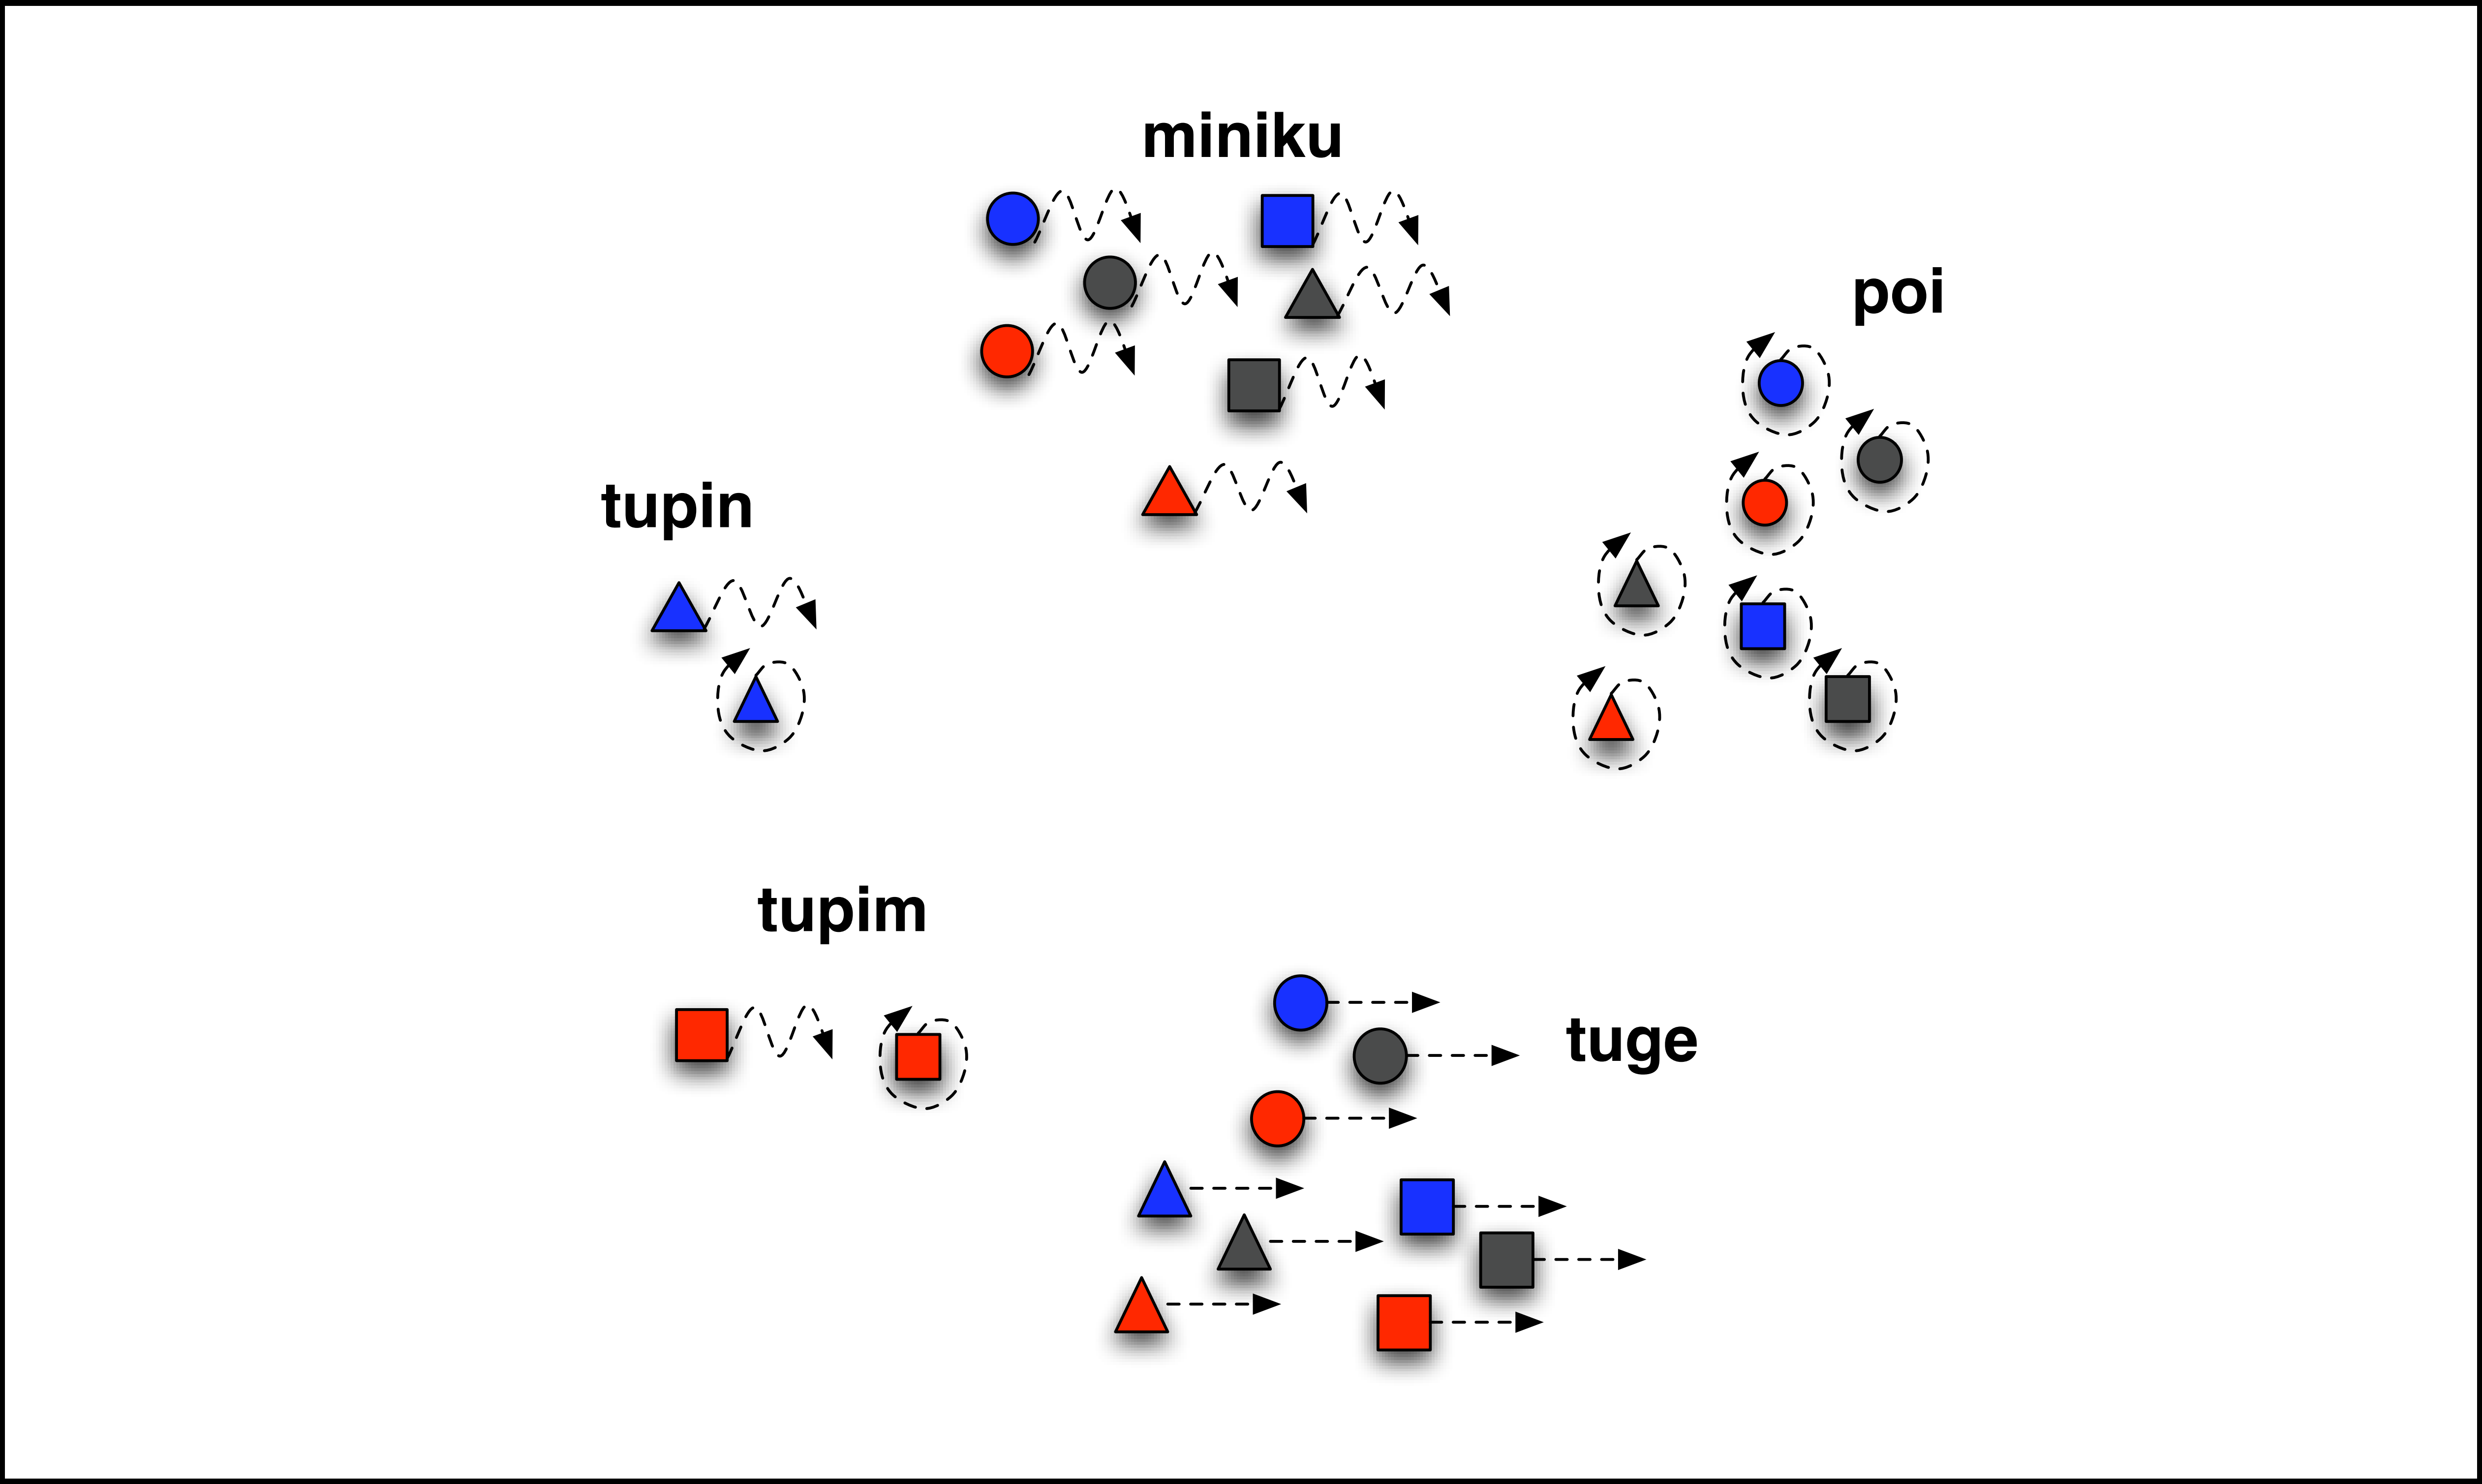

AFTER GENERATION 8

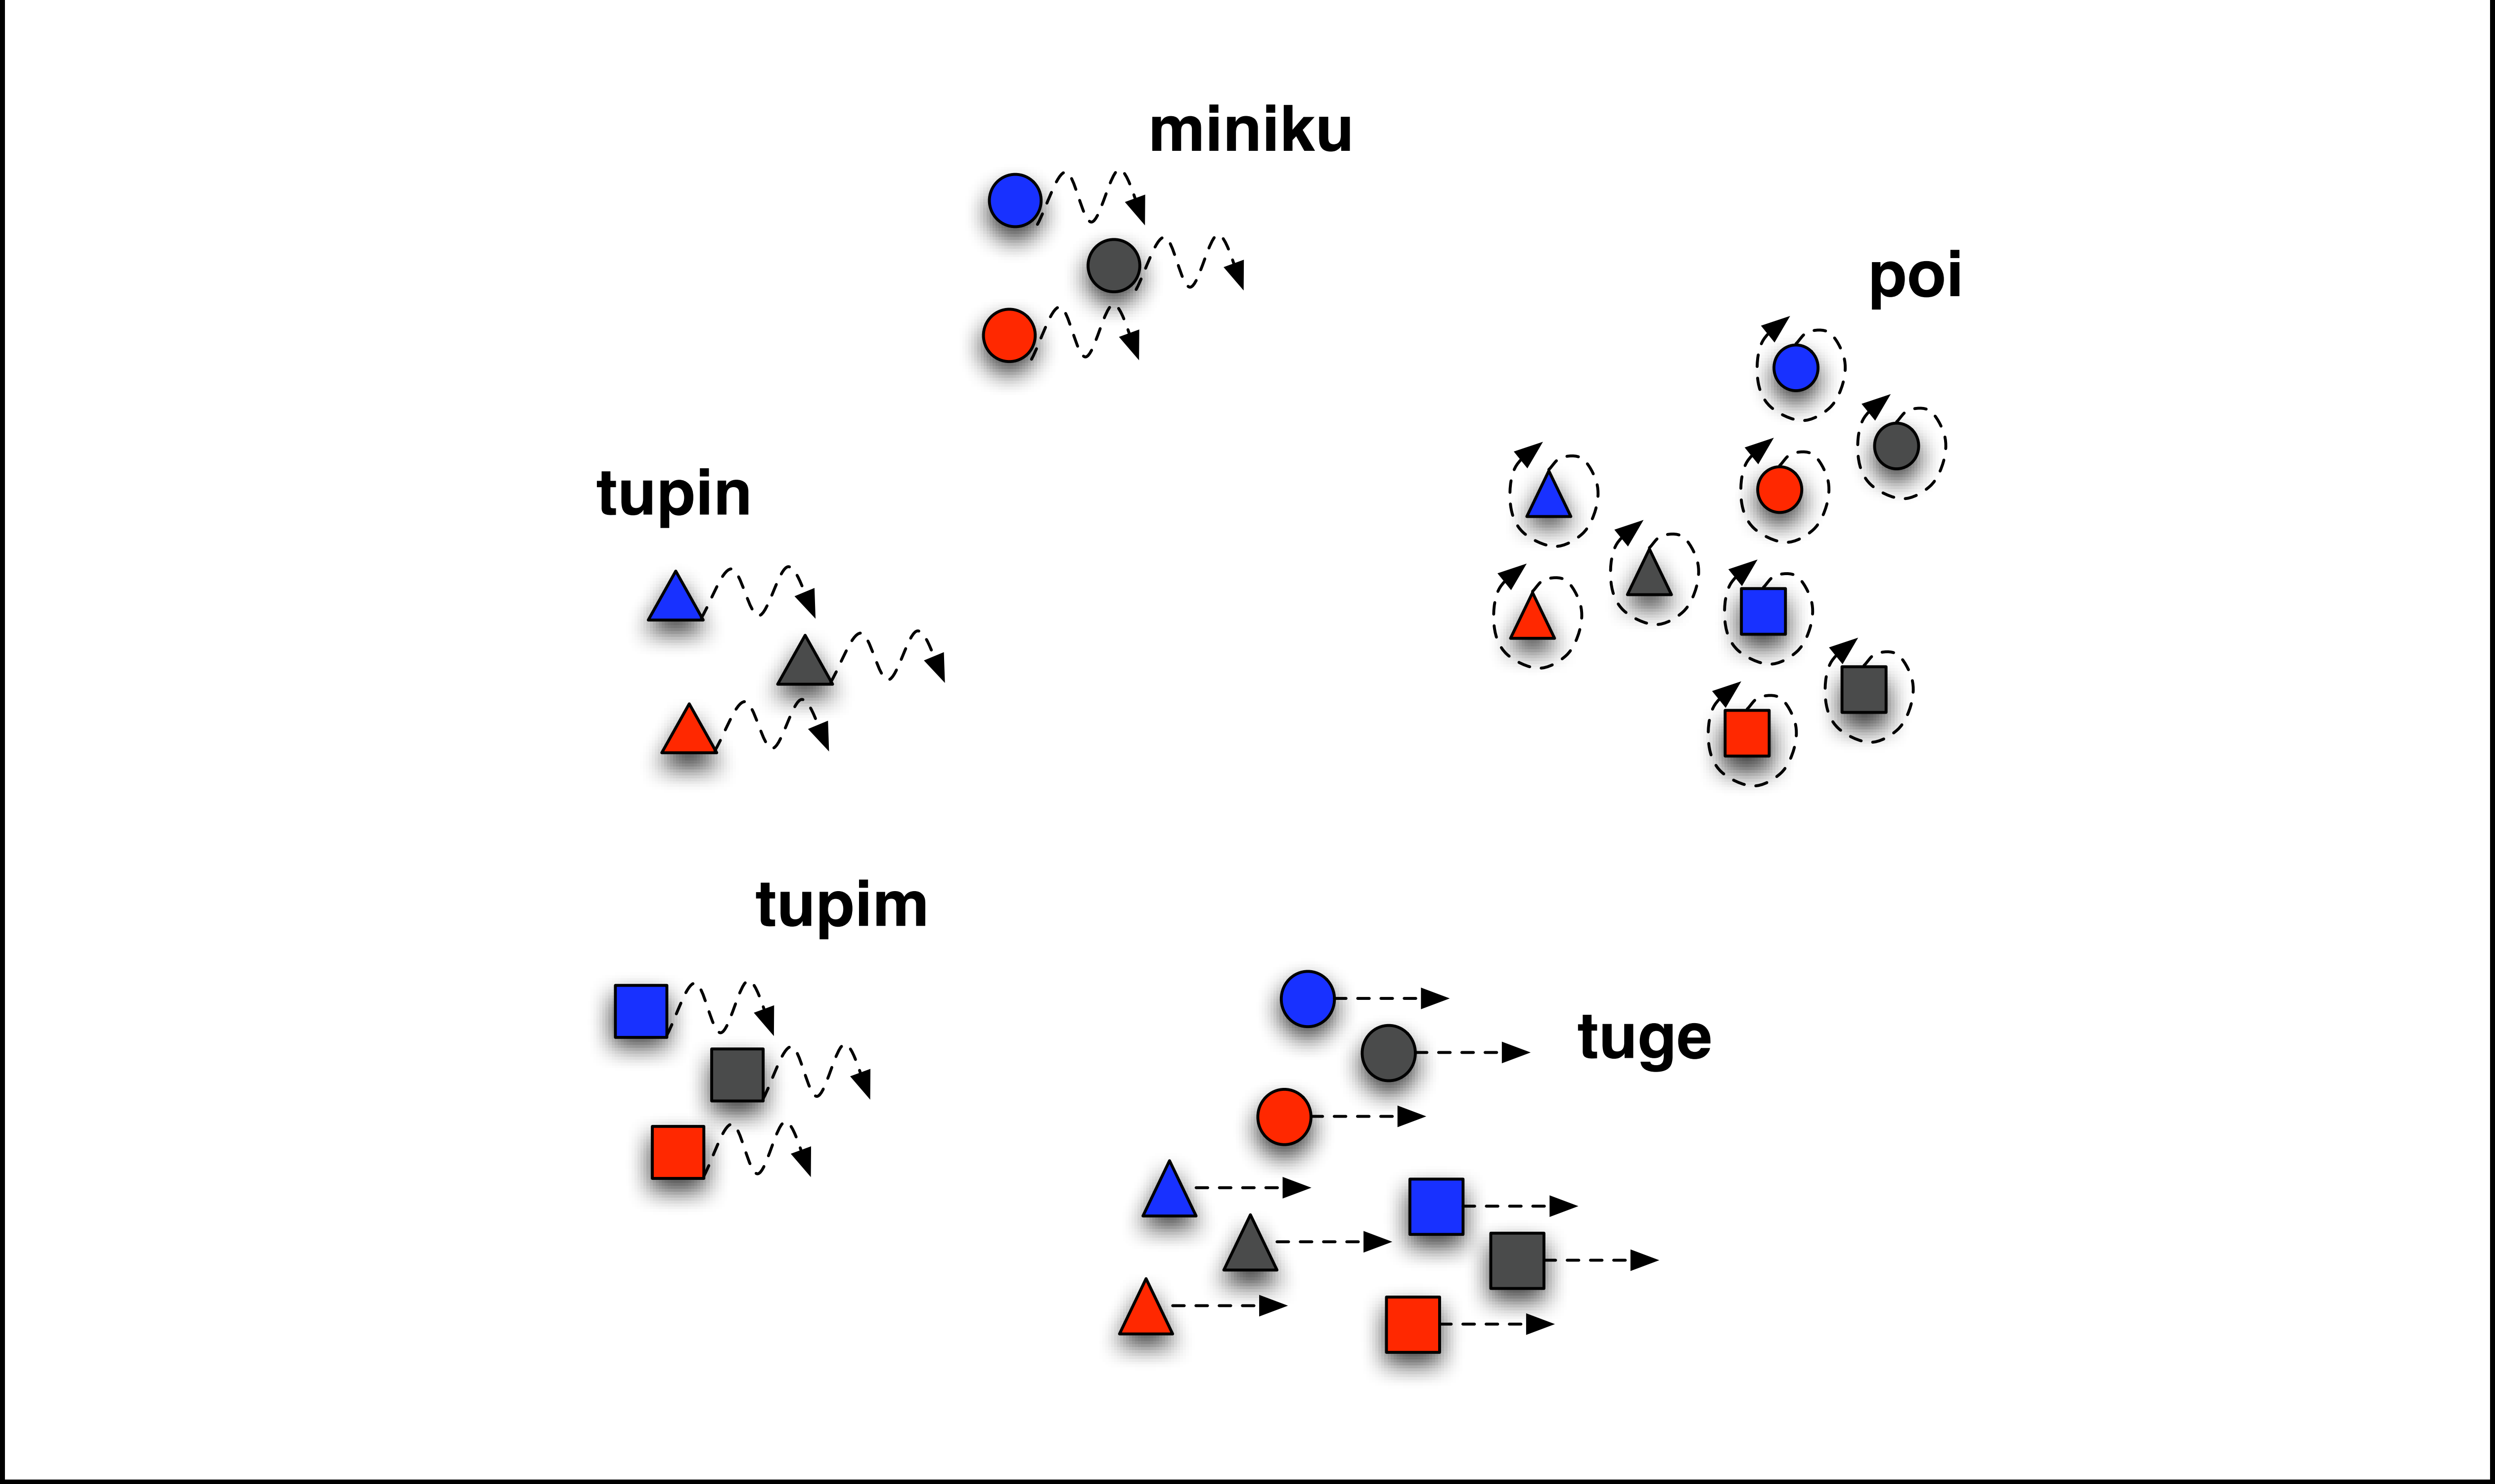

AFTER GENERATION 9

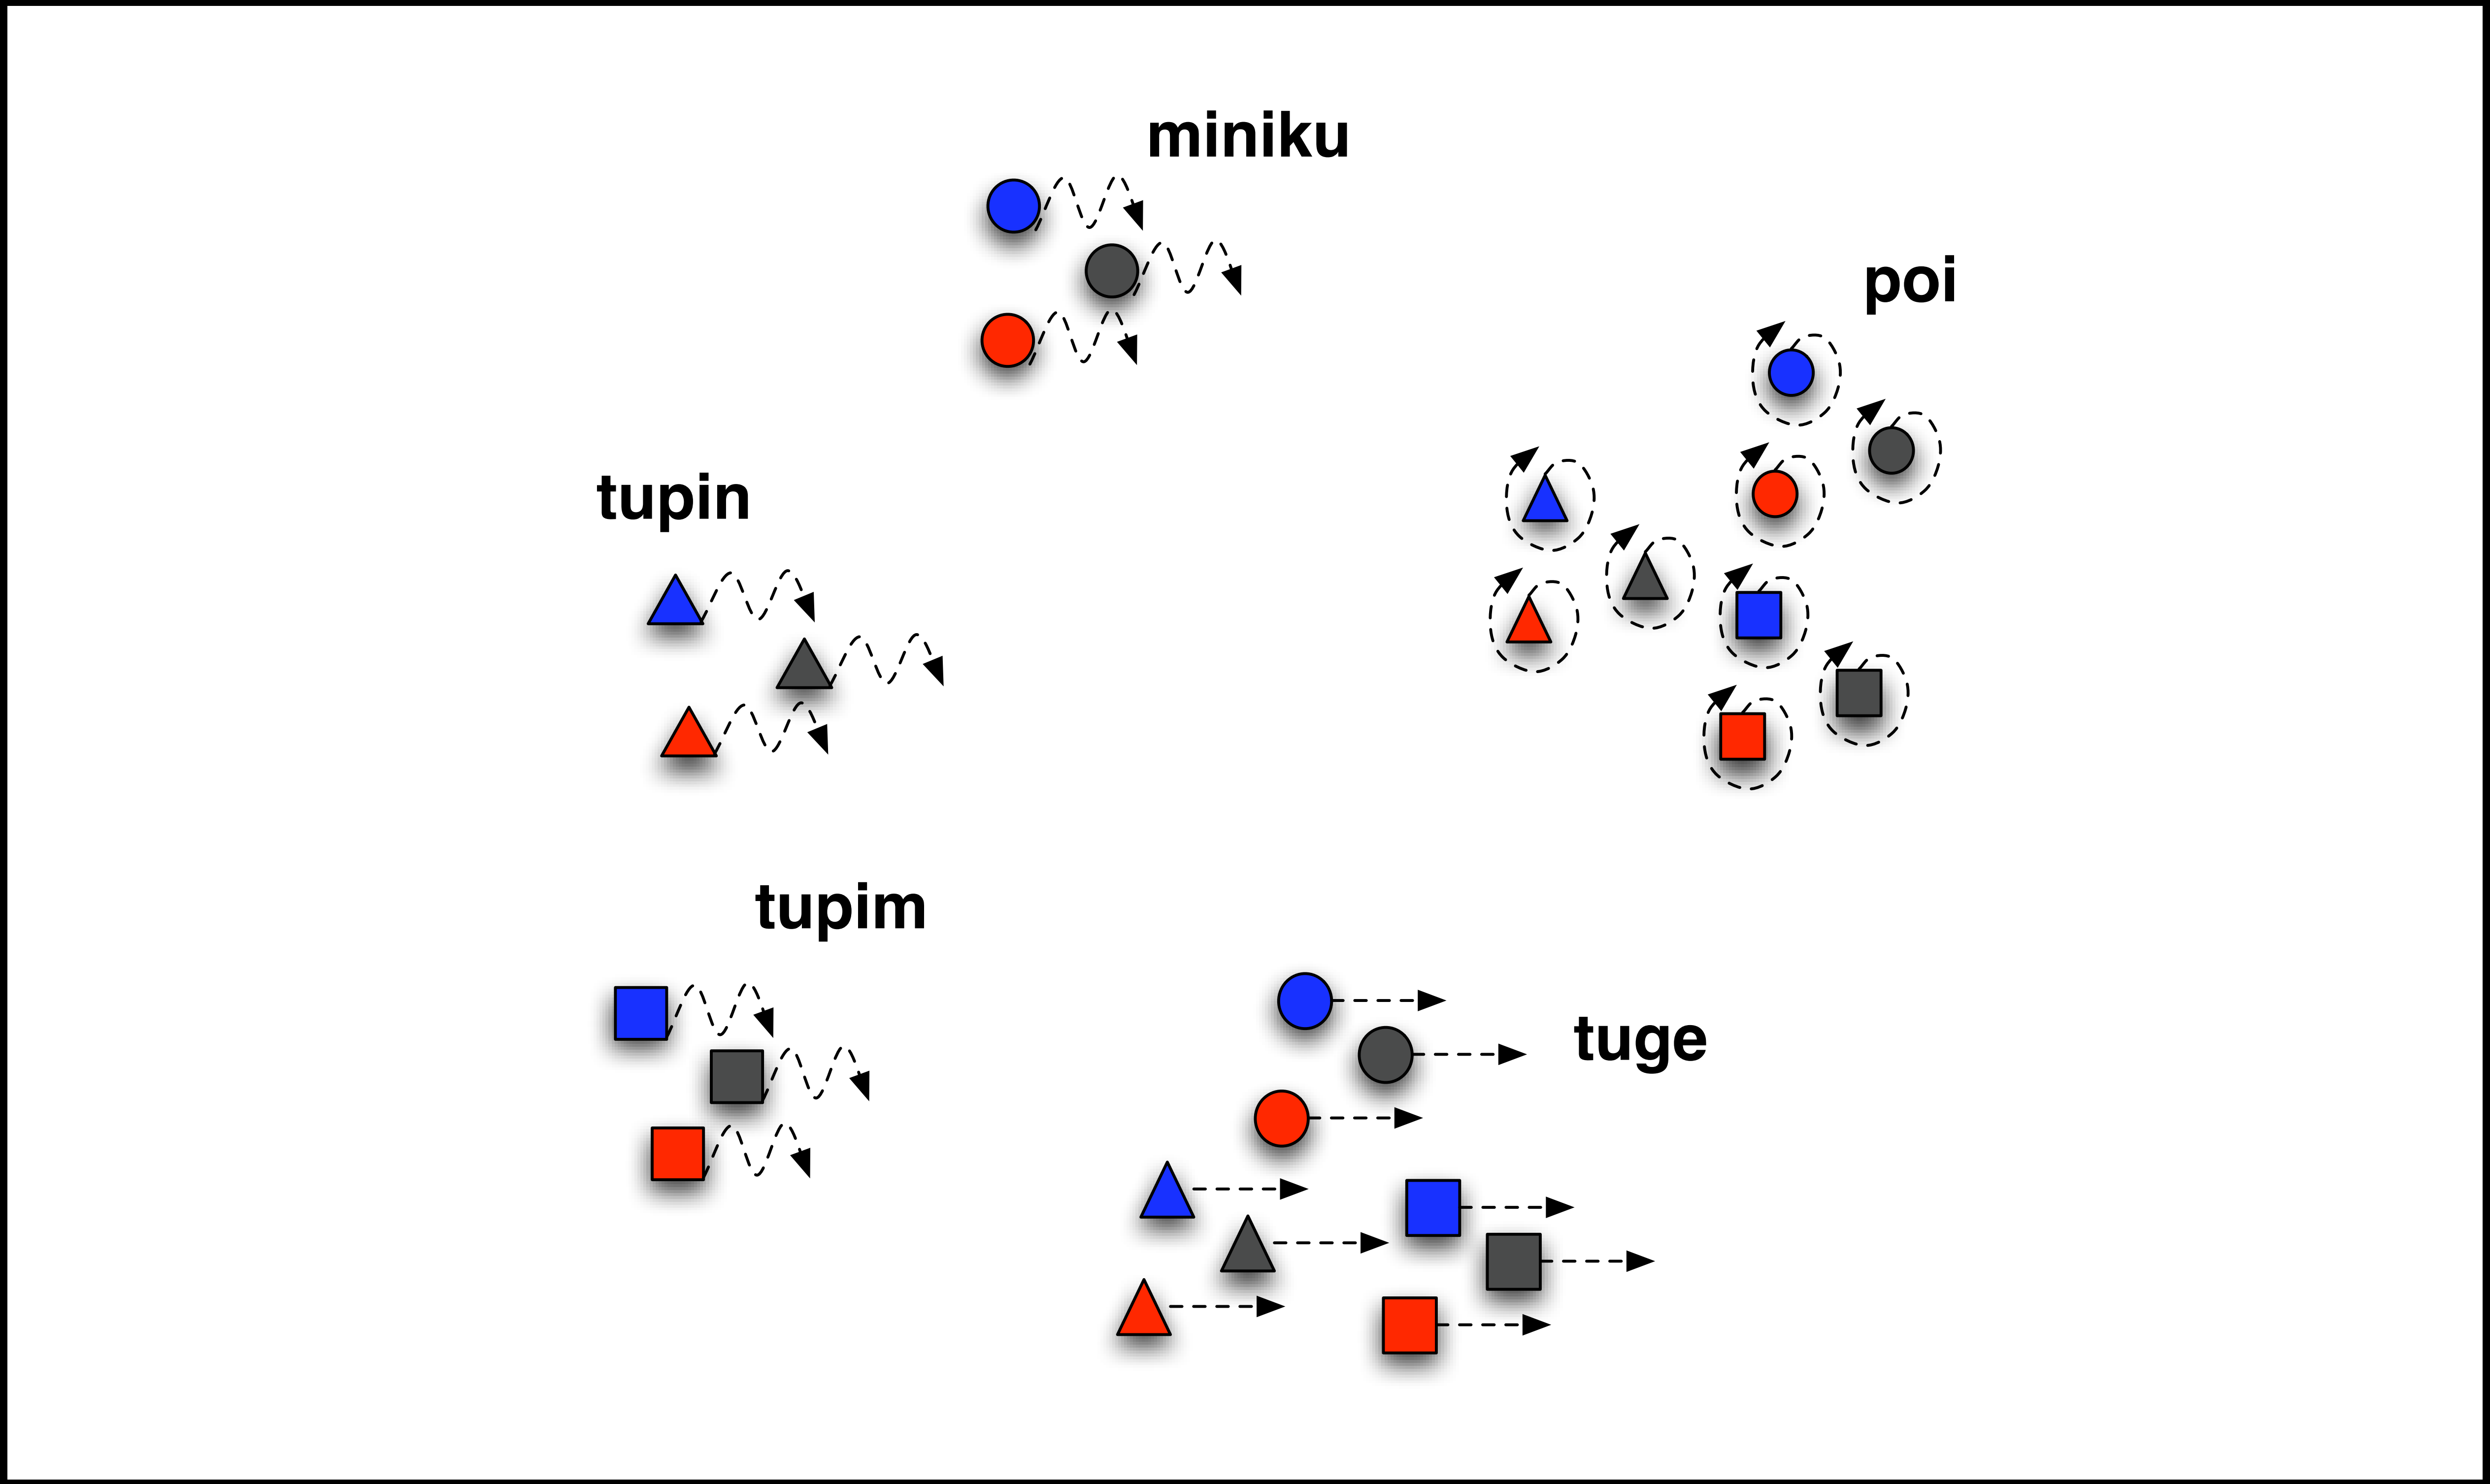

AFTER GENERATION 10

DEGENERATE LANGUAGE

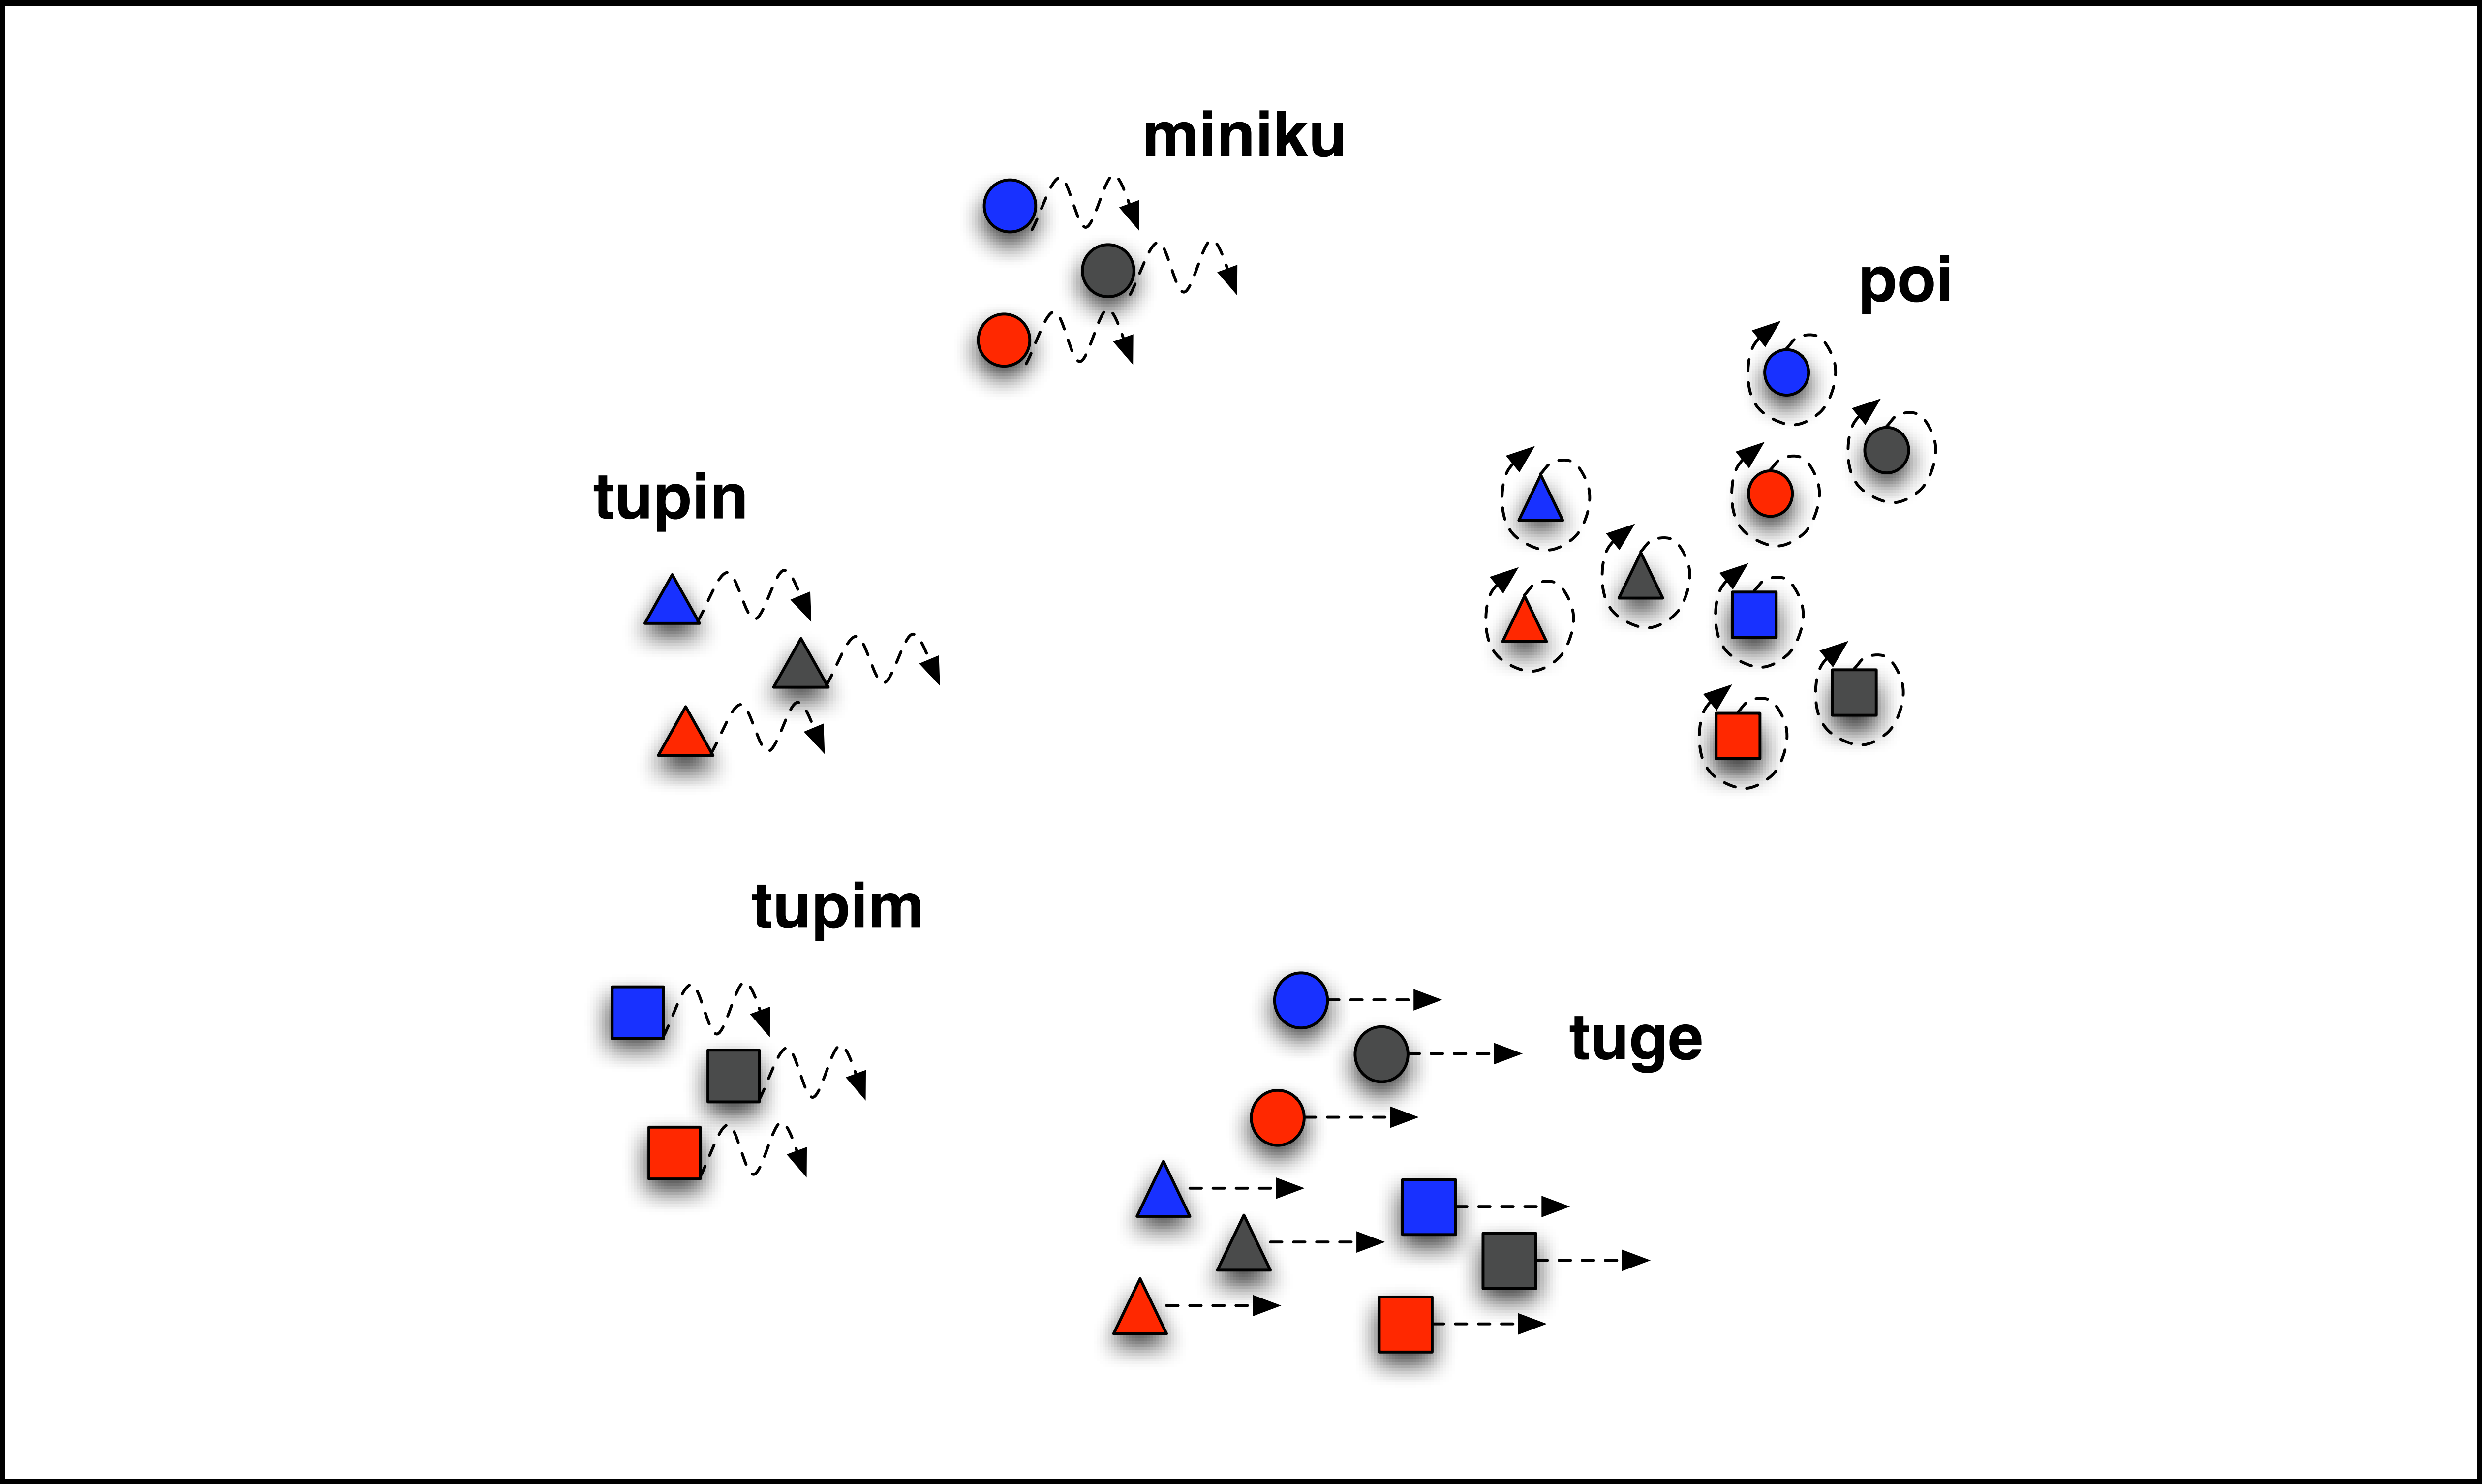

SOMETHING NOT QUITE RIGHT HERE...

Easiest behaviour to learn:
do the same thing in every situation

Easiest language to learn:
one word for every meaning
(n.b. highly *compressible*)

Need to add pressure for *expressivity* to avoid **degenerate** languages

http://www.jeremytheartist.com

ADDING COMMUNICATION

- A new version of the experiment
- Two participants at each generation

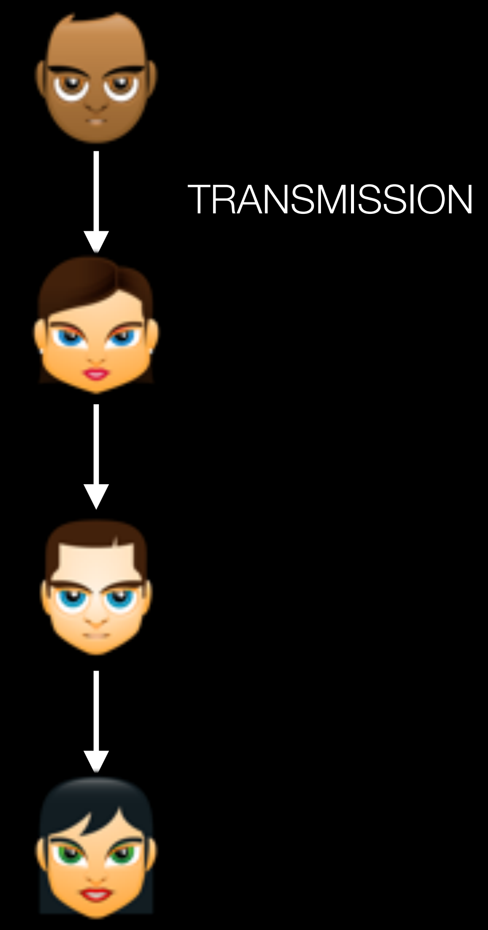

previous design - single learner at each generation

Kirby, Tamariz, Cornish & Smith (in prep)

ADDING COMMUNICATION

- A new version of the experiment
- Two participants at each generation
- Pairs interact to try and pick a particular picture from a distractor array of 6
- The output of the communication game is used to train the next dyad

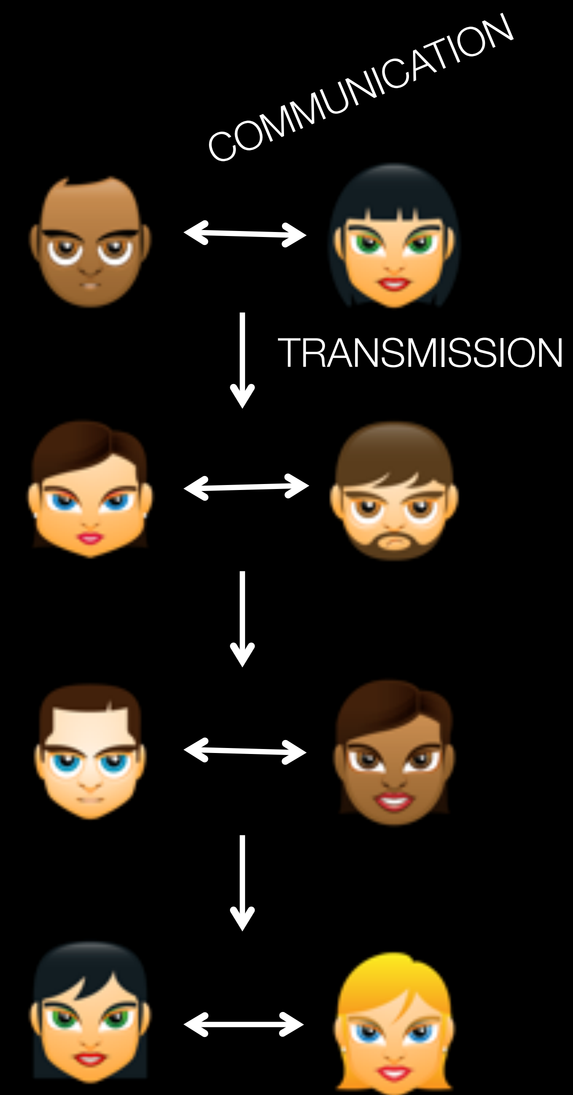

new design - dyads interact at each generation

Kirby, Tamariz, Cornish & Smith (in prep)

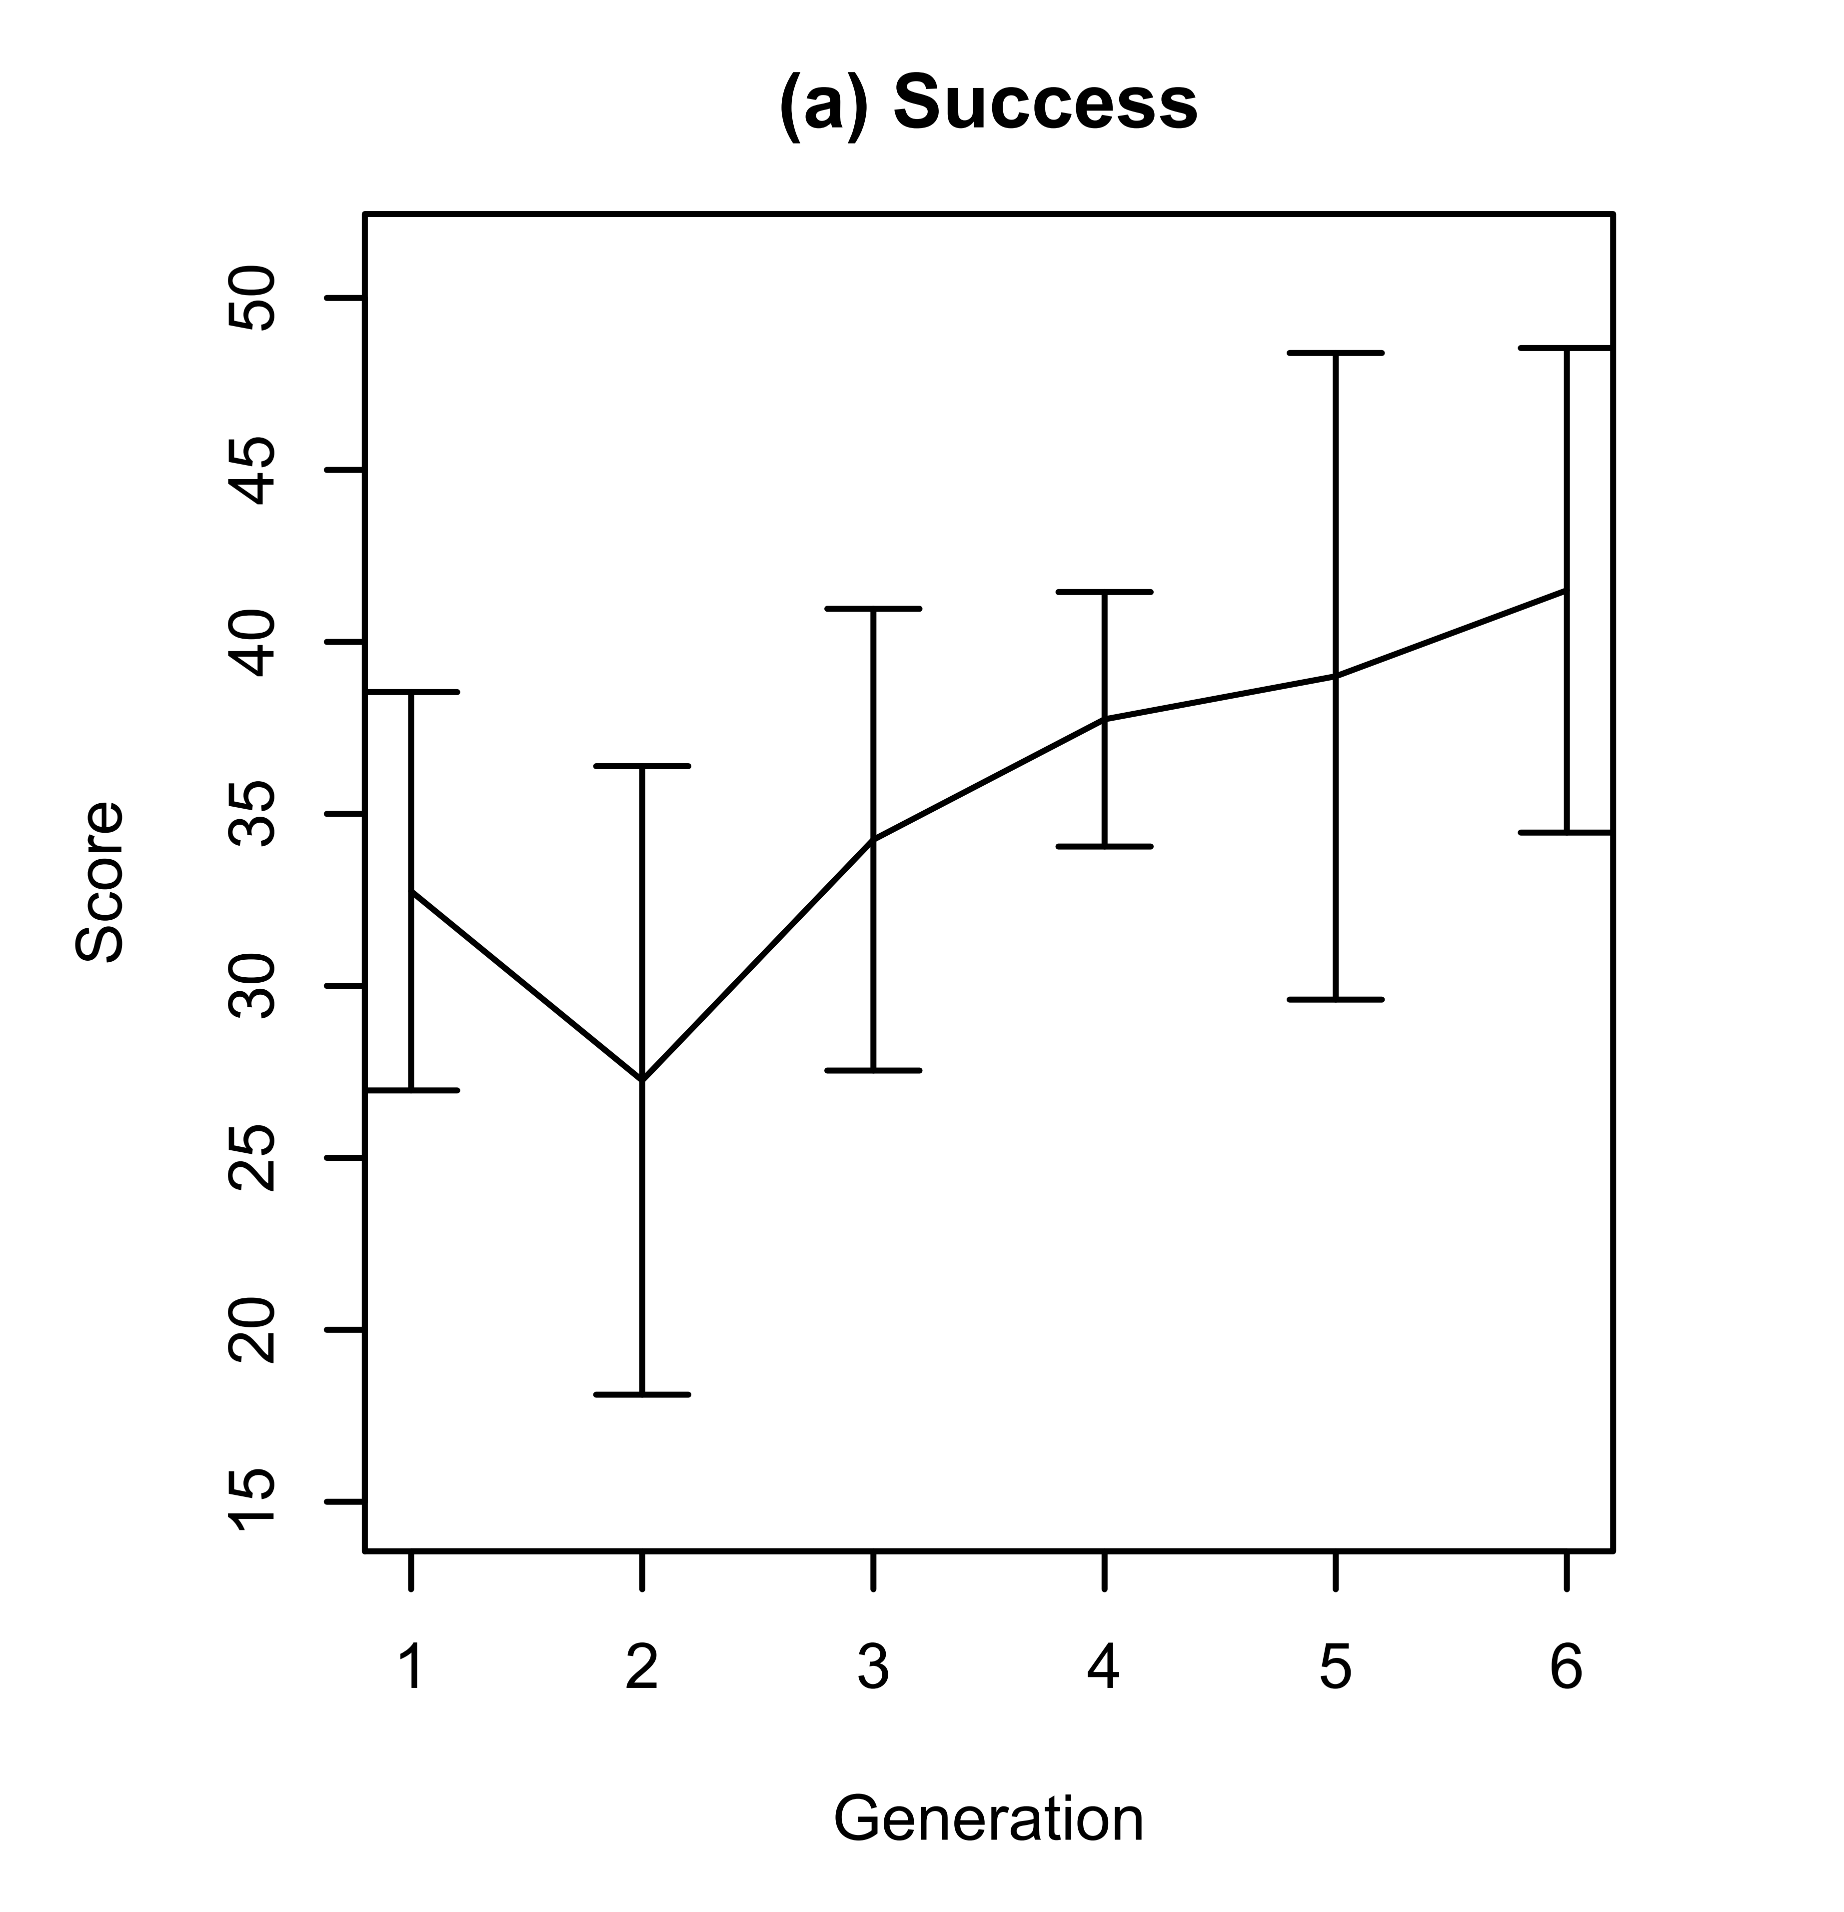
(a) Success
Score
50
45
40
35
30
25
20
15
1
2
3
4
5
6
Generation

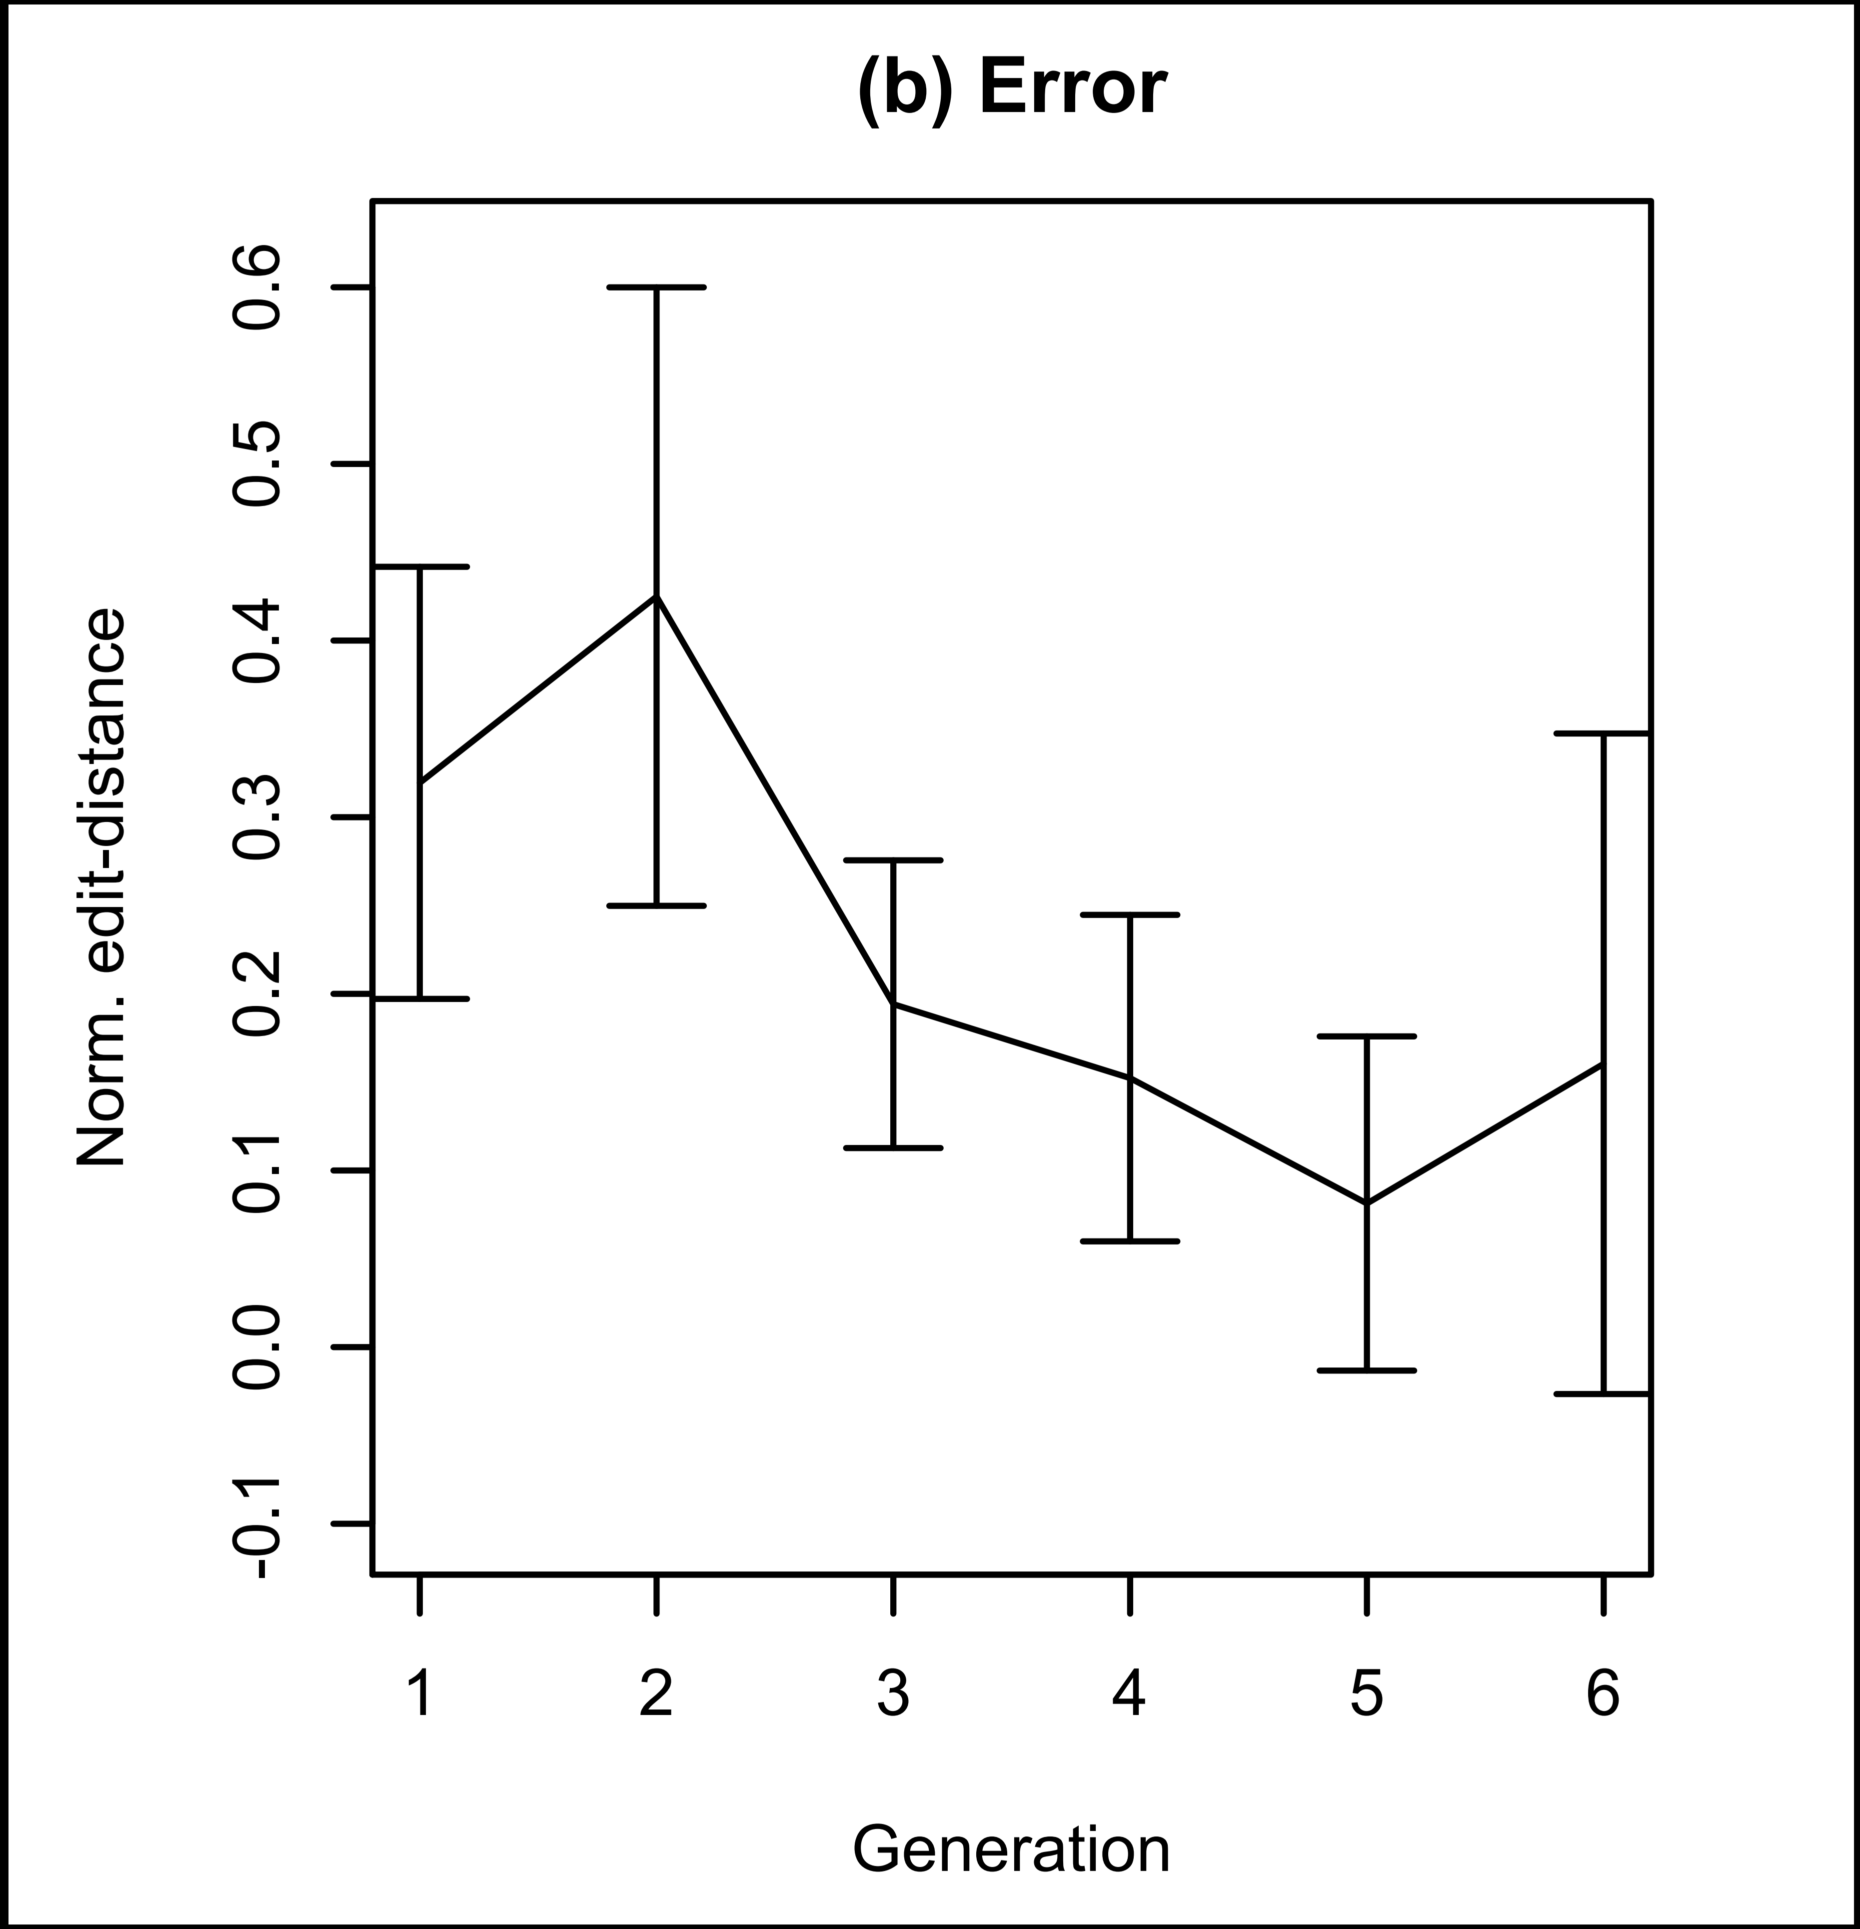
(b) Error
Norm. edit-distance
-0.1
0.0
0.1
0.2
0.3
0.4
0.5
0.6
1
2
3
4
5
6
Generation

INITIAL **HOLISTIC** RANDOM LANGUAGE

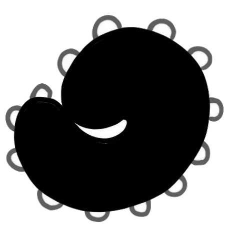 megemume	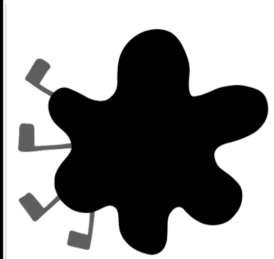 megi	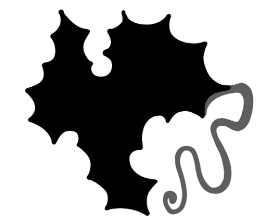 lameme
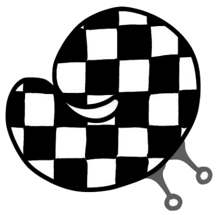 mugimemu	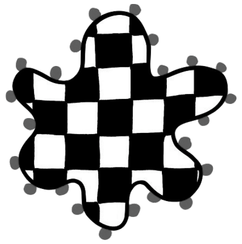 giwulami	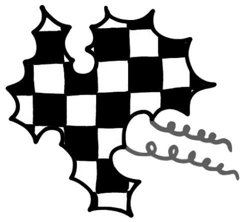 nomenoge
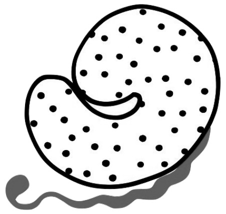 wugi	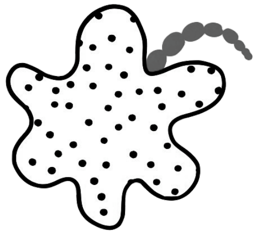 wumume	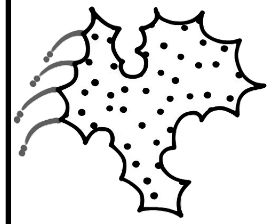 gemulawu
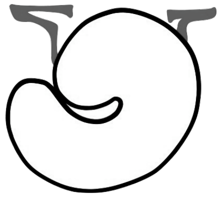 lamege	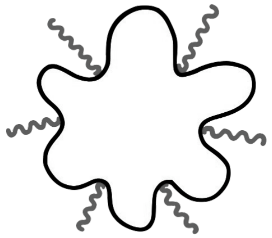 wulamugi	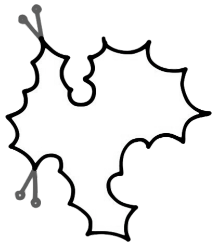 megiwuwa

SIX GENERATIONS LATER:
STRUCTURED (COMPOSITIONAL) LANGUAGE

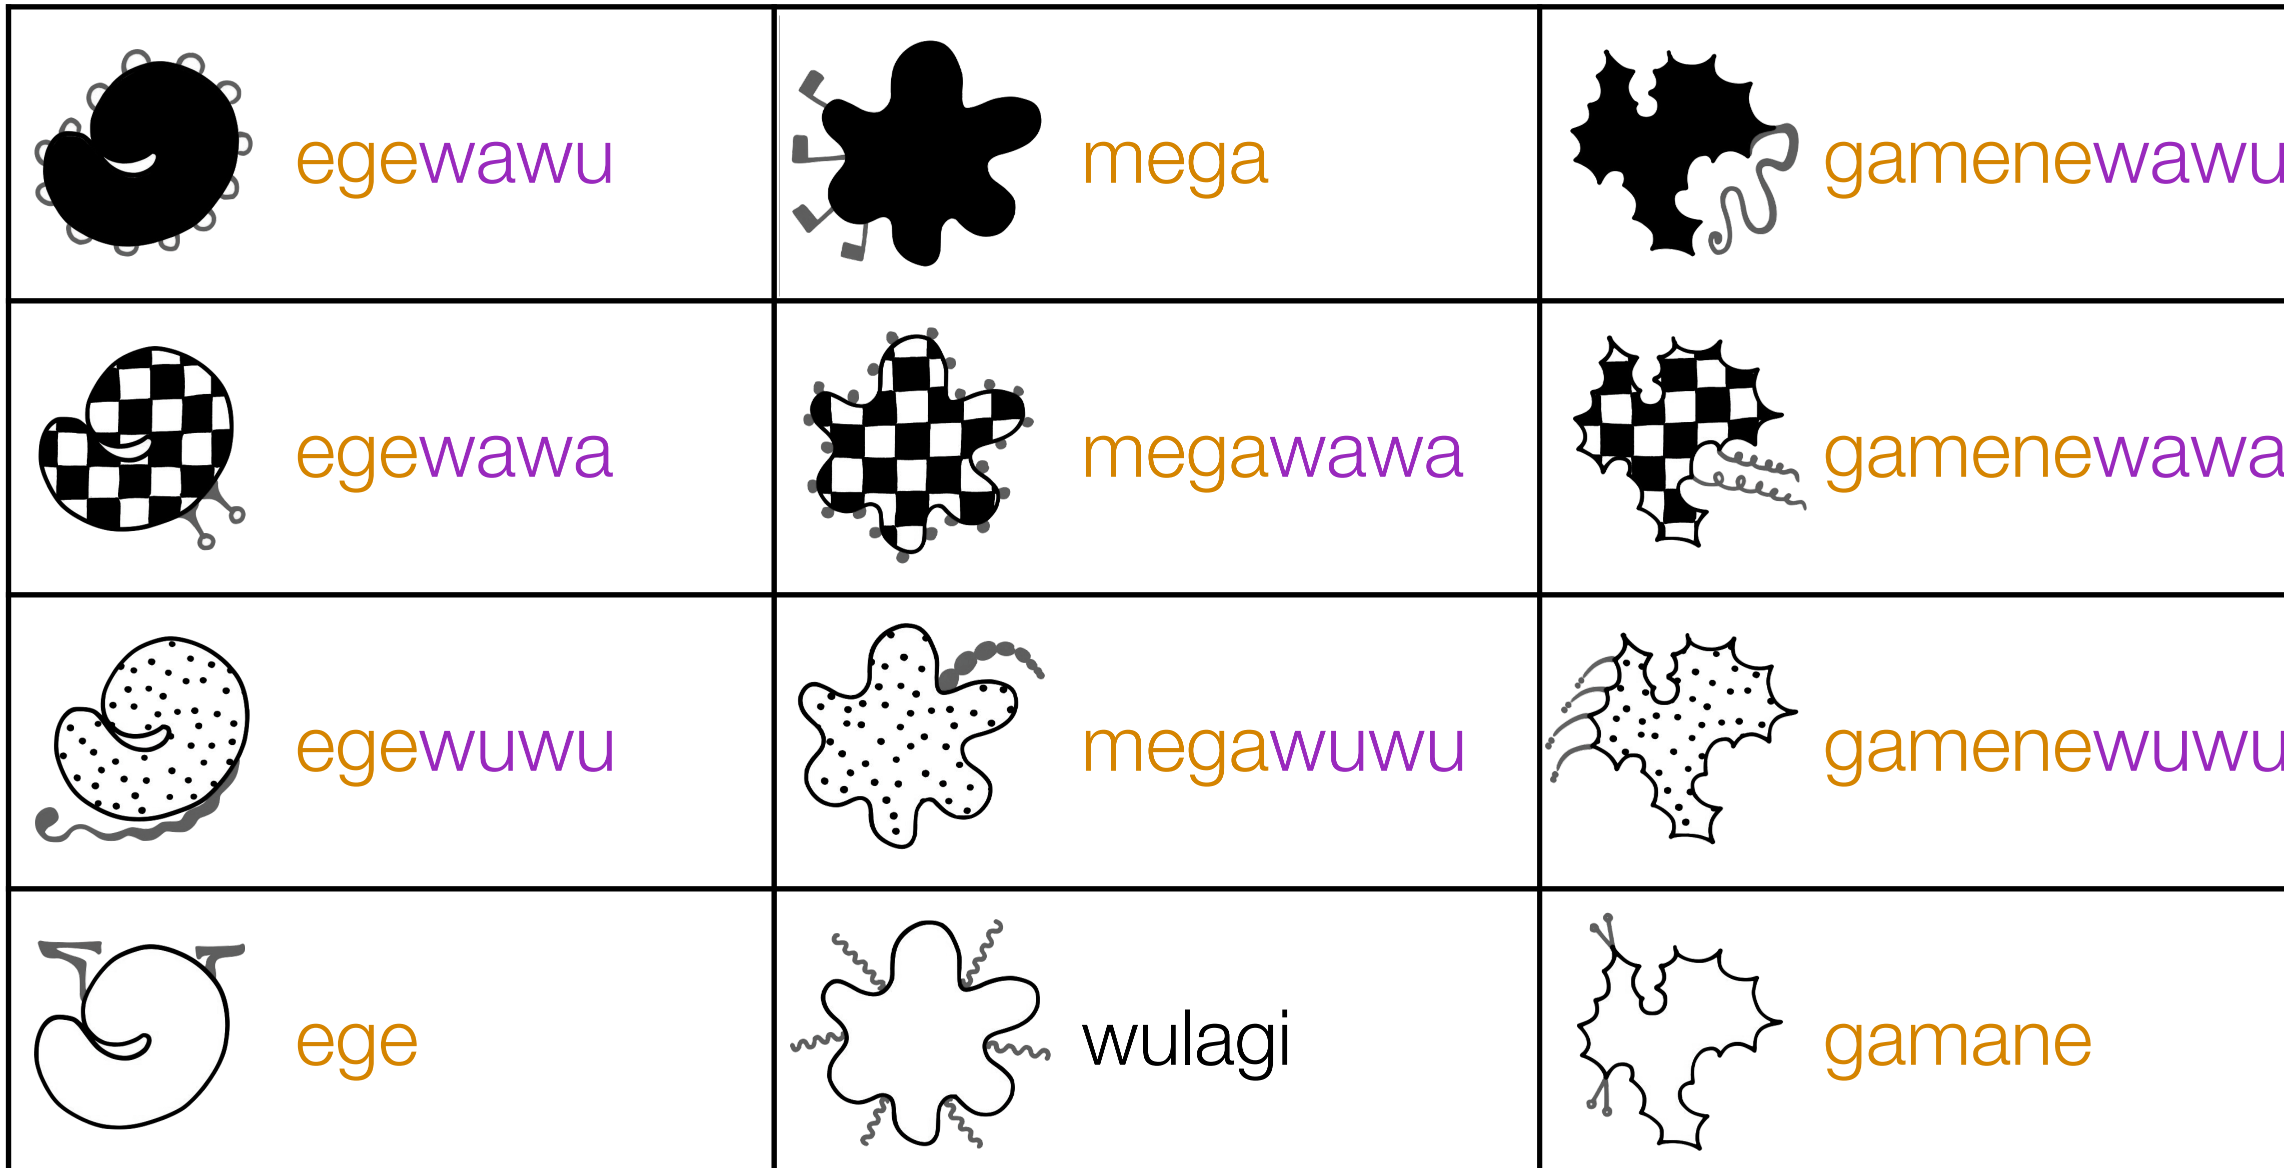

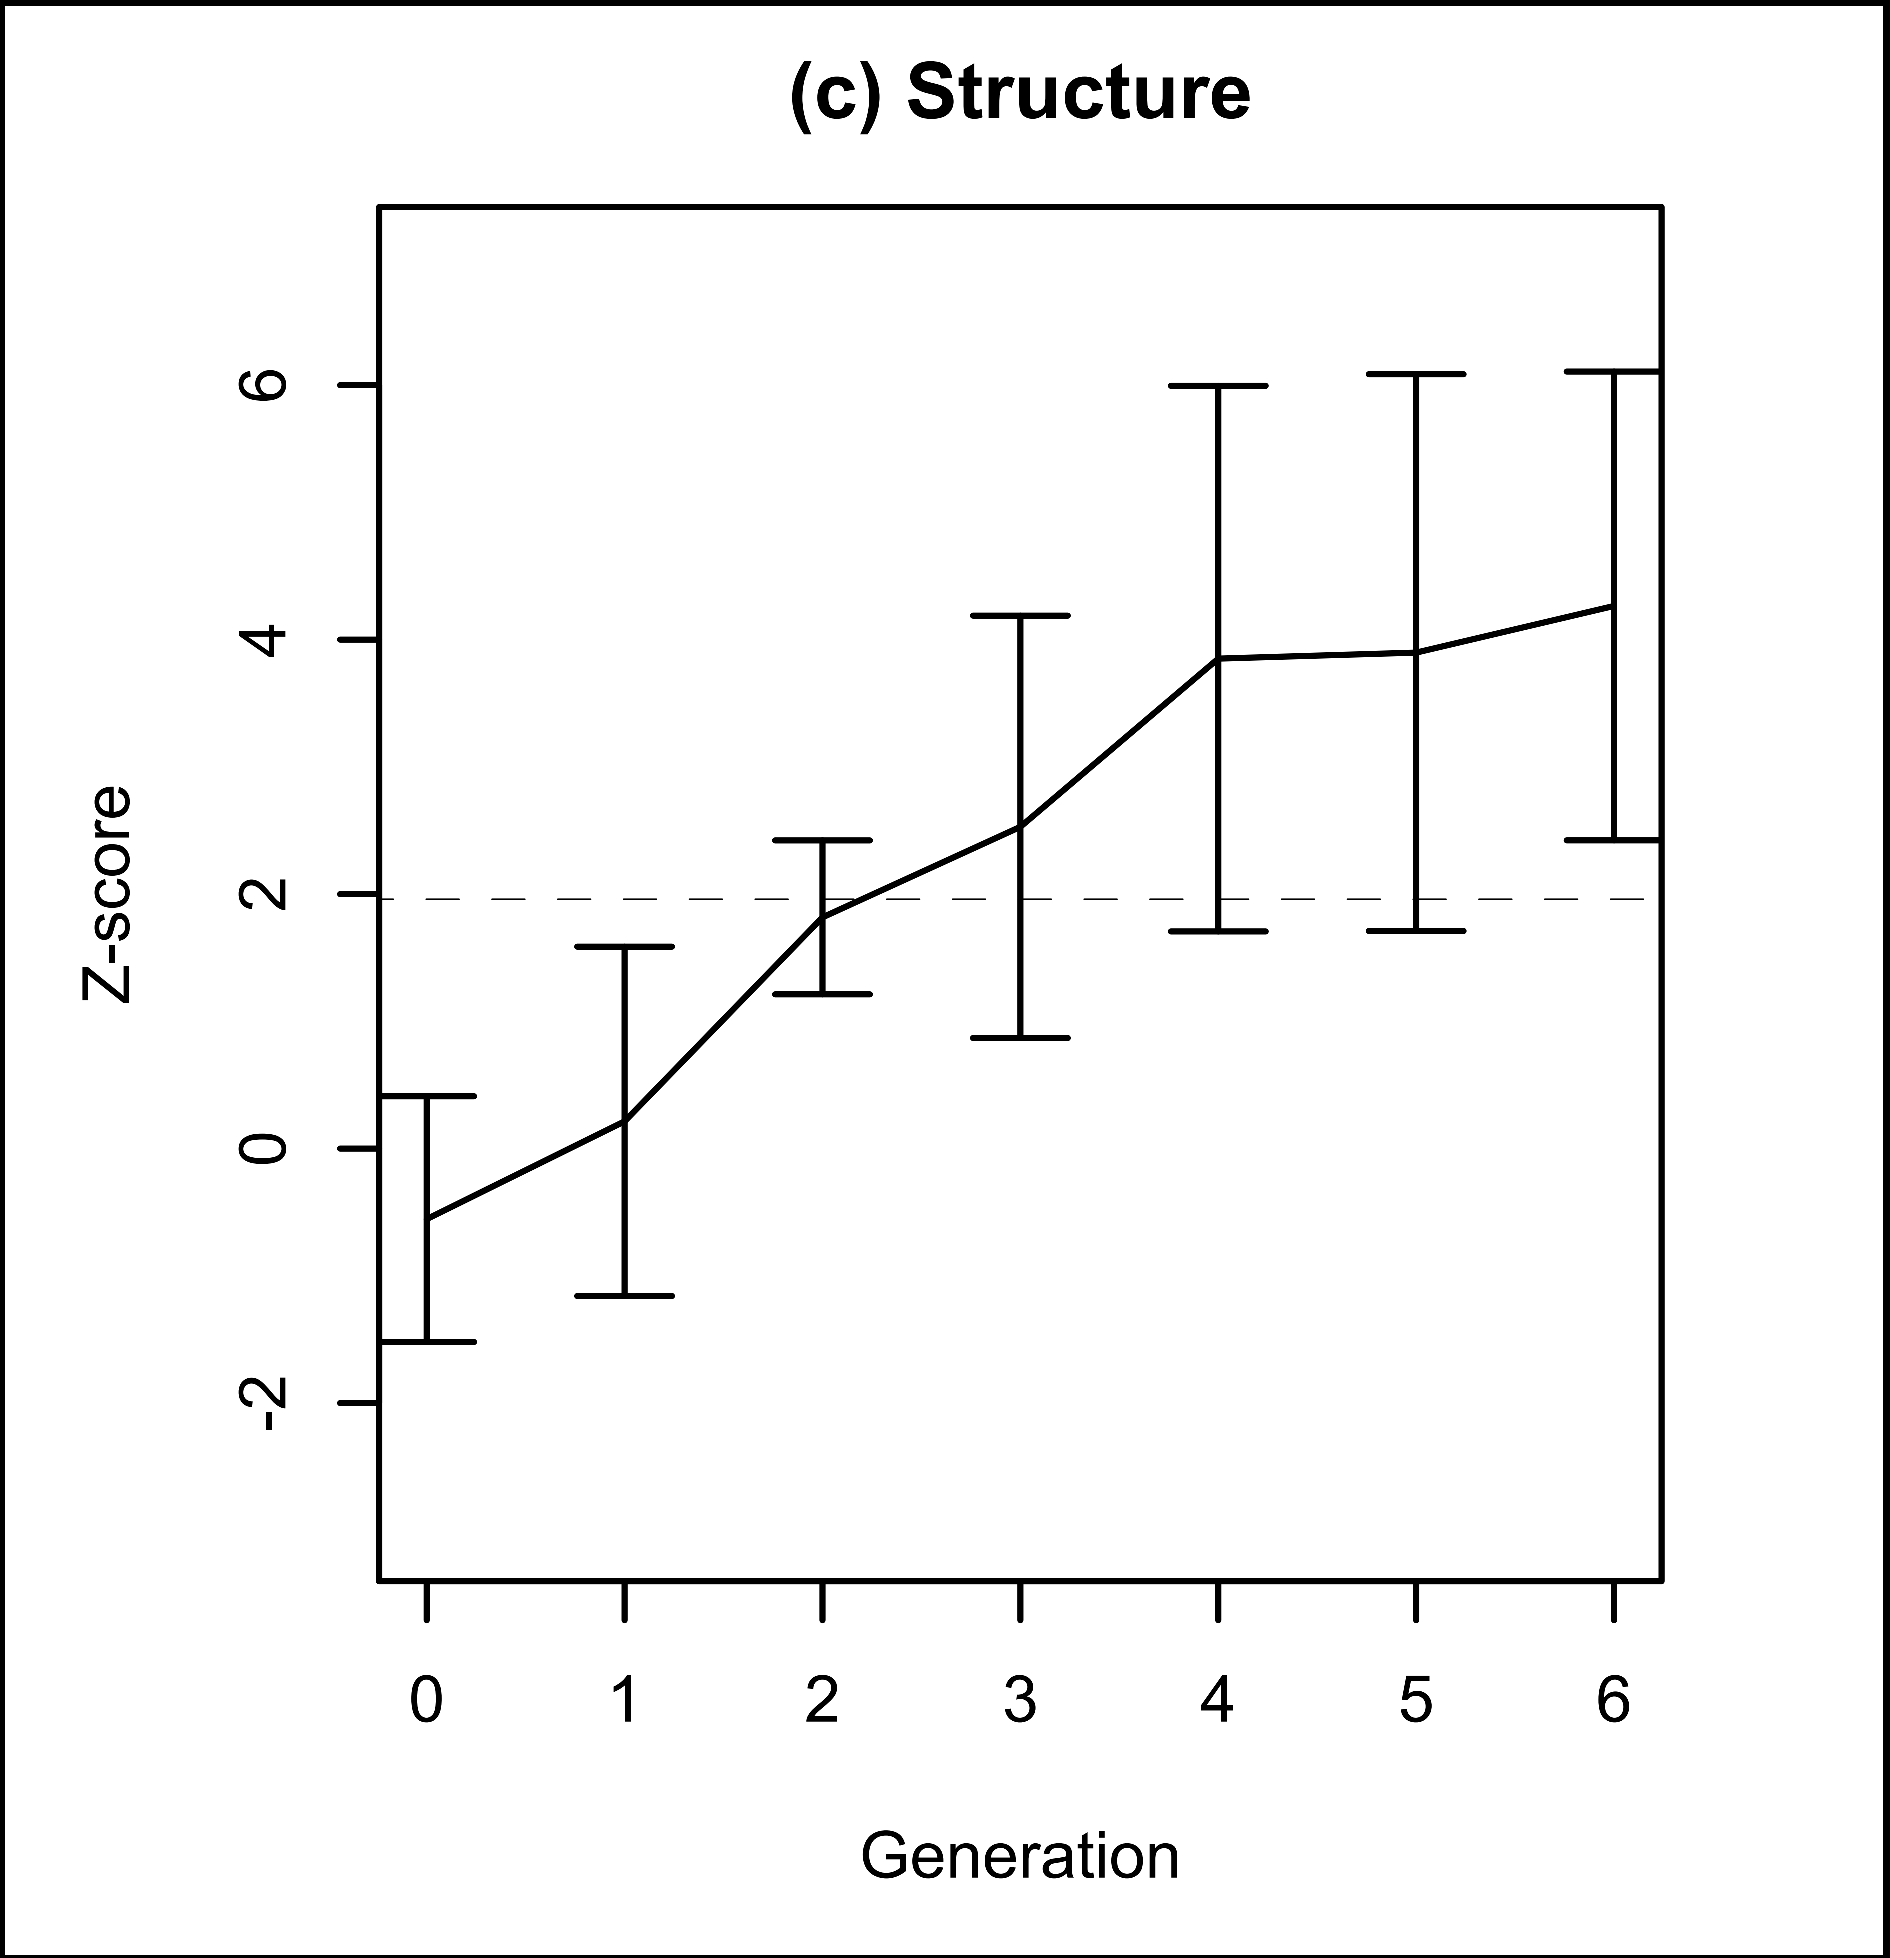

(c) Structure
Z-score
Generation
-2
0
2
4
6
0
1
2
3
4
5
6

WHY DO WE GET THIS RESULT?

Pressure from learning/transmission →

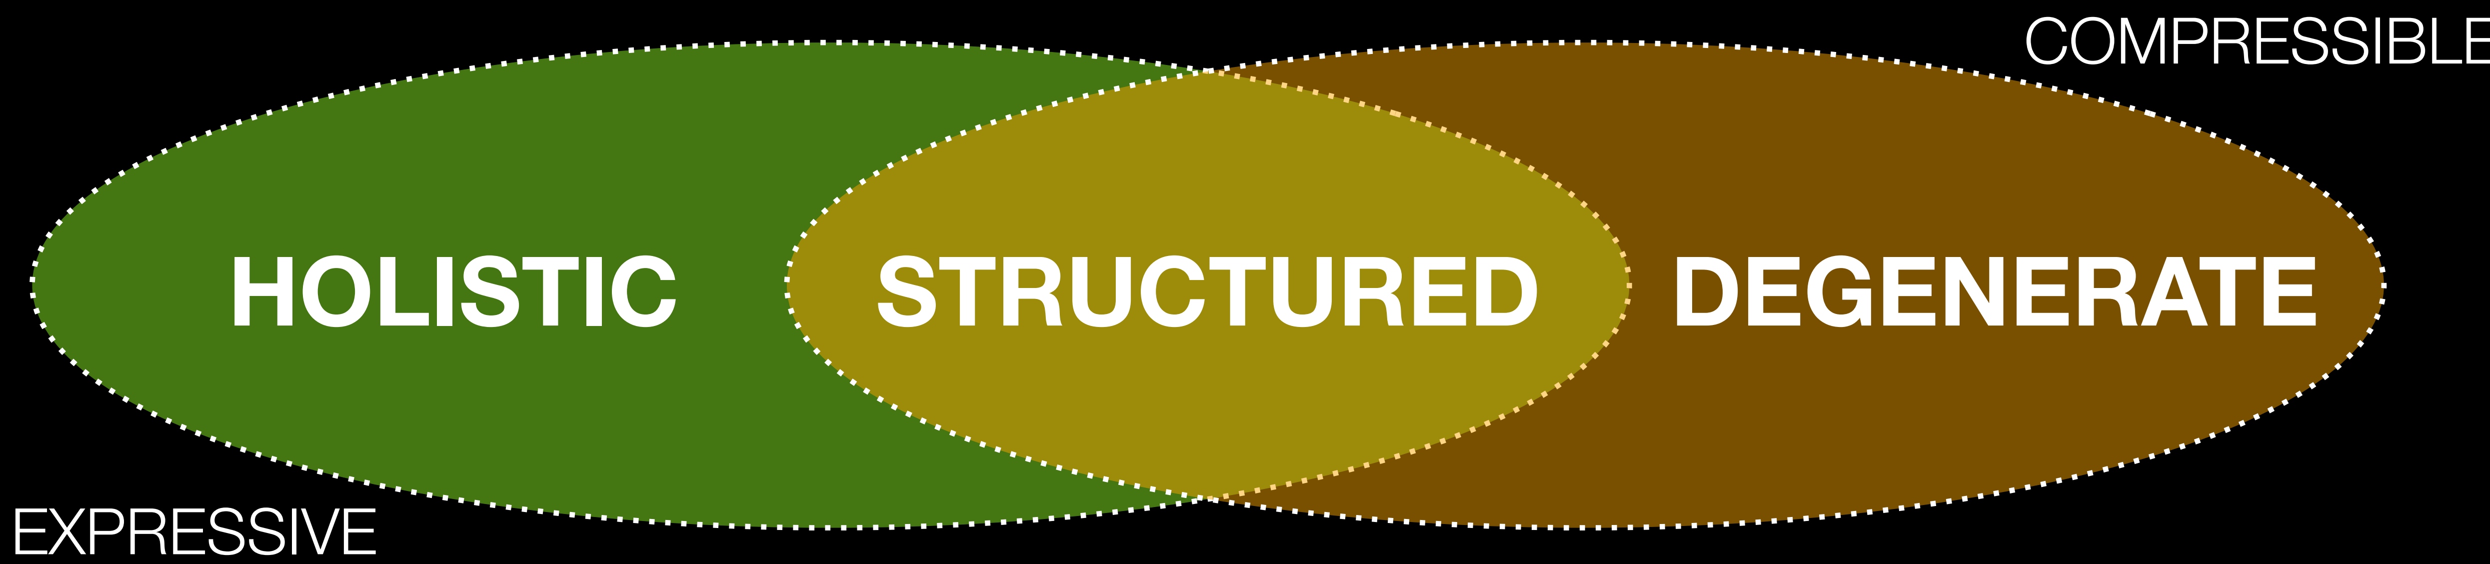

← Pressure from communication

REMOVING TRANSMISSION TO NAIVE LEARNERS

- We can test the hypothesis by changing the population structure
- Simply use the same pair of participants repeatedly
- No turnover of population - i.e. no transmission to naive learners

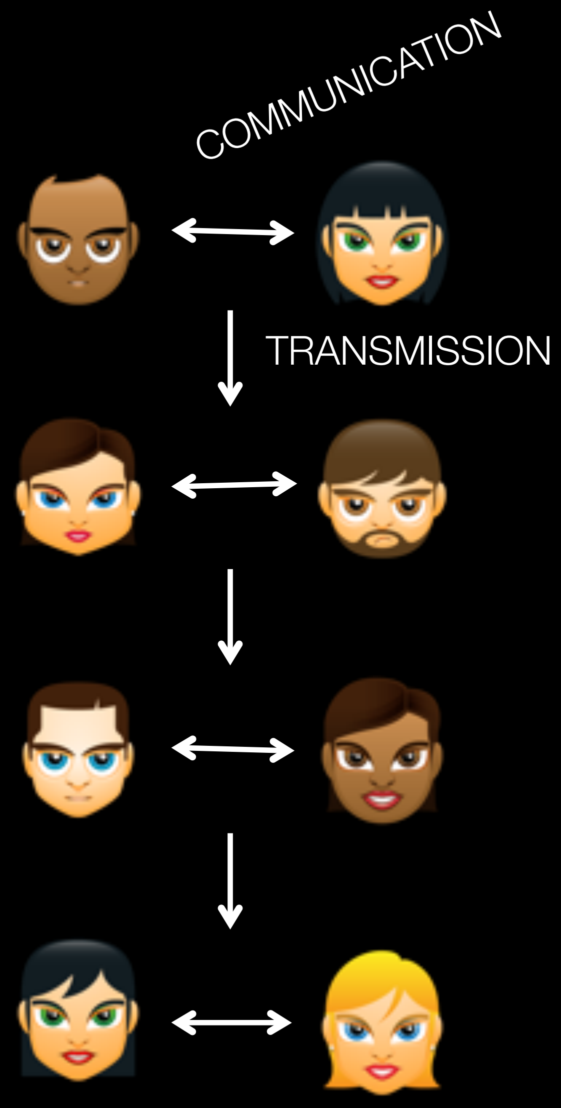

new dyad interact at each "generation"

REMOVING TRANSMISSION TO NAIVE LEARNERS

- We can test the hypothesis by changing the population structure
- Simply use the same pair of participants repeatedly
- No turnover of population - i.e. no transmission to naive learners
- Should shift the compressibility/ expressivity balance

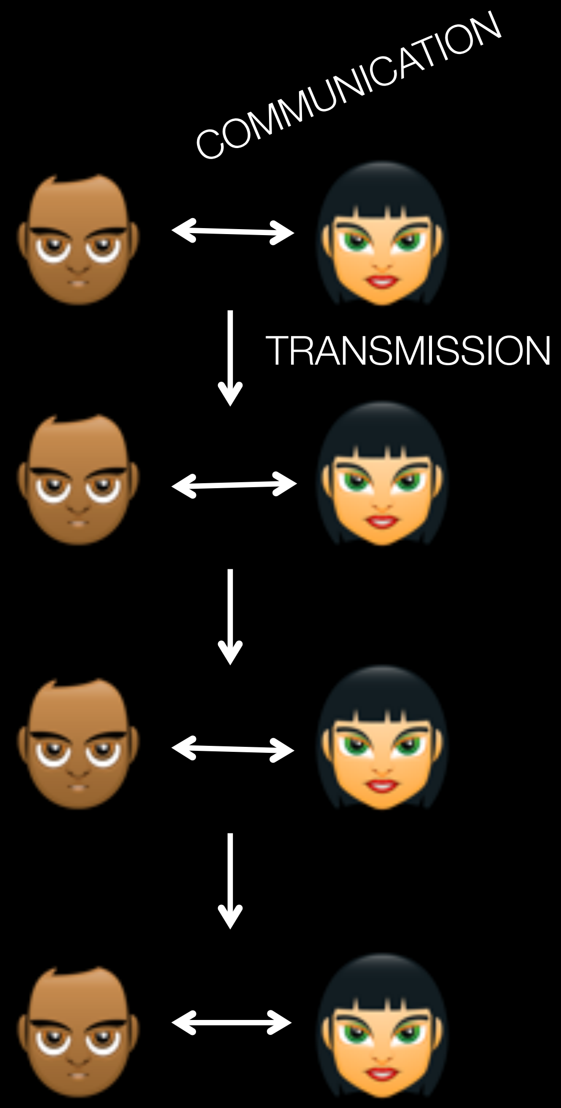

same participants interact repeatedly

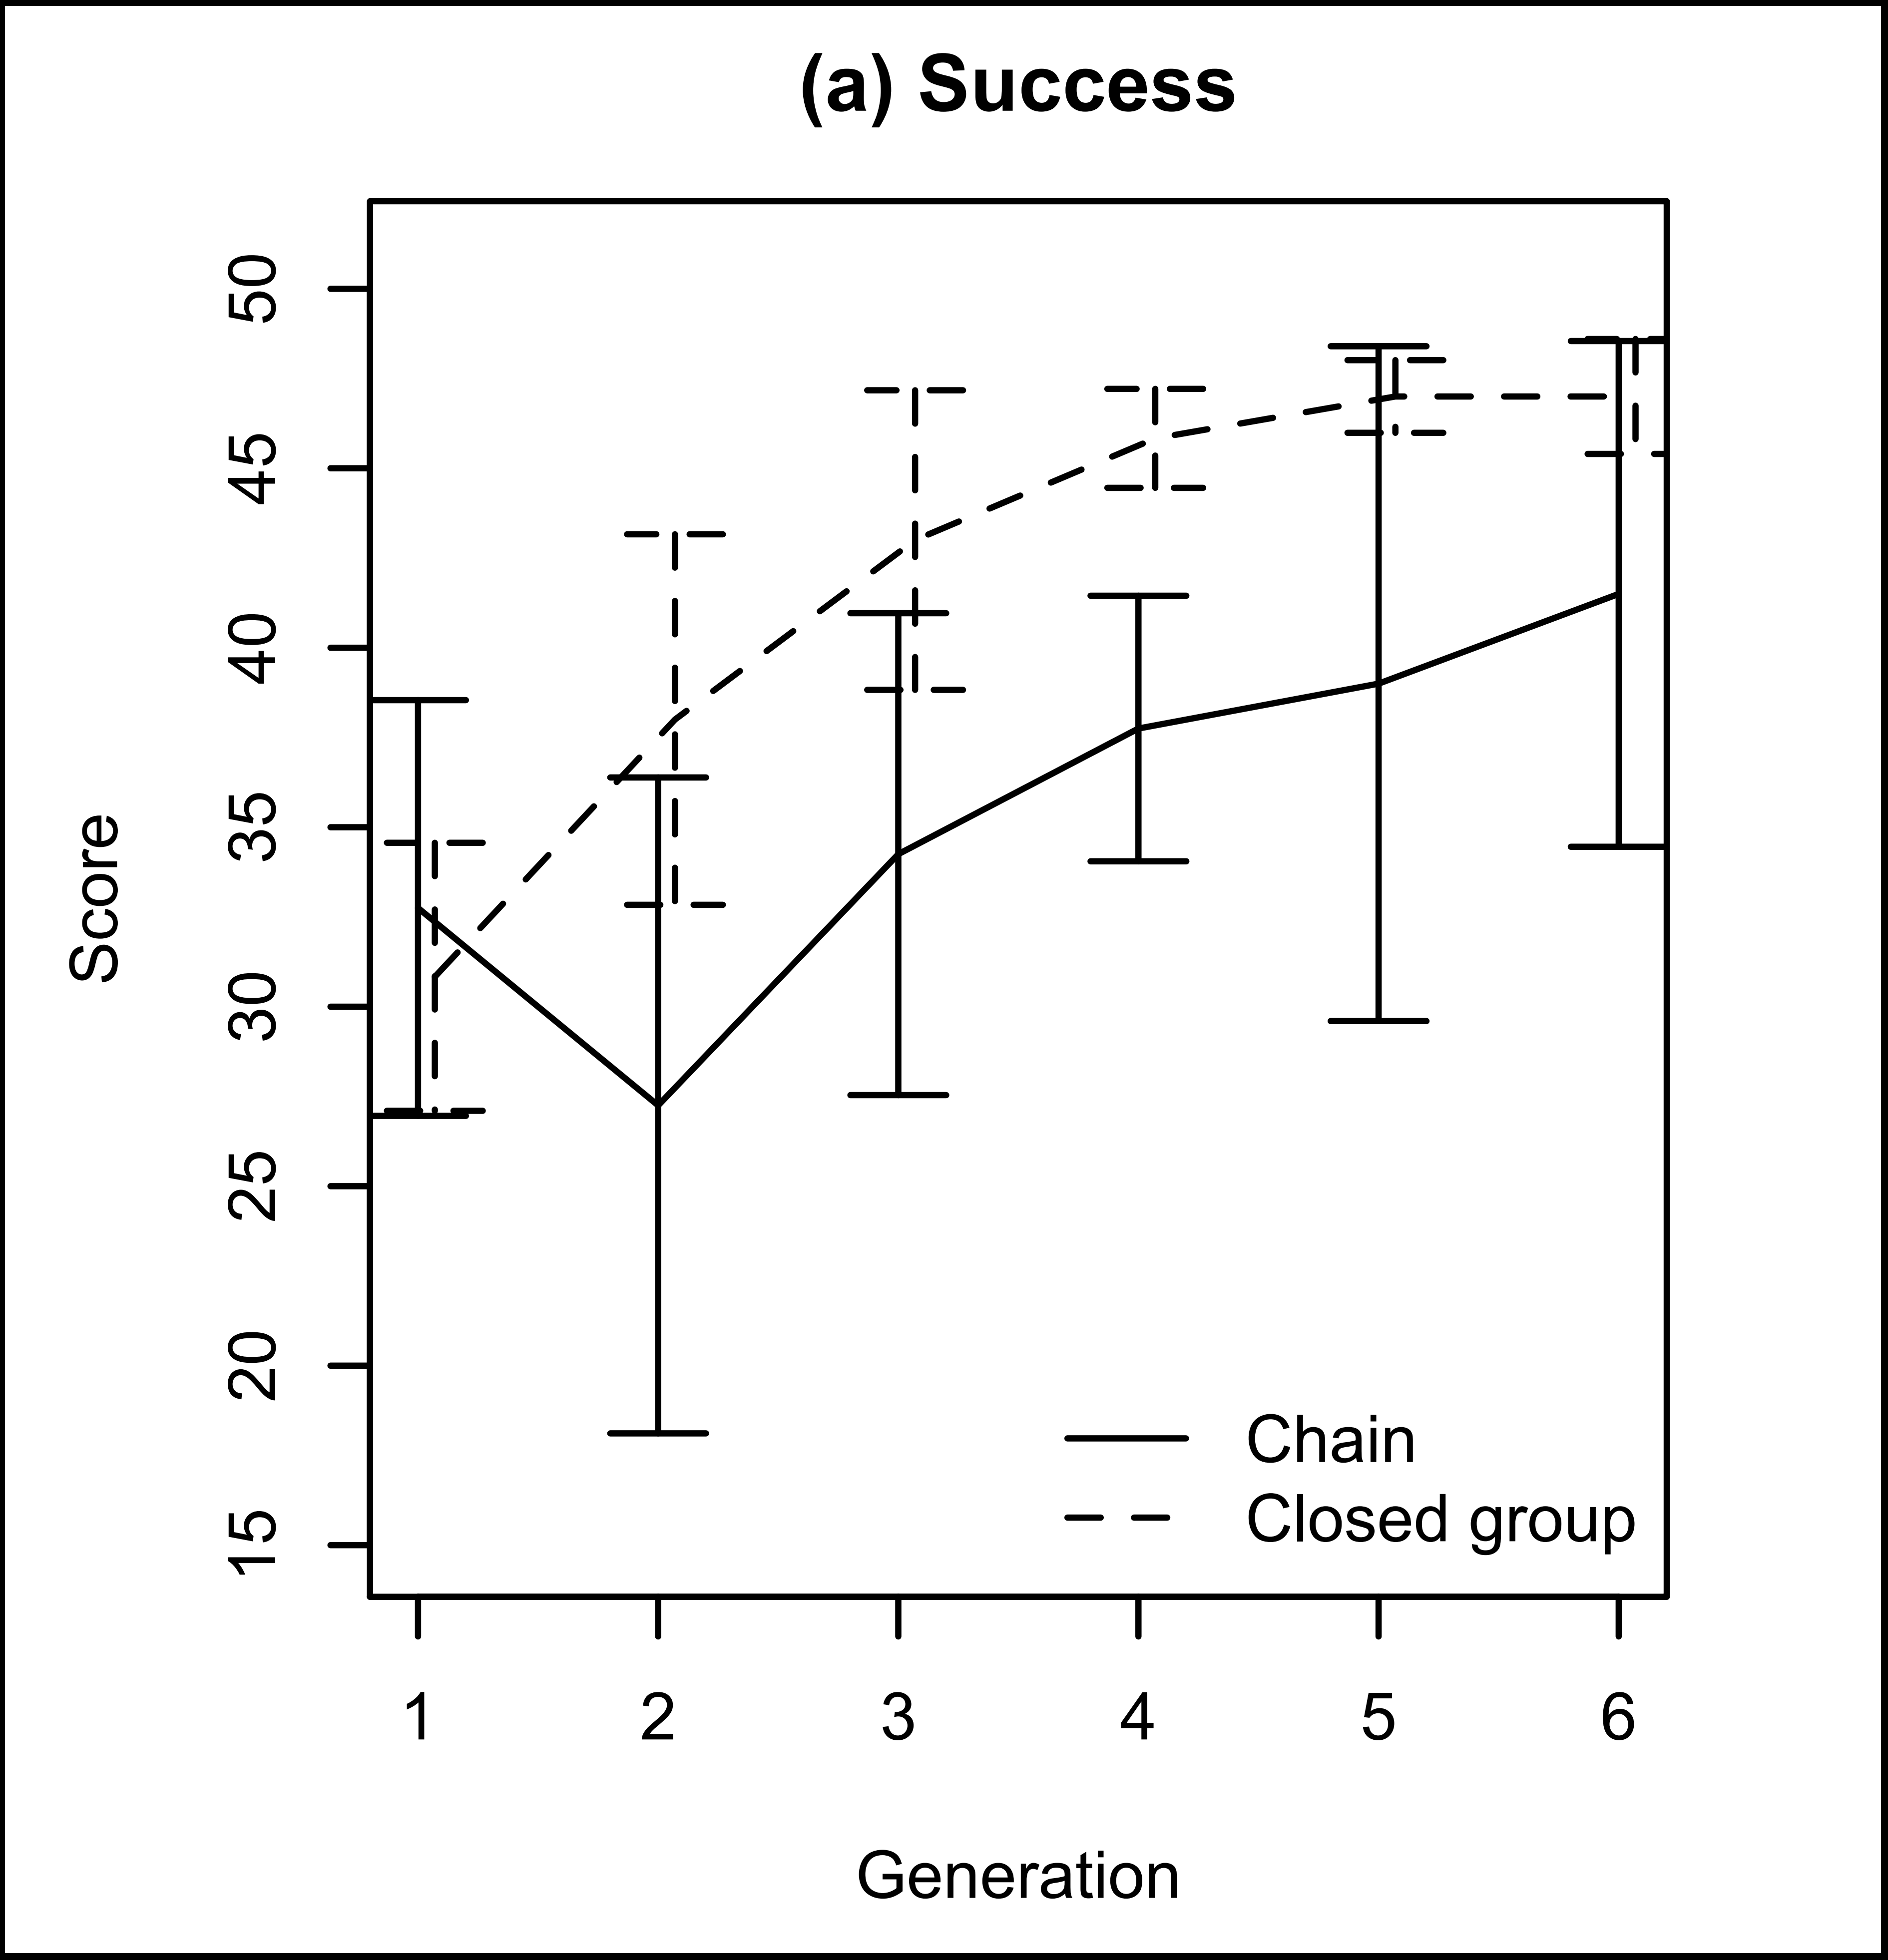
(a) Success
Score
50
45
40
35
30
25
20
15
Generation
1
2
3
4
5
6
Chain
Closed group

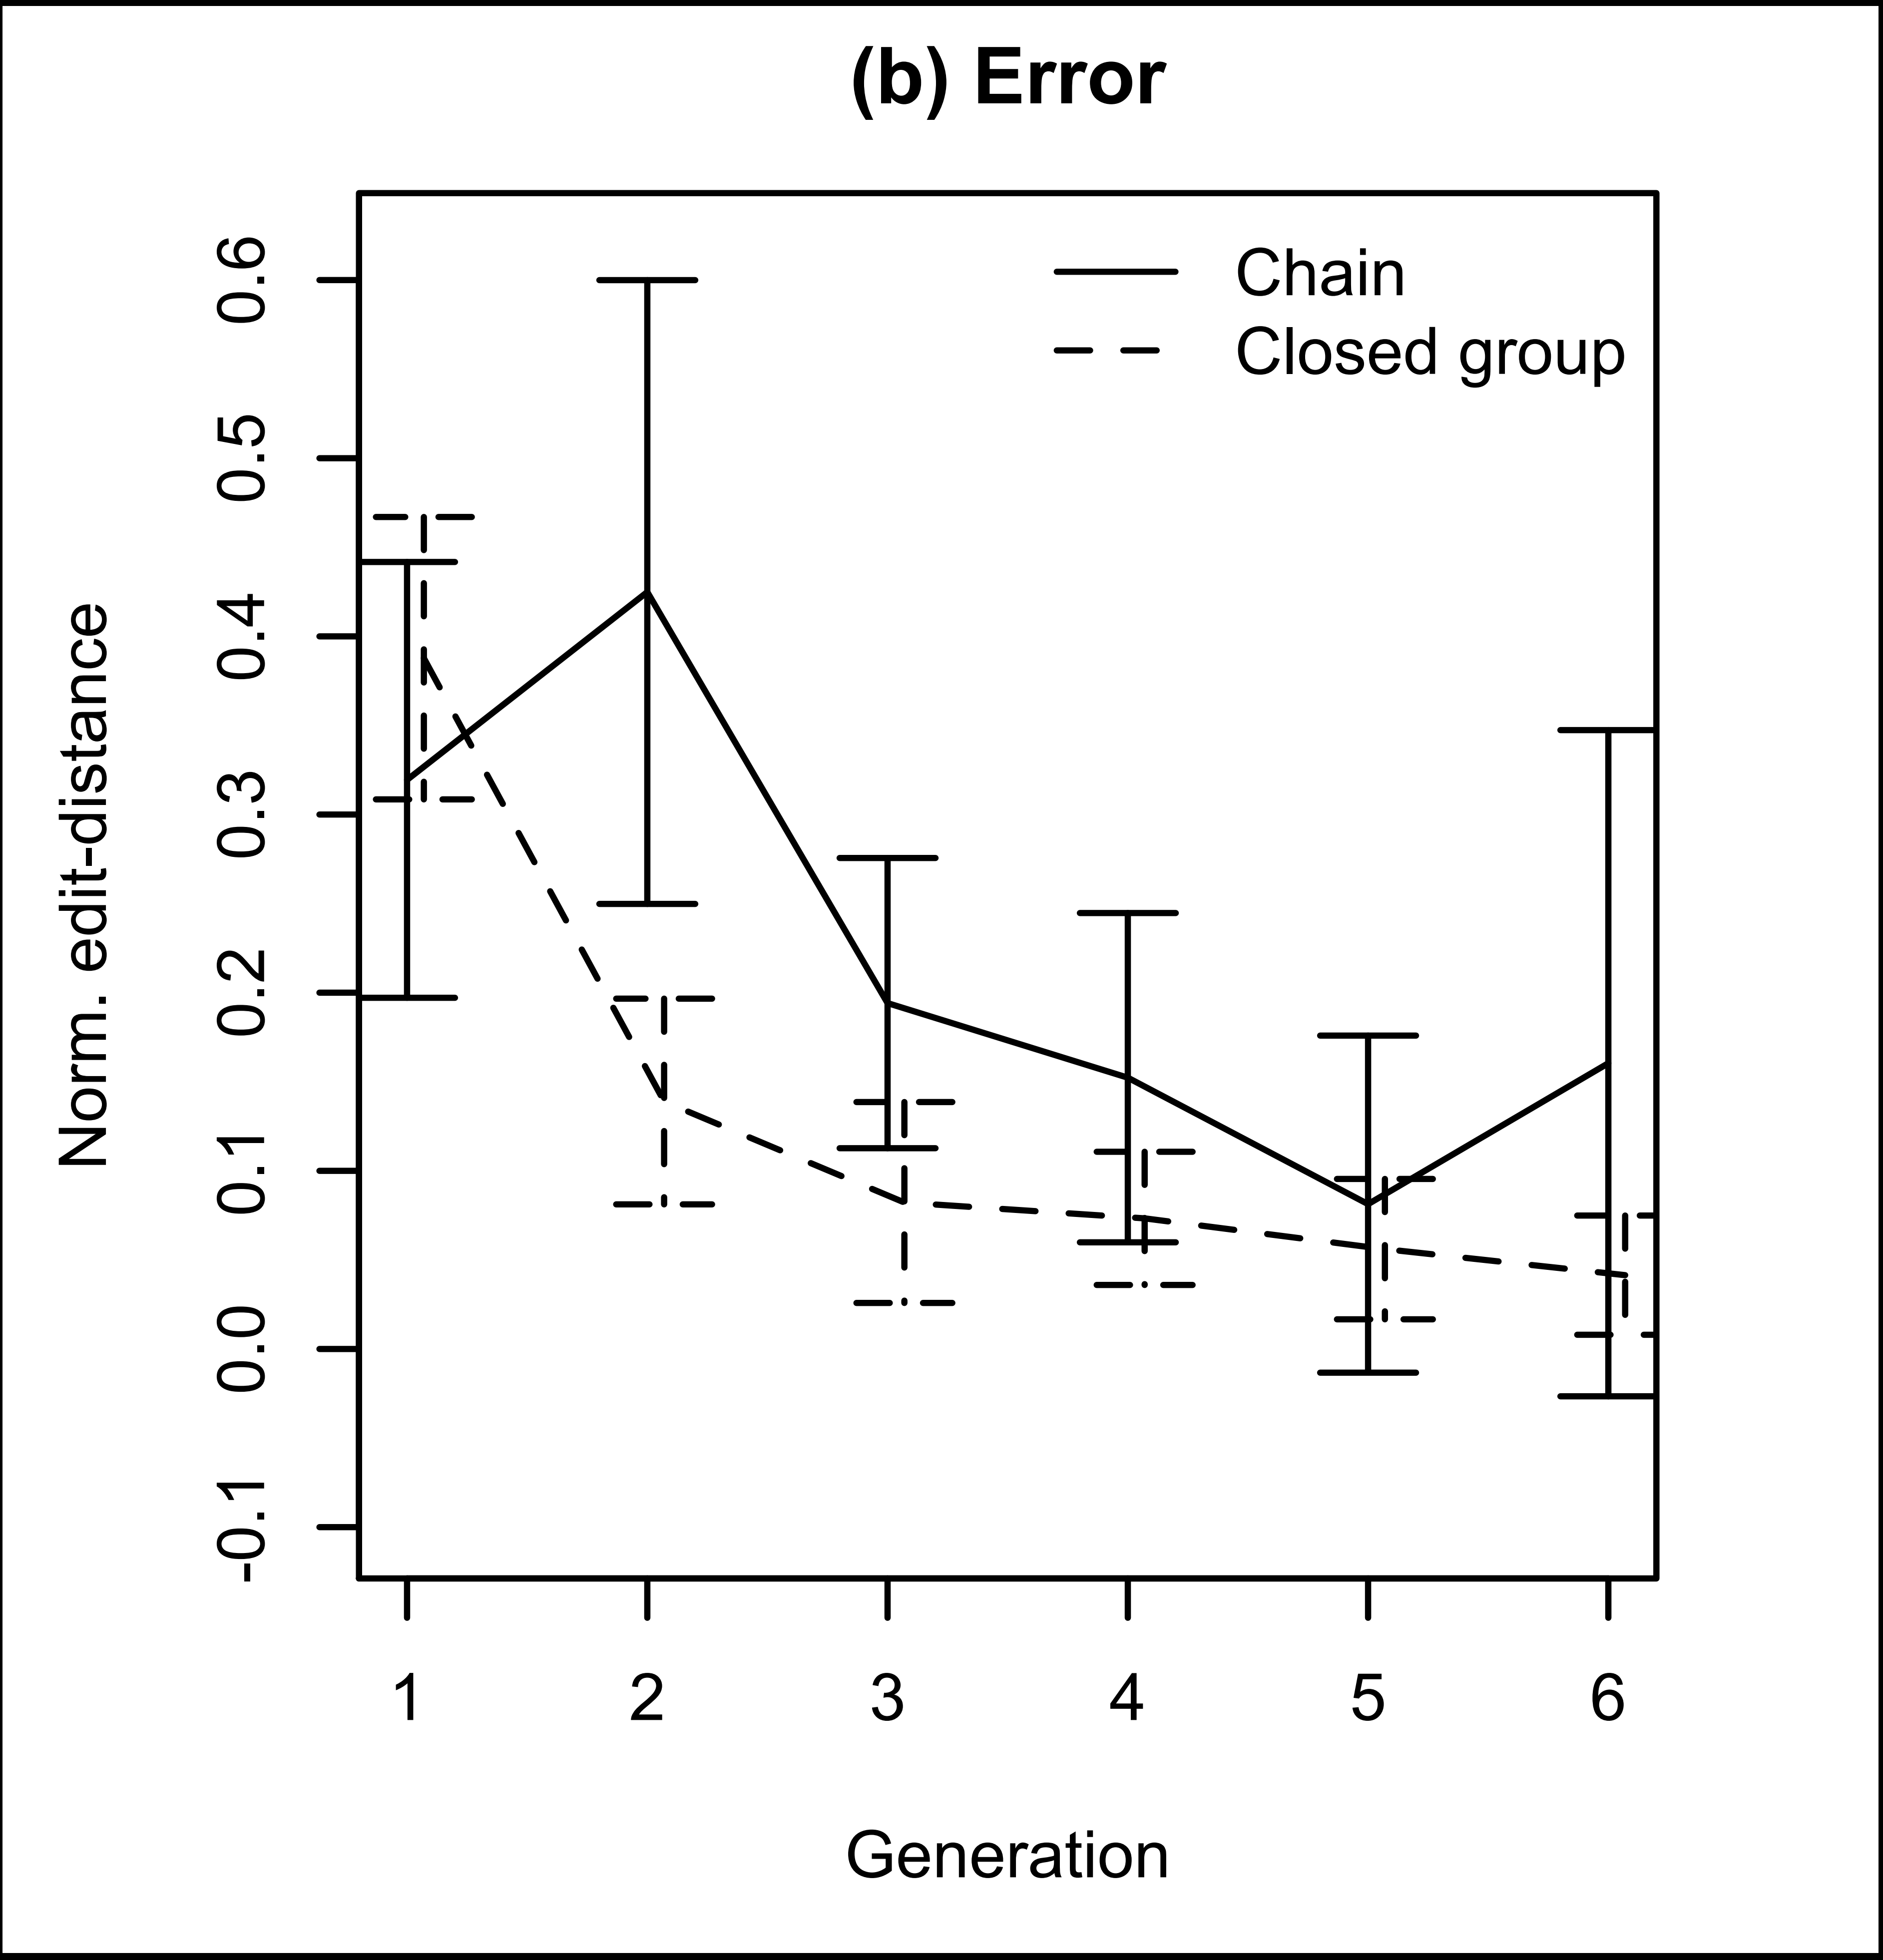
(b) Error
Chain
Closed group
Norm. edit-distance
0.6
0.5
0.4
0.3
0.2
0.1
0.0
-0.1
1
2
3
4
5
6
Generation

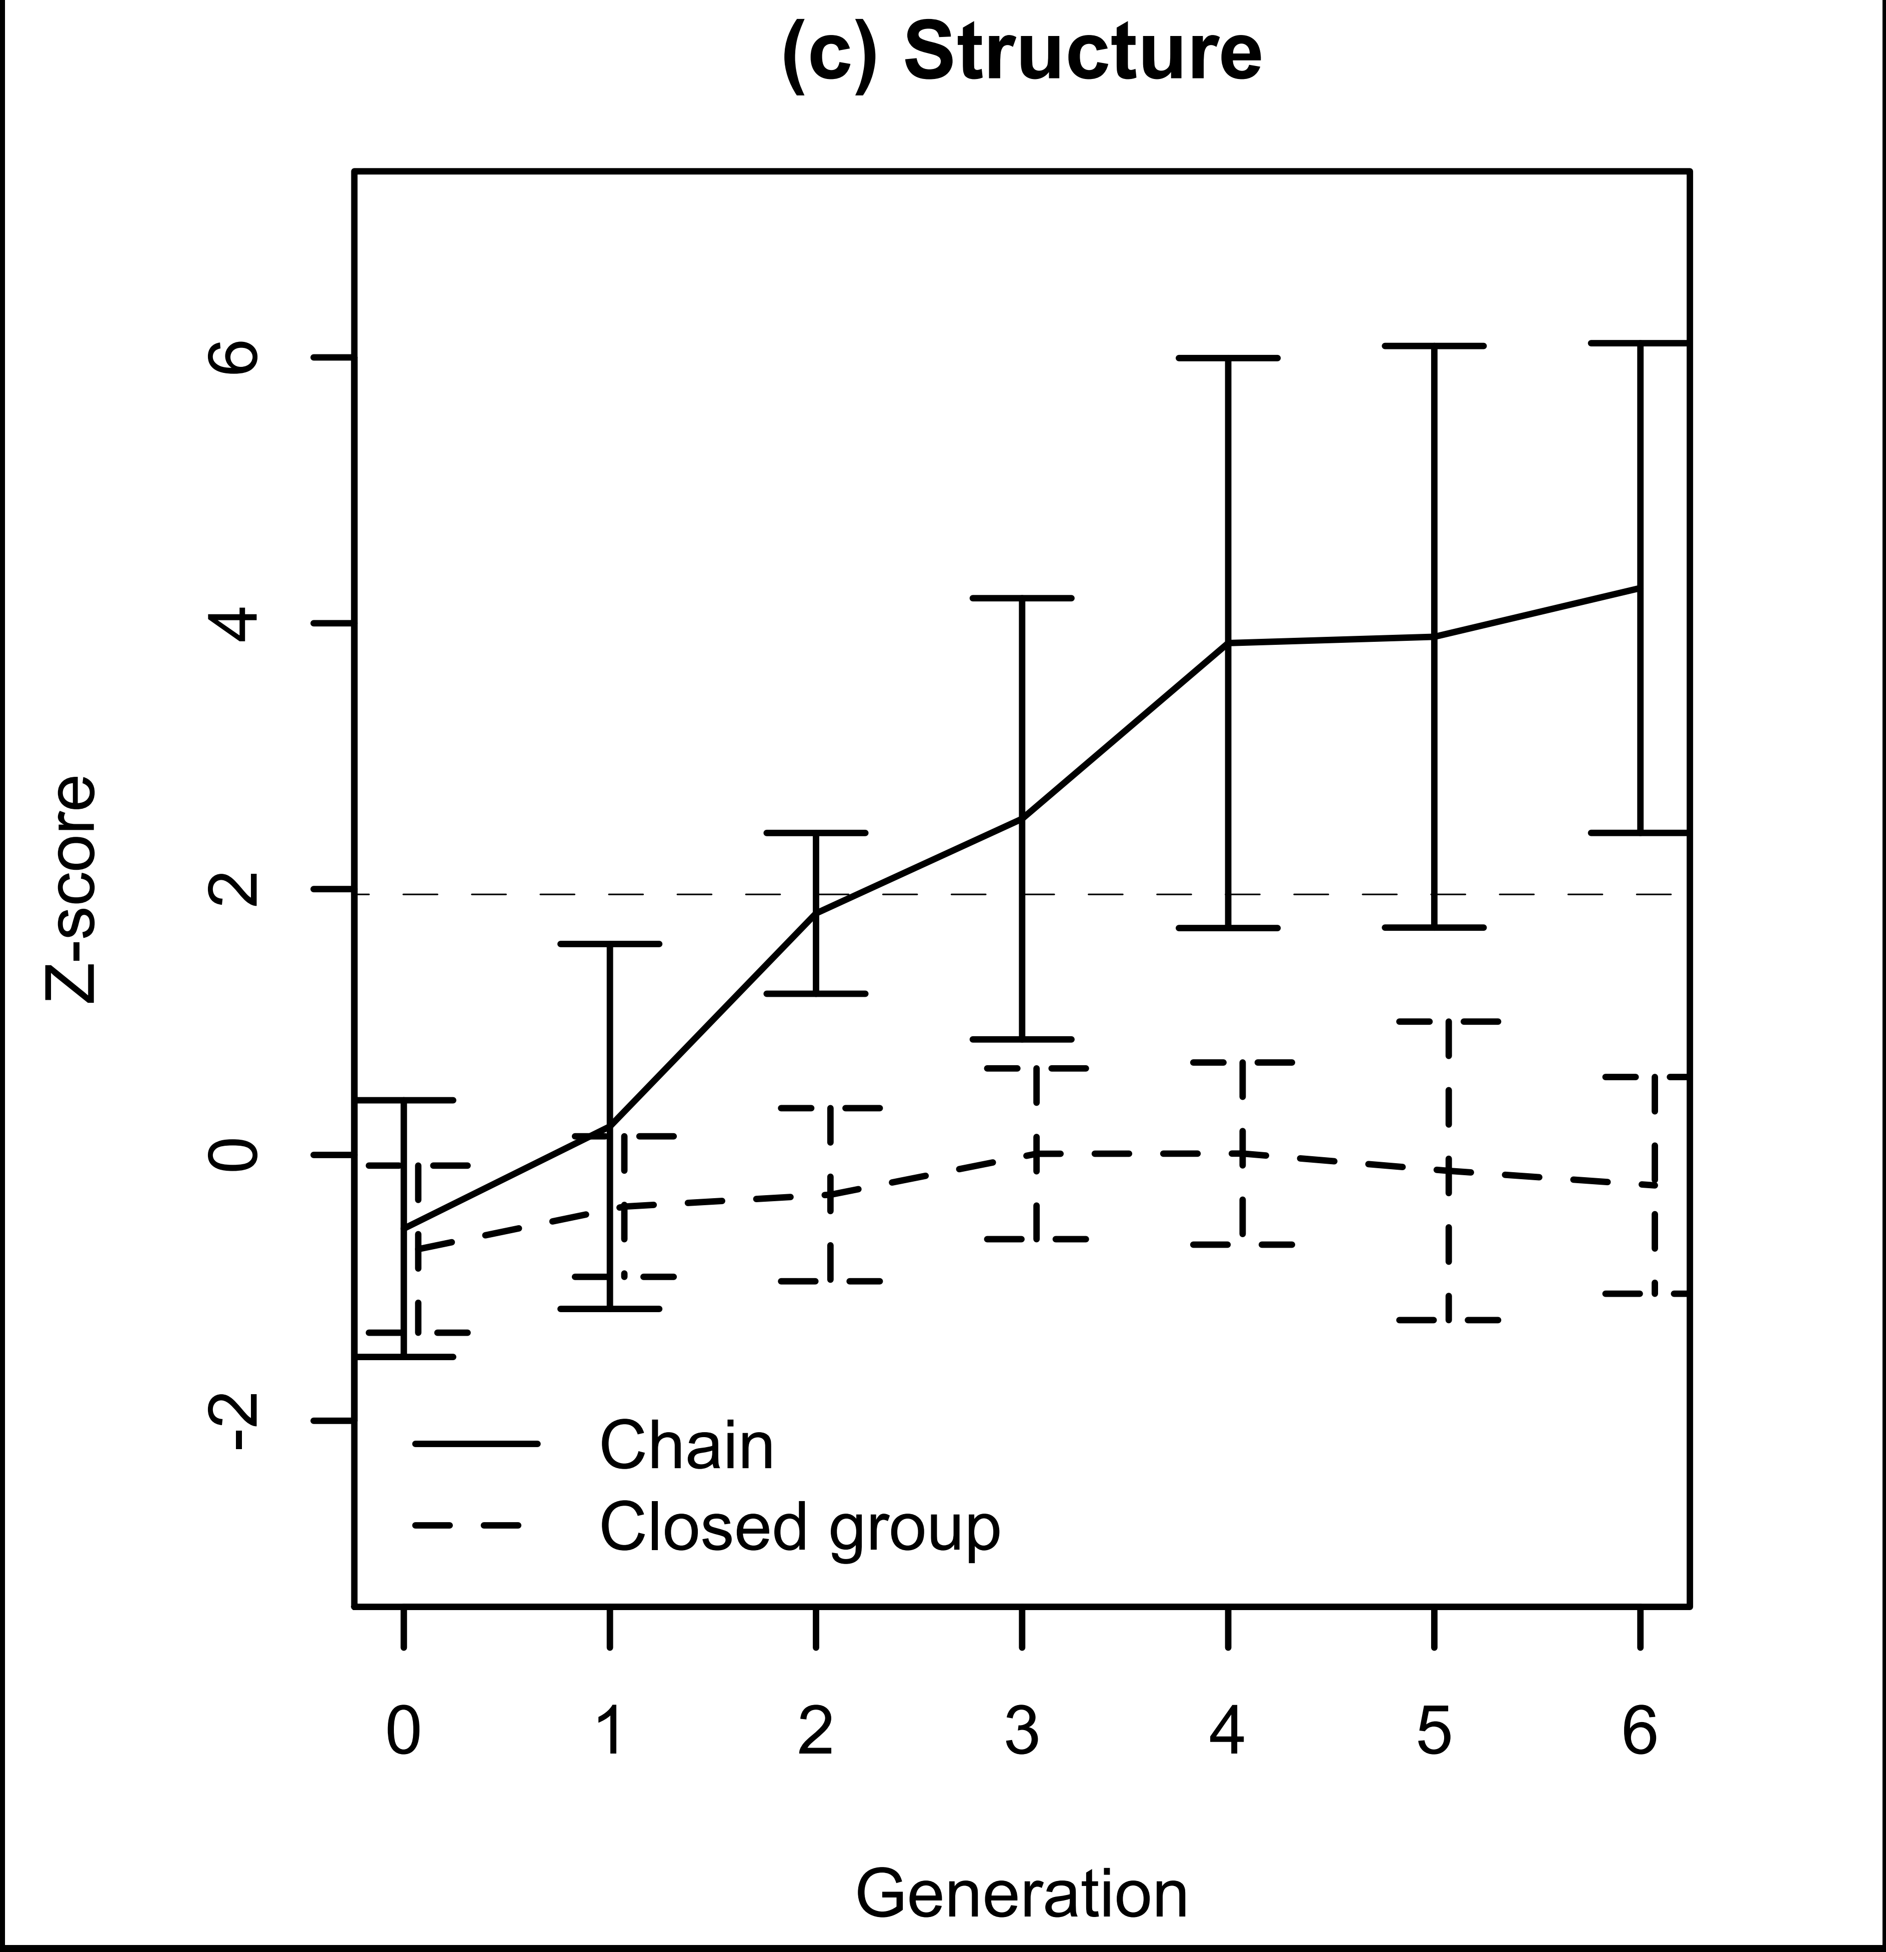

(c) Structure
Z-score
6
4
2
0
-2
0
1
2
3
4
5
6
Generation
Chain
Closed group

INITIAL RANDOM LANGUAGE

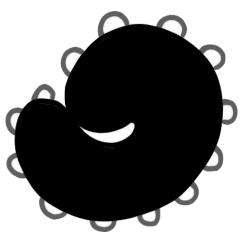 mokimu	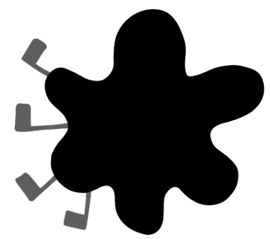 moko	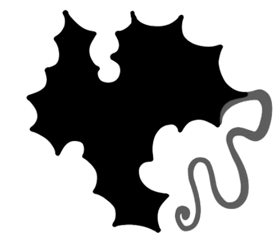 konu
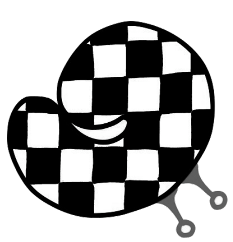 kimuwahu	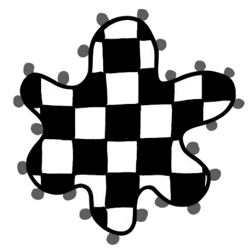 wahuhu	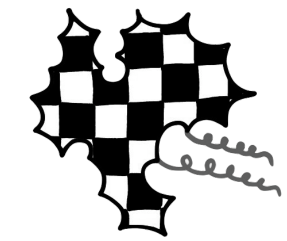 lawa
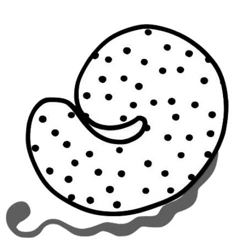 kinuki	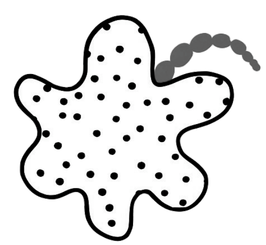 wekihu	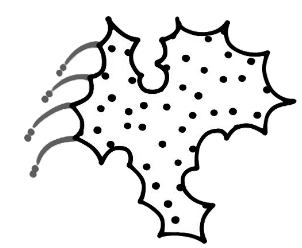 mohumu
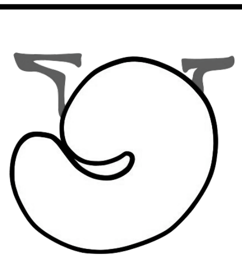 mukimuwa	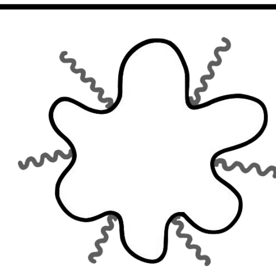 numu	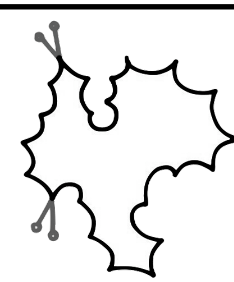 wakimu

AFTER SIX ROUNDS:
HOLISTIC LANGUAGE

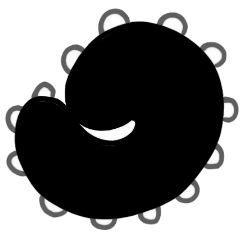 manunumoko	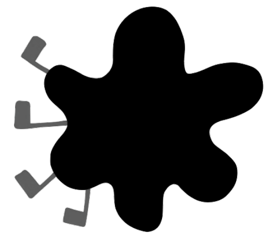 moko	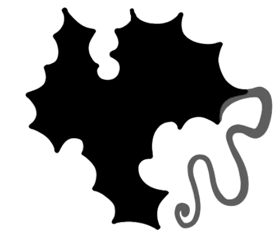 konu
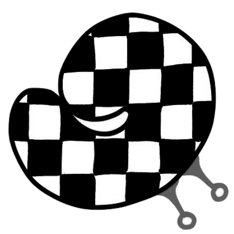 wekihumanunu	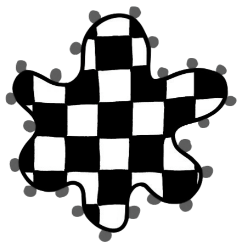 mokowekihu	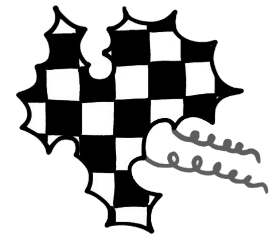 lawa
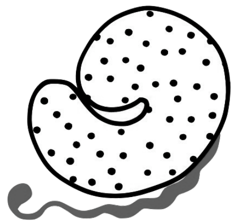 makihu	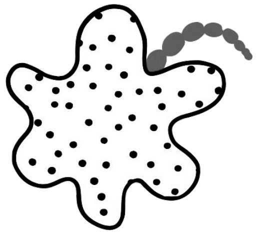 mahiku	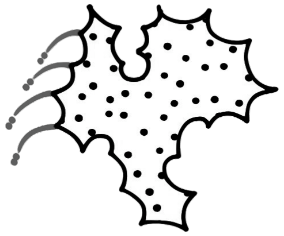 wekihulawa
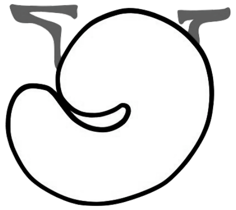 manunumonu	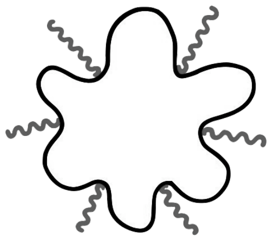 nomu	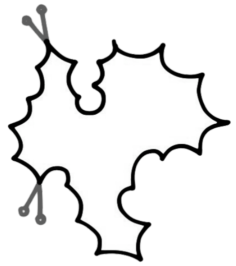 wekihu

SUMMARY OF EXPERIMENTS

- Transmission to naive learners → Compressibility pressure

  Communication → Expressivity pressure

  - [+TRANSMISSION, -COMMUNICATION] → **Degenerate** language
  - [+TRANSMISSION, +COMMUNICATION] → **Structured** language
  - [-TRANSMISSION, +COMMUNICATION] → **Holistic** language

- Adaptive structure derives directly from cultural evolution

WHAT'S REQUIRED FOR STRUCTURE?

A general pressure for simplicity/compressibility

A task that favours discrimination

A population structure that includes communicating with naive learners

Cultural evolution by iterated learning does the rest

PART 2:
FROM EXPERIMENTS TO SIMULATION

SIMULATING ITERATED LEARNING

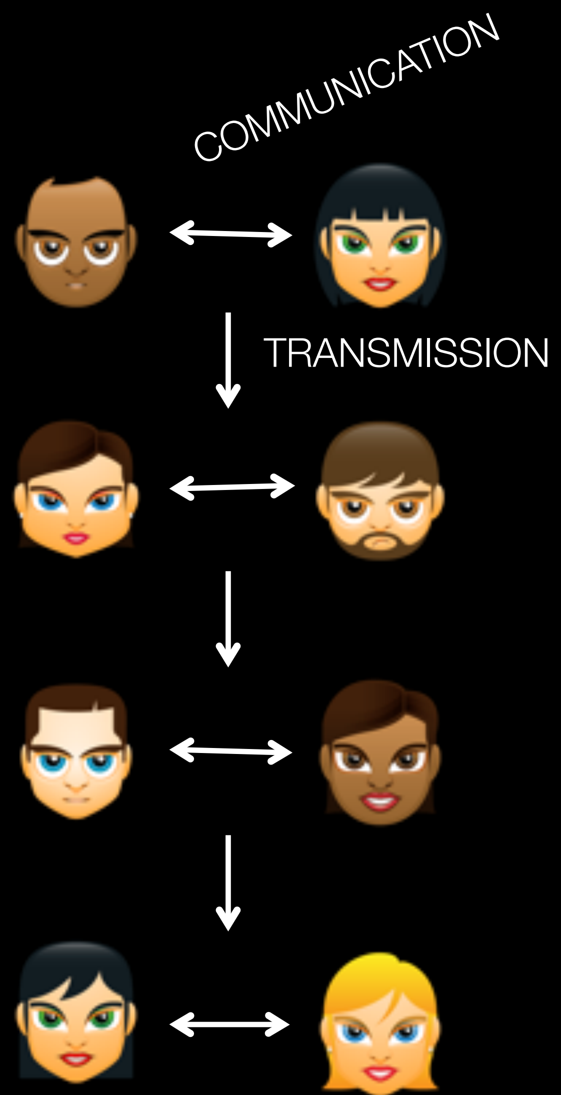

- Allows us to rule out influence of participants' native language
- Also allows us to explore:
    - Other aspects of learners biases (e.g. strength of preference for simplicity)
    - Other population structures (beyond chains and small closed groups)
    - Memory
    - …

SIMULATING ITERATED LEARNING

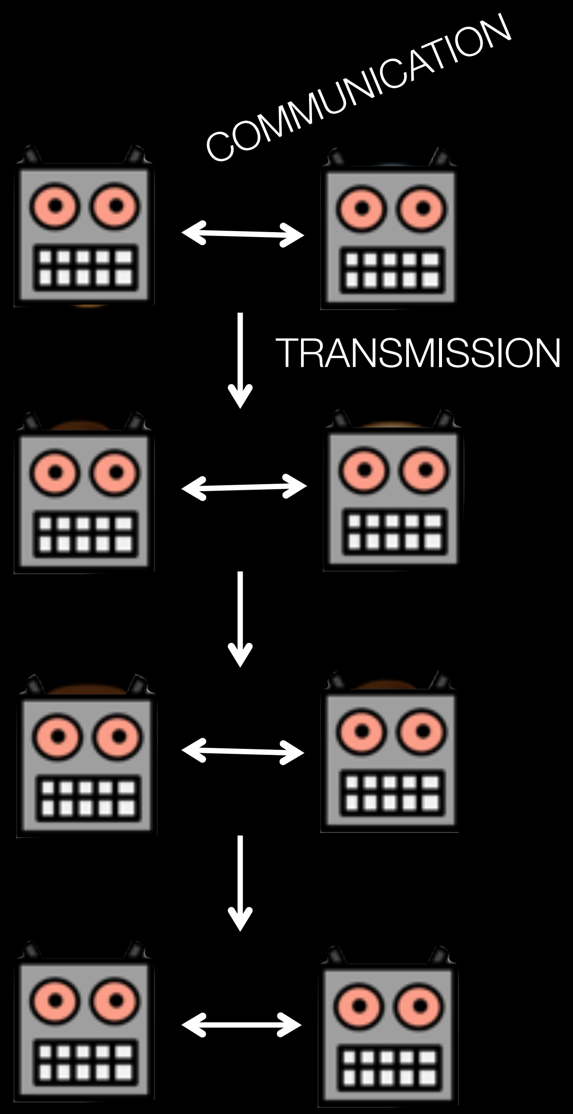

- Allows us to rule out influence of participants' native language
- Also allows us to explore:
  - Other aspects of learners biases (e.g. strength of preference for simplicity)
  - Other population structures (beyond chains and small closed groups)
  - Memory
  - …

LEARNING AS BAYESIAN INFERENCE

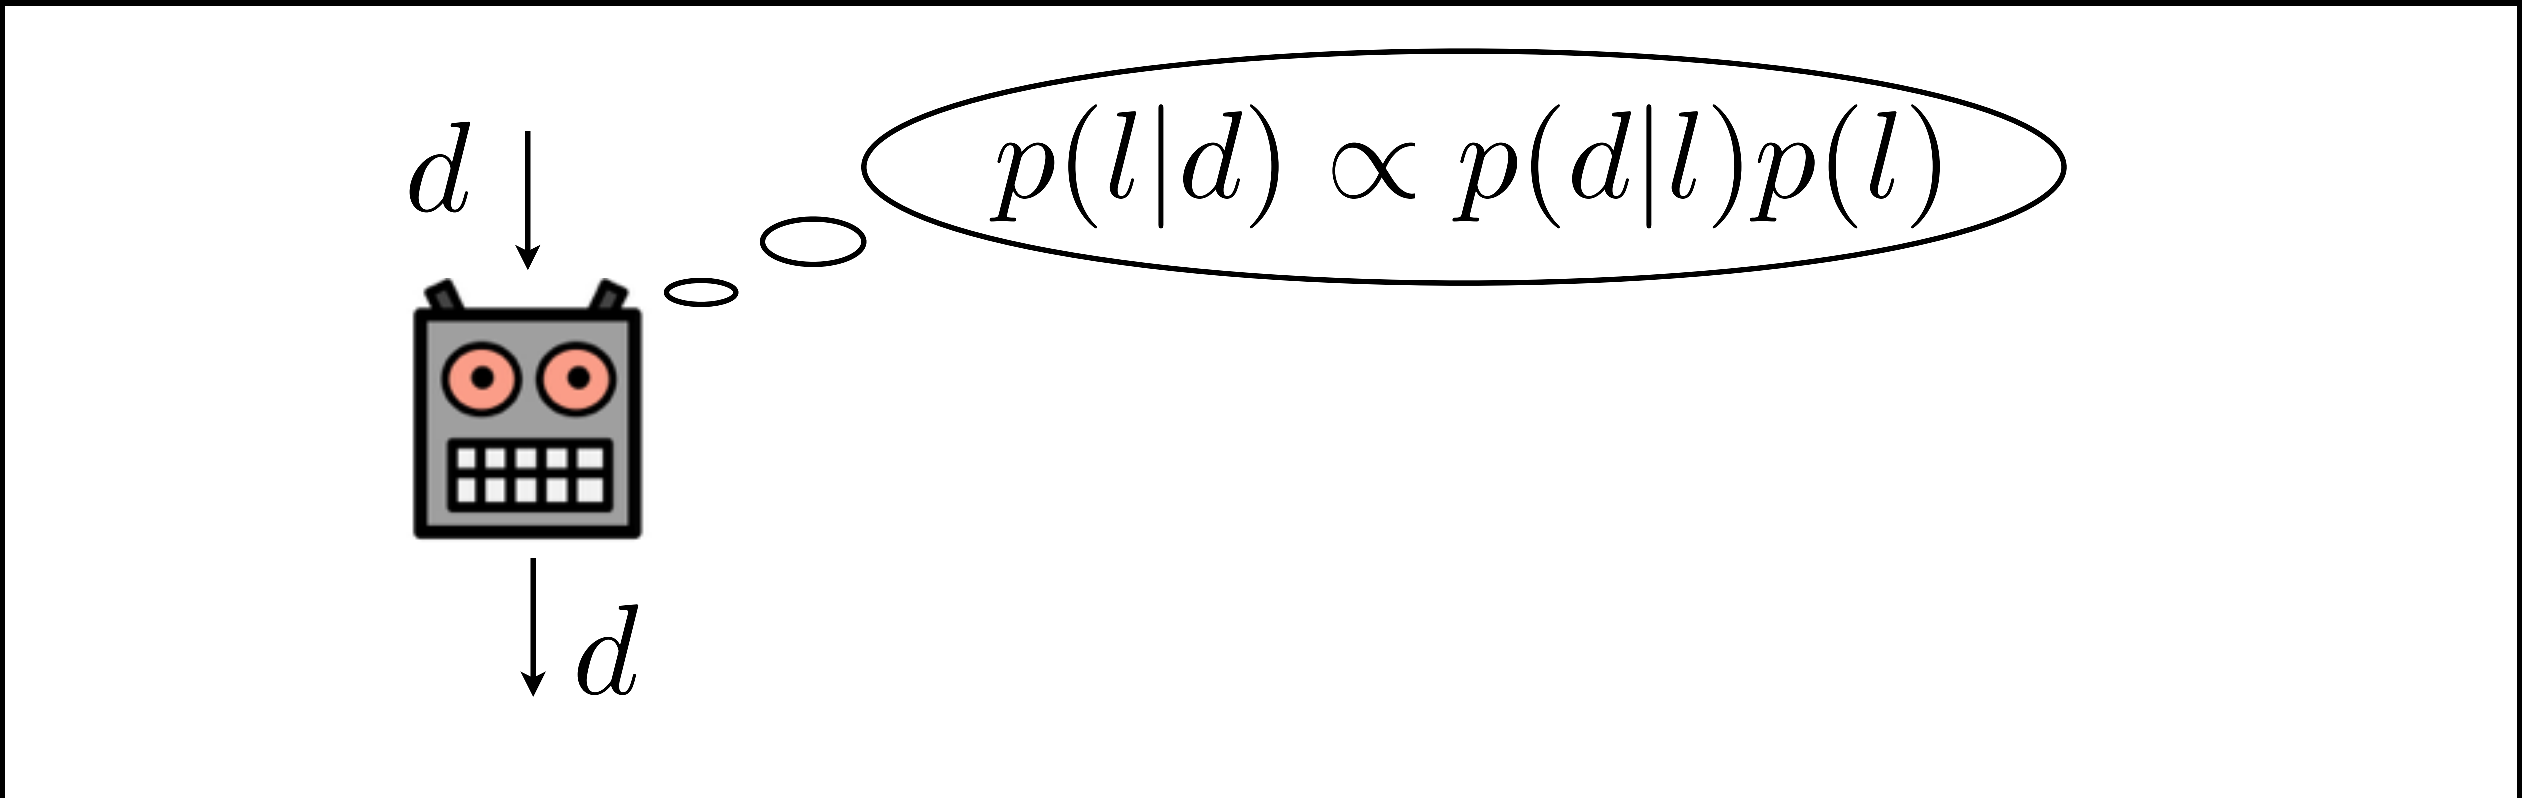

Data: Meaning-signal pairs

Languages: specify signal for each meaning

Likelihood p(d|l): avoid ambiguous utterances

Prior p(l): prefer compressible languages (shorter coding length)

LANGUAGES

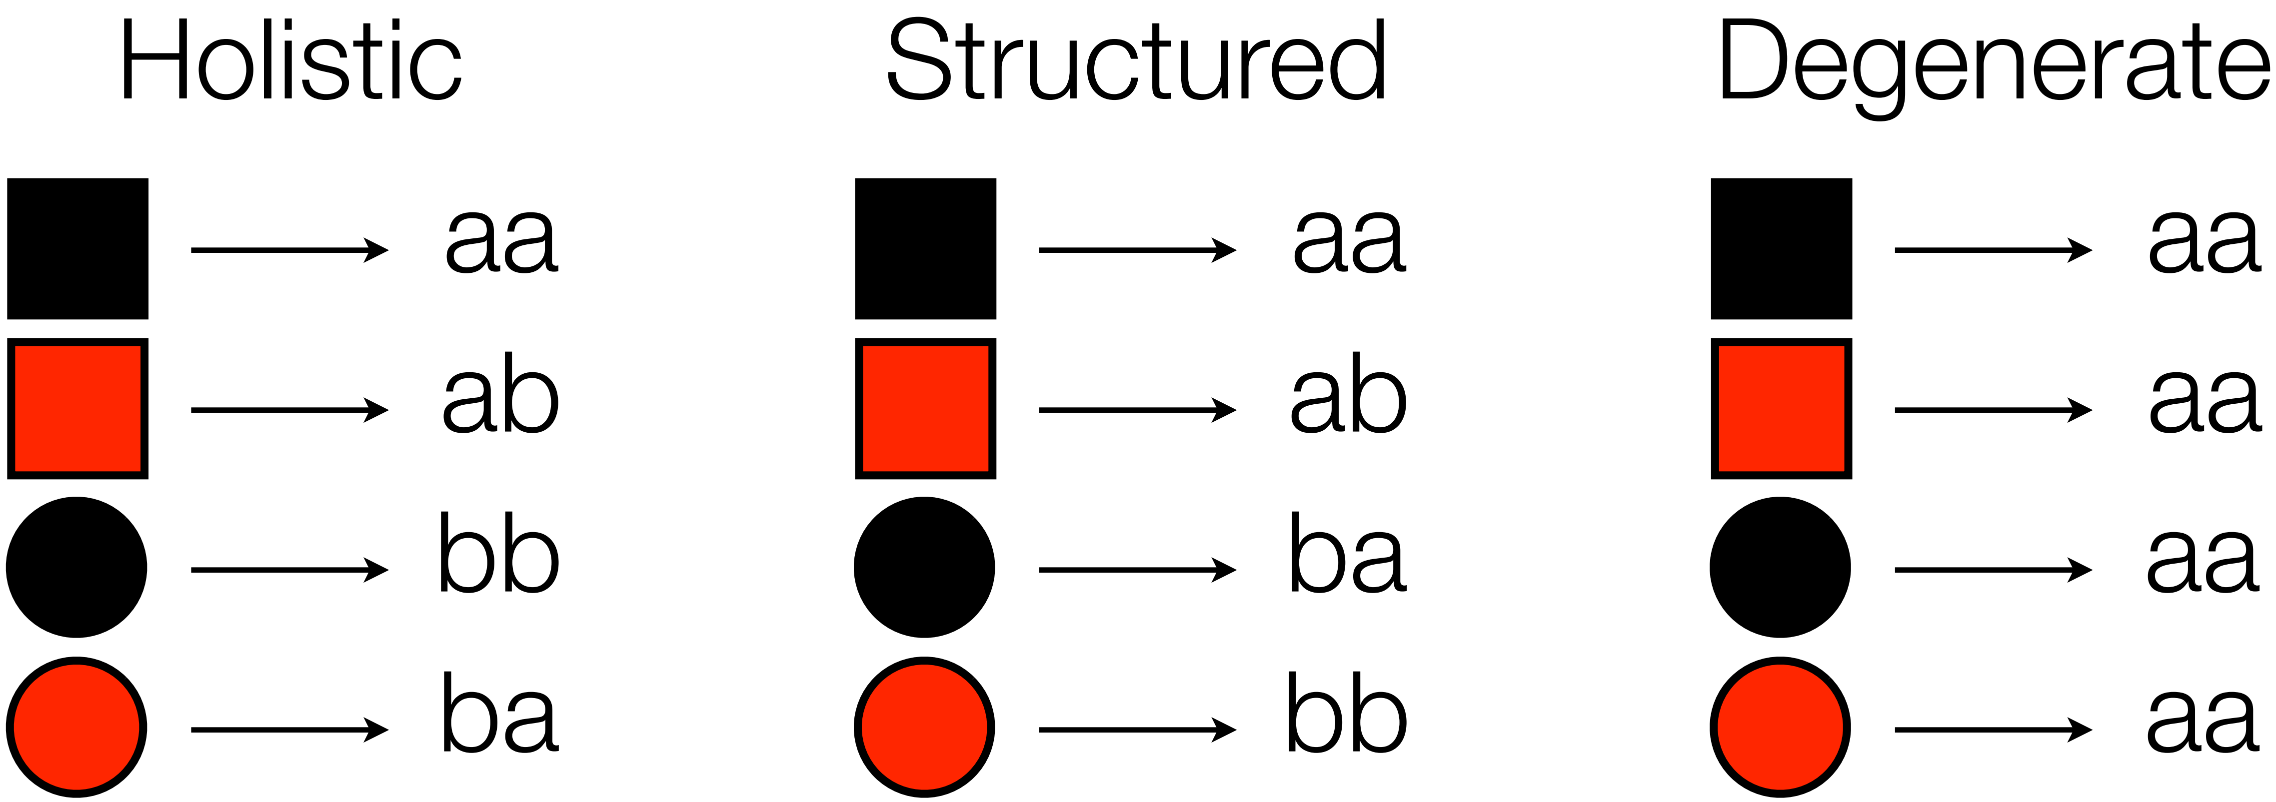
Holistic
aa
ab
bb
ba
Structured
aa
ab
ba
bb
Degenerate
aa
aa
aa
aa

LIKELIHOOD

$$p(signal|meaning, l)$$

My language

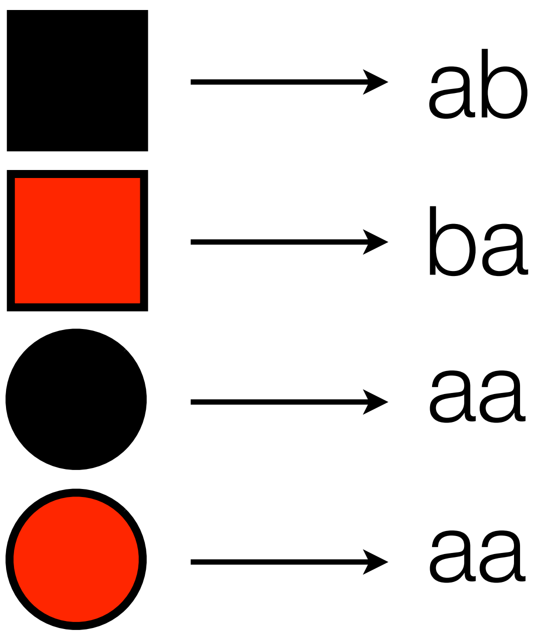

$$p(aa|\bigcirc, \text{my language}) \propto 1 - \epsilon$$

$$p(ba|\bigcirc, \text{my language}) \propto \epsilon/3$$

$$p(bb|\bigcirc, \text{my language}) \propto \epsilon/3$$

$$p(ab|\bigcirc, \text{my language}) \propto \epsilon/3$$

LIKELIHOOD

$$p(signal|meaning, l).\,[1/ambiguity(signal|l)]^{\alpha}$$

My language

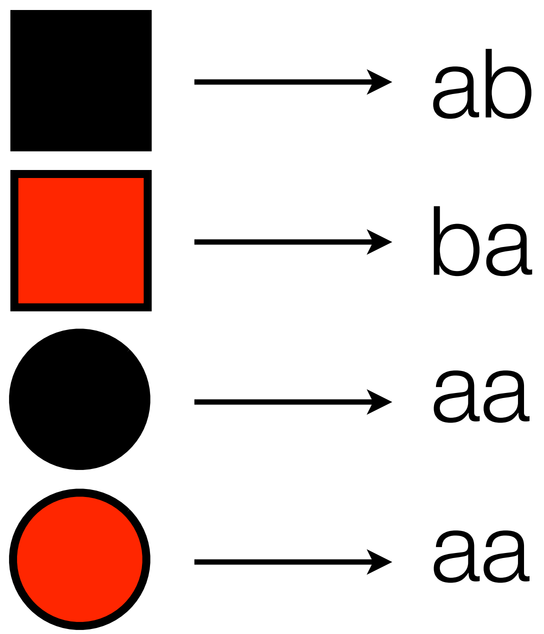

$$p(aa|\bigcirc, \text{my language}) \propto (1-\epsilon)(1/2)^{\alpha}$$

$$p(ba|\bigcirc, \text{my language}) \propto \epsilon/3$$

$$p(bb|\bigcirc, \text{my language}) \propto \epsilon/3$$

$$p(ab|\bigcirc, \text{my language}) \propto \epsilon/3$$

A COMPRESSION-BASED PRIOR

- Our hypothesis: learners prefer simpler languages
- In the model: assign p(l) depending on **coding length** of grammars
  - Complex grammars, low prior probability
  - Simple grammars, high prior probability

CODING LENGTH

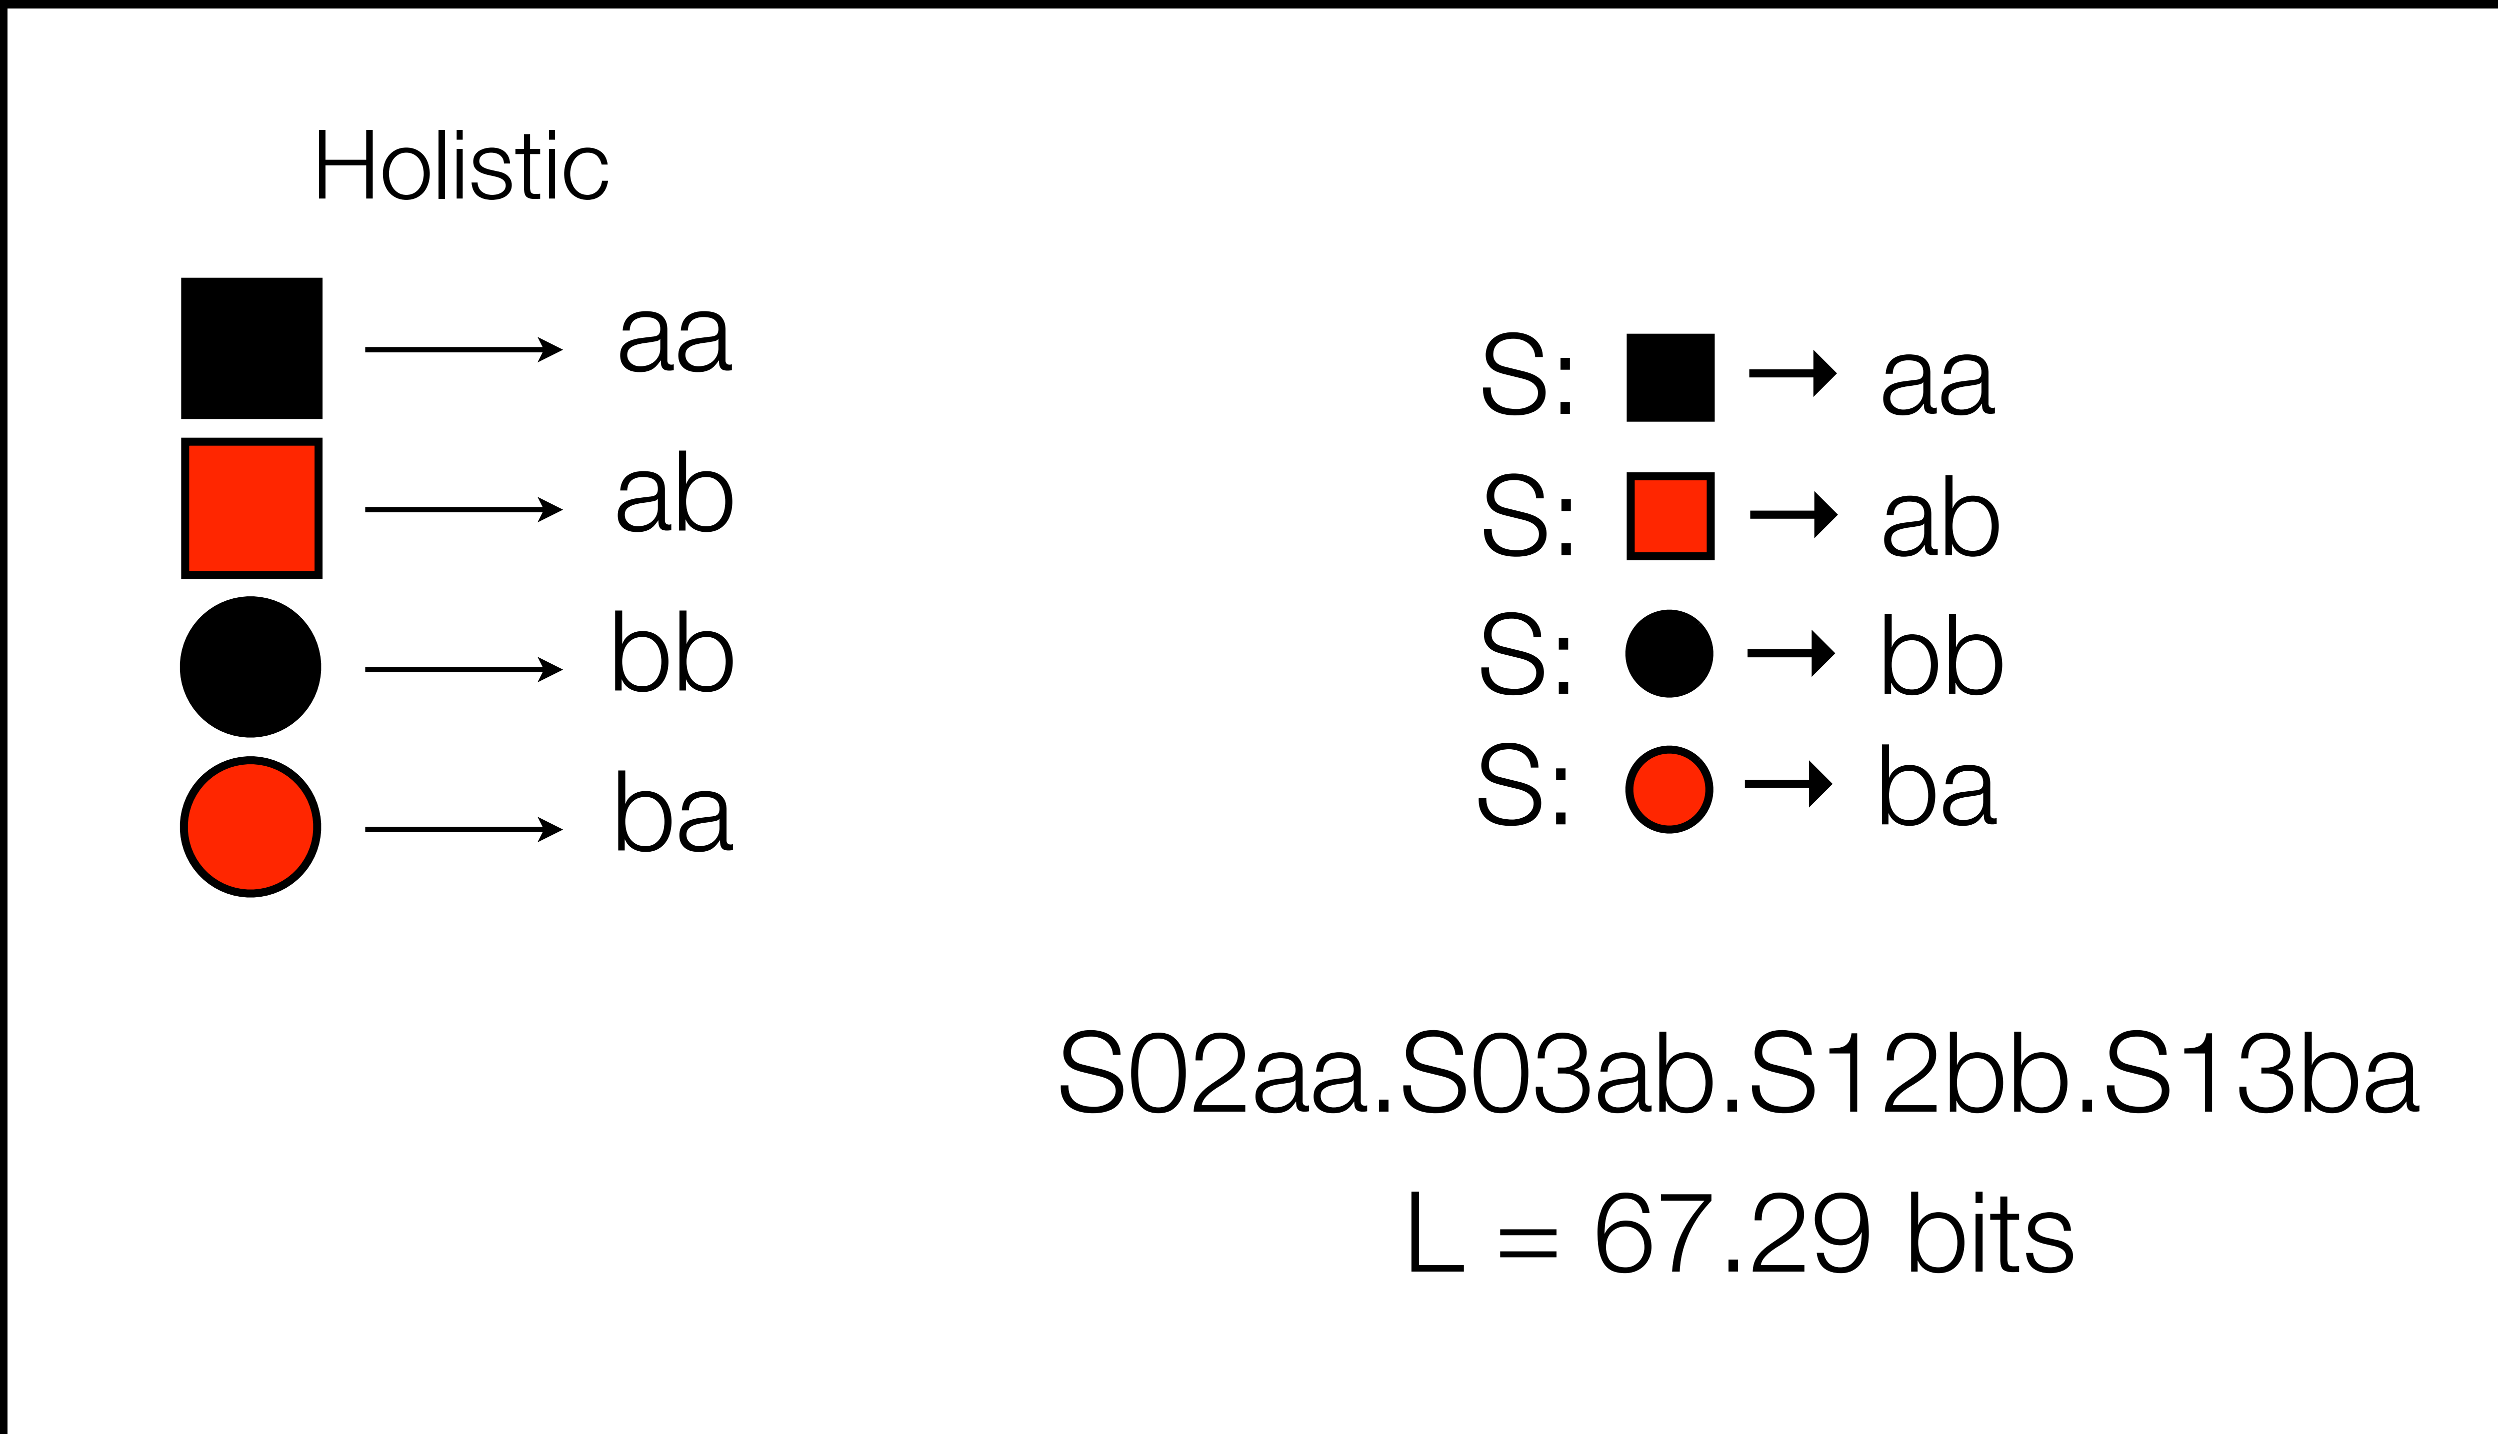

CODING LENGTH

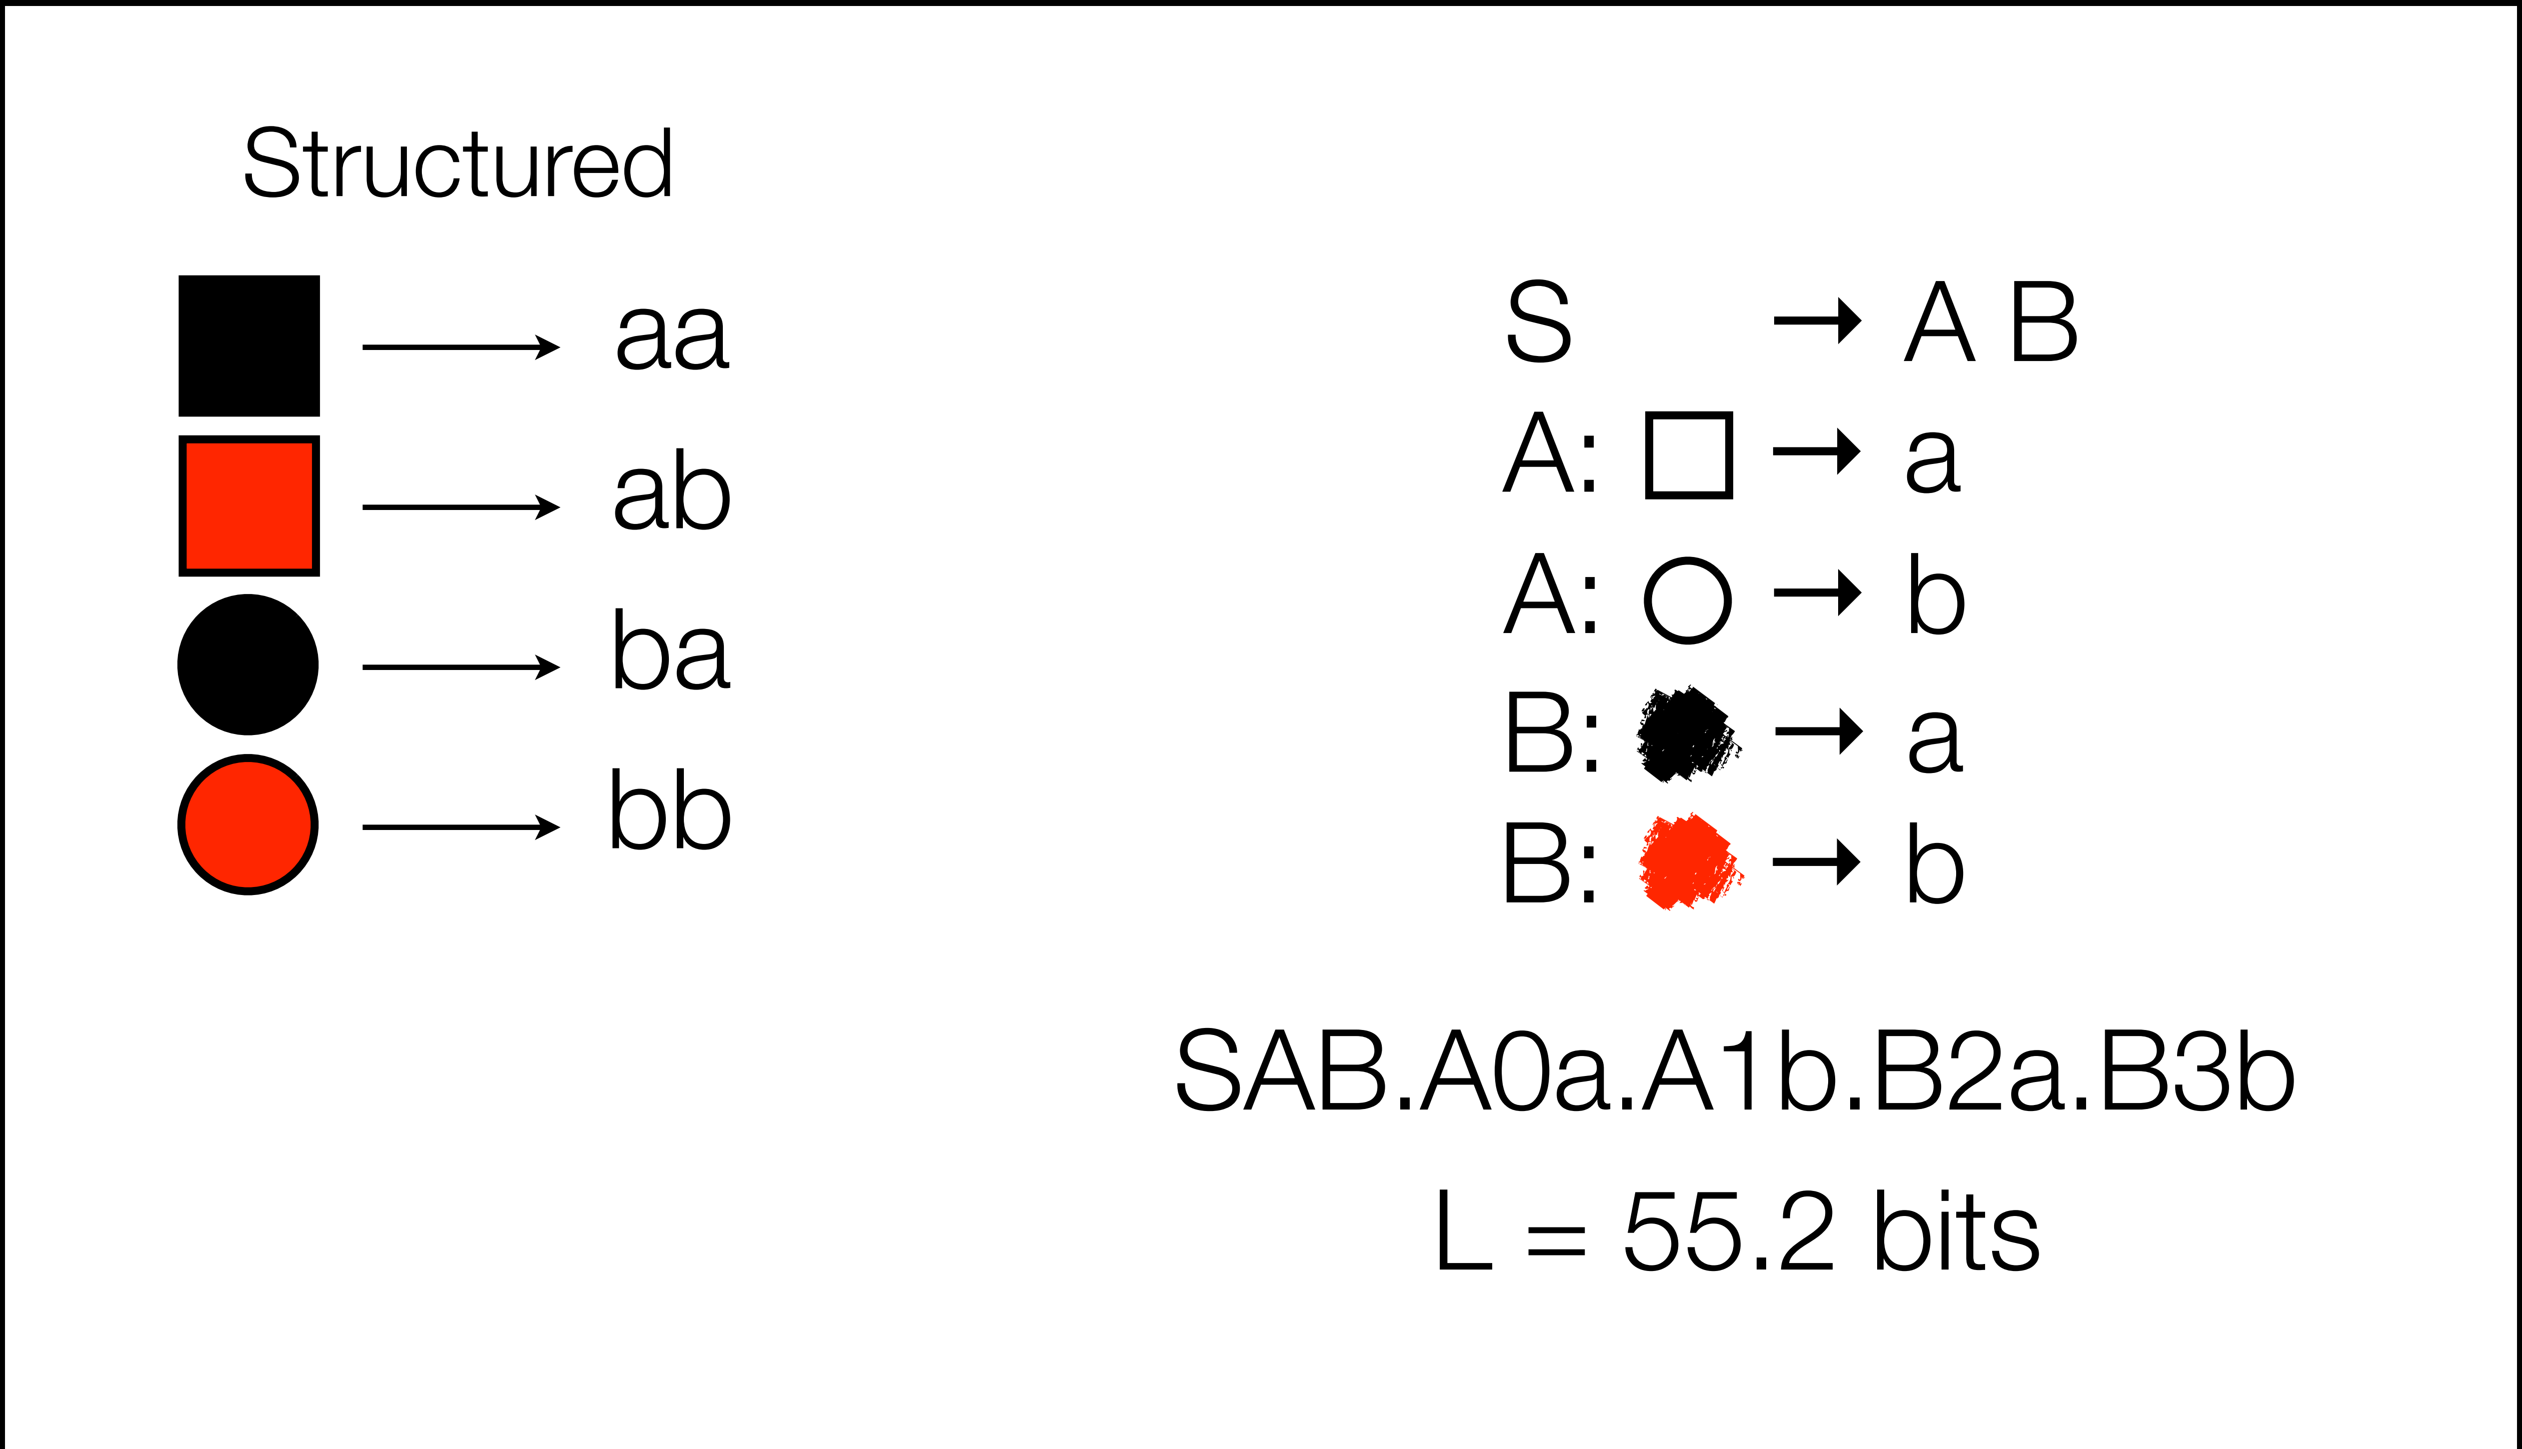

CODING LENGTH

Degenerate

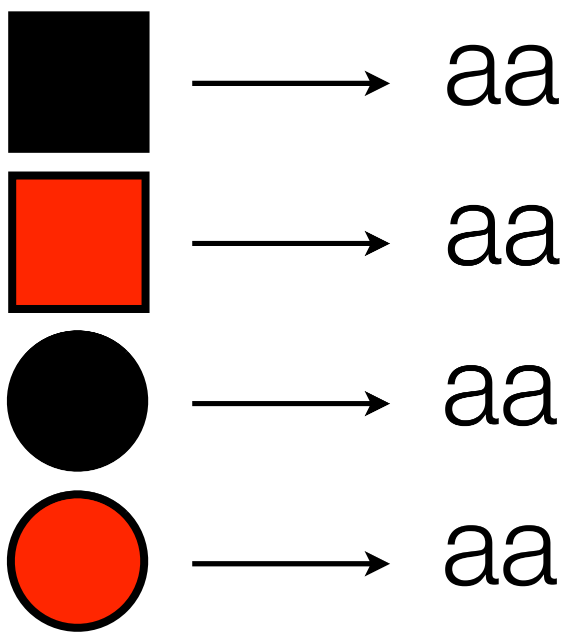

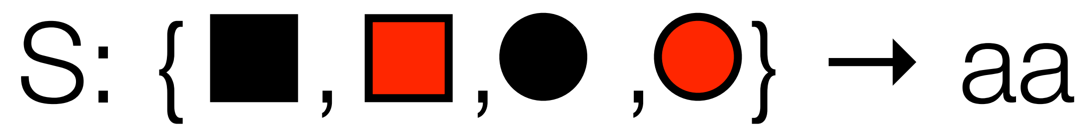

S02,03,12,13aa

L = 38.55 bits

FROM CODING LENGTH TO P(L)

Language	L	$p(l) = 2^{-L}$
S02,03,12,13aa	38.55	0.399
SAB.A0a.A1b.B2a.B3b	55.2	0.00000388
S02aa.S03ab.S12bb.S13ba	67.29	0.0000000009

THE PRIOR

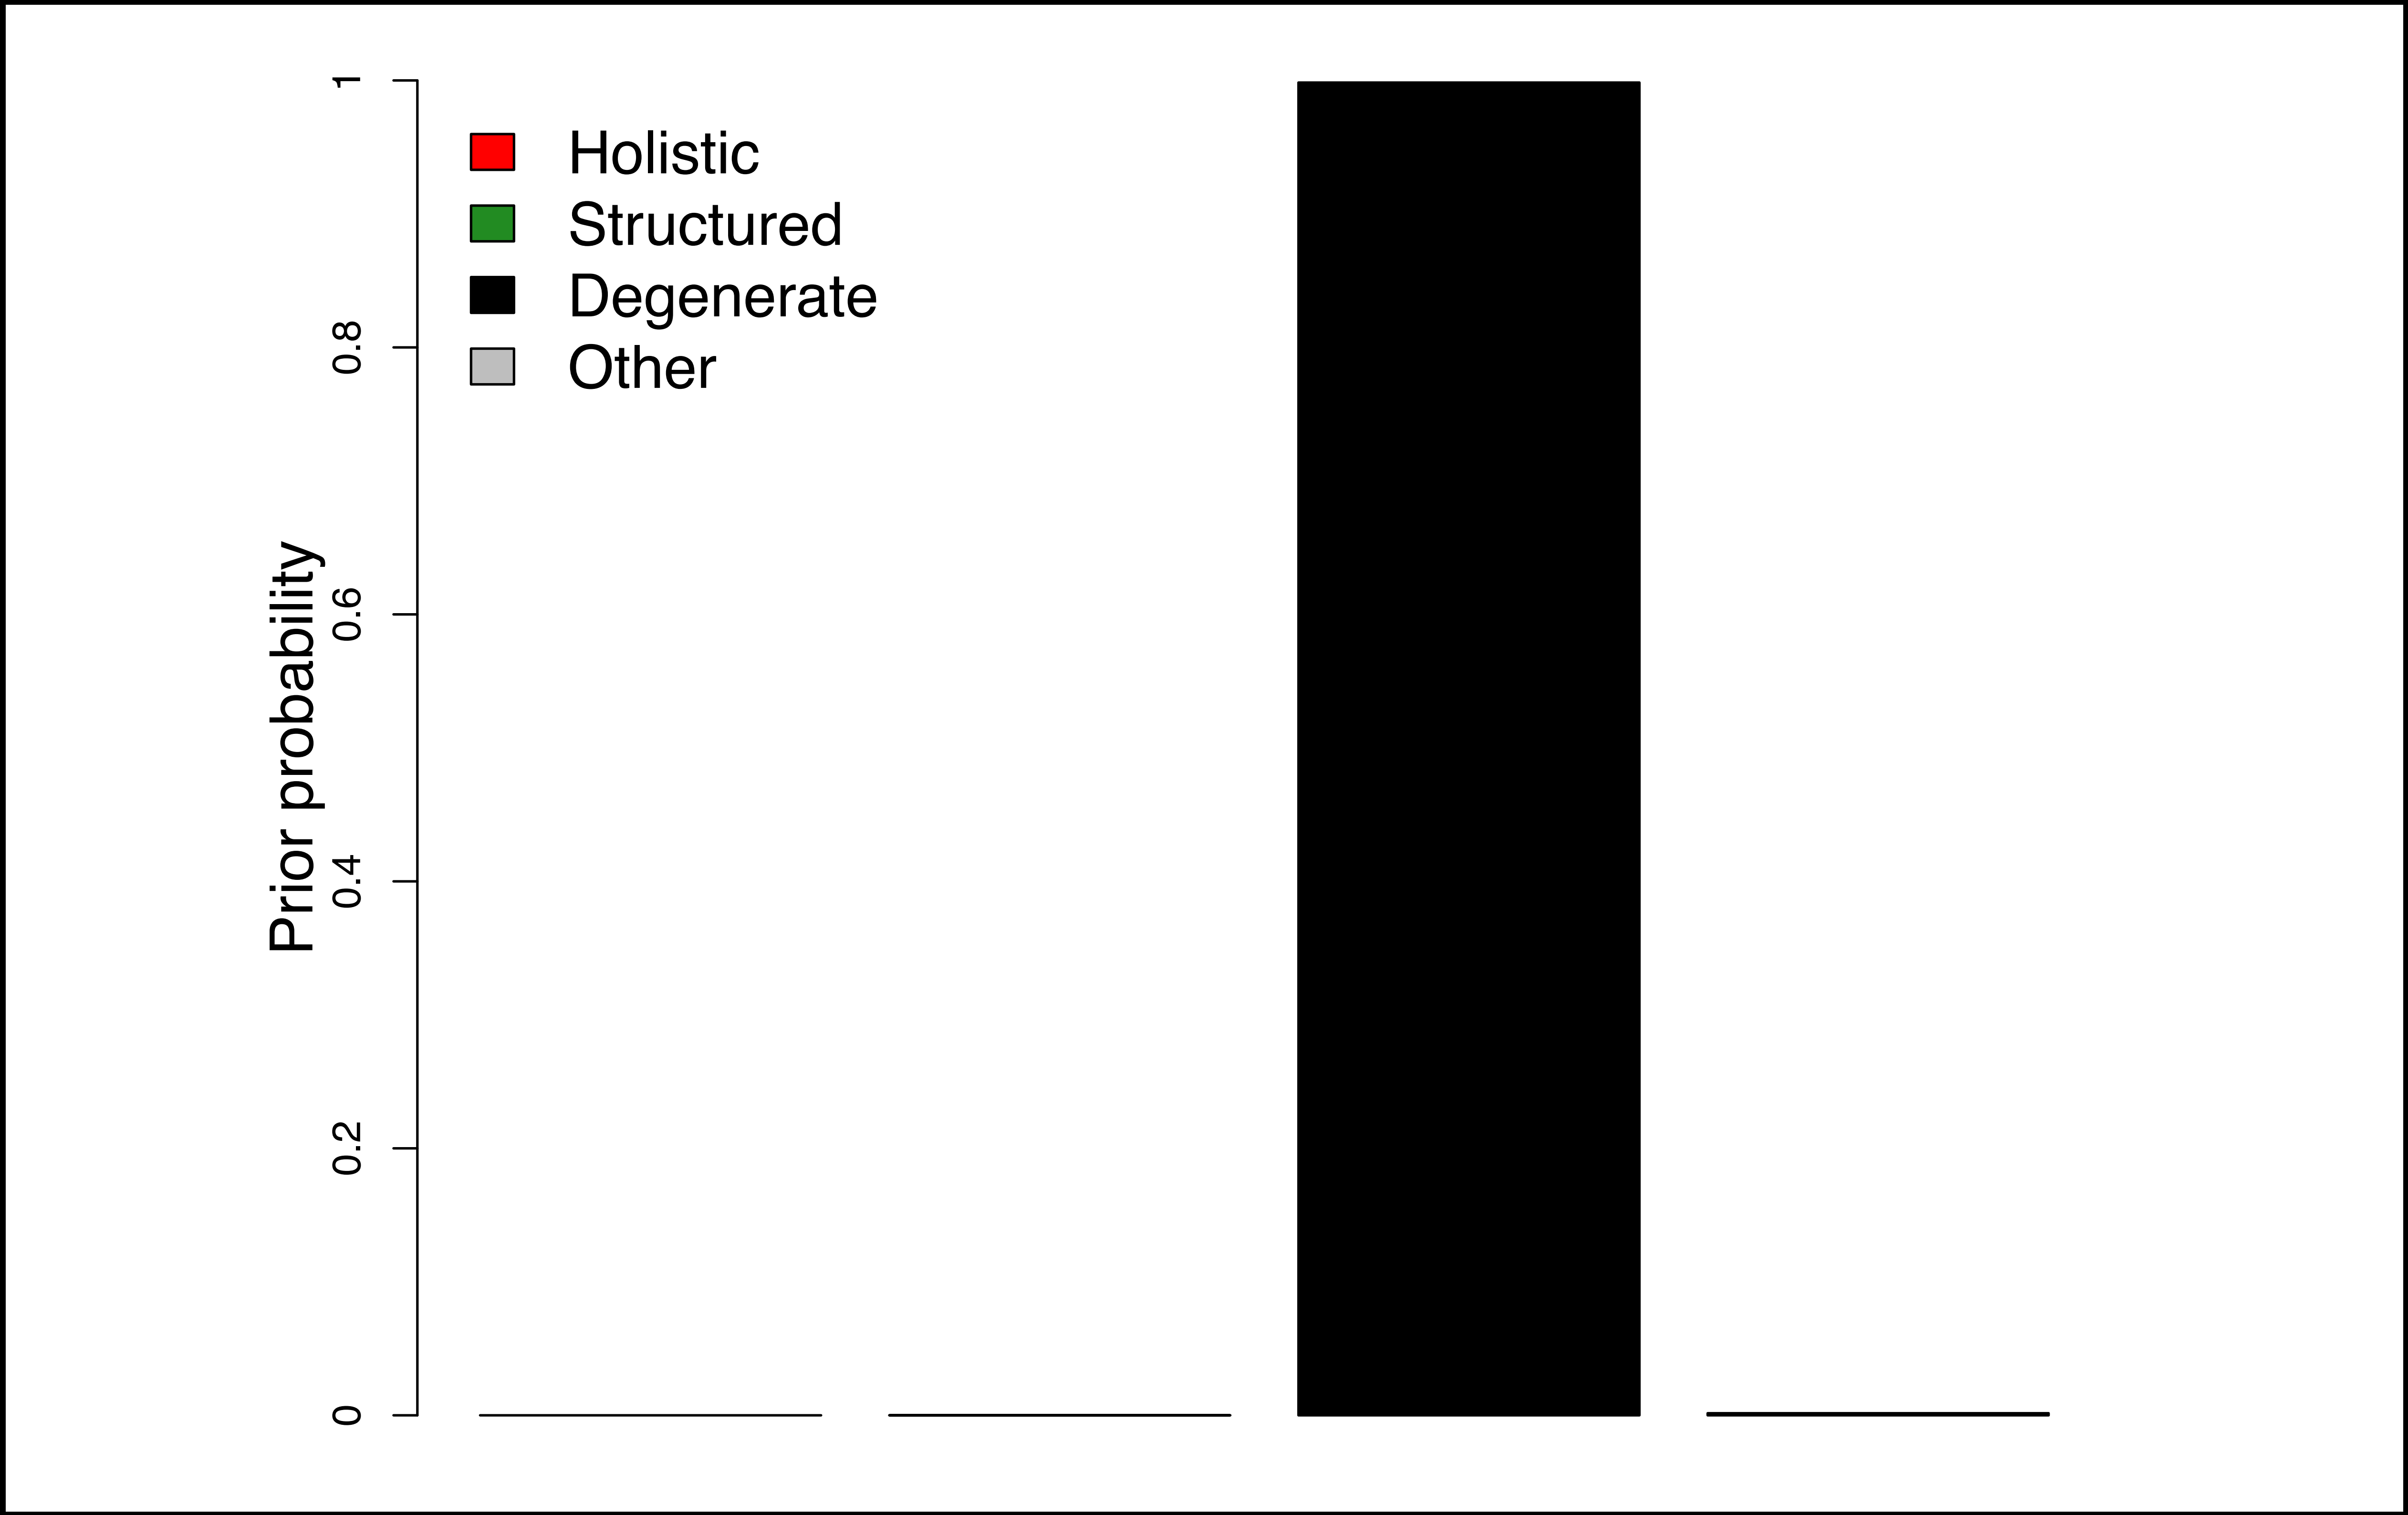

Holistic
Structured
Degenerate
Other
Prior probability
0
0.2
0.4
0.6
0.8
1

LEARNING AS BAYESIAN INFERENCE

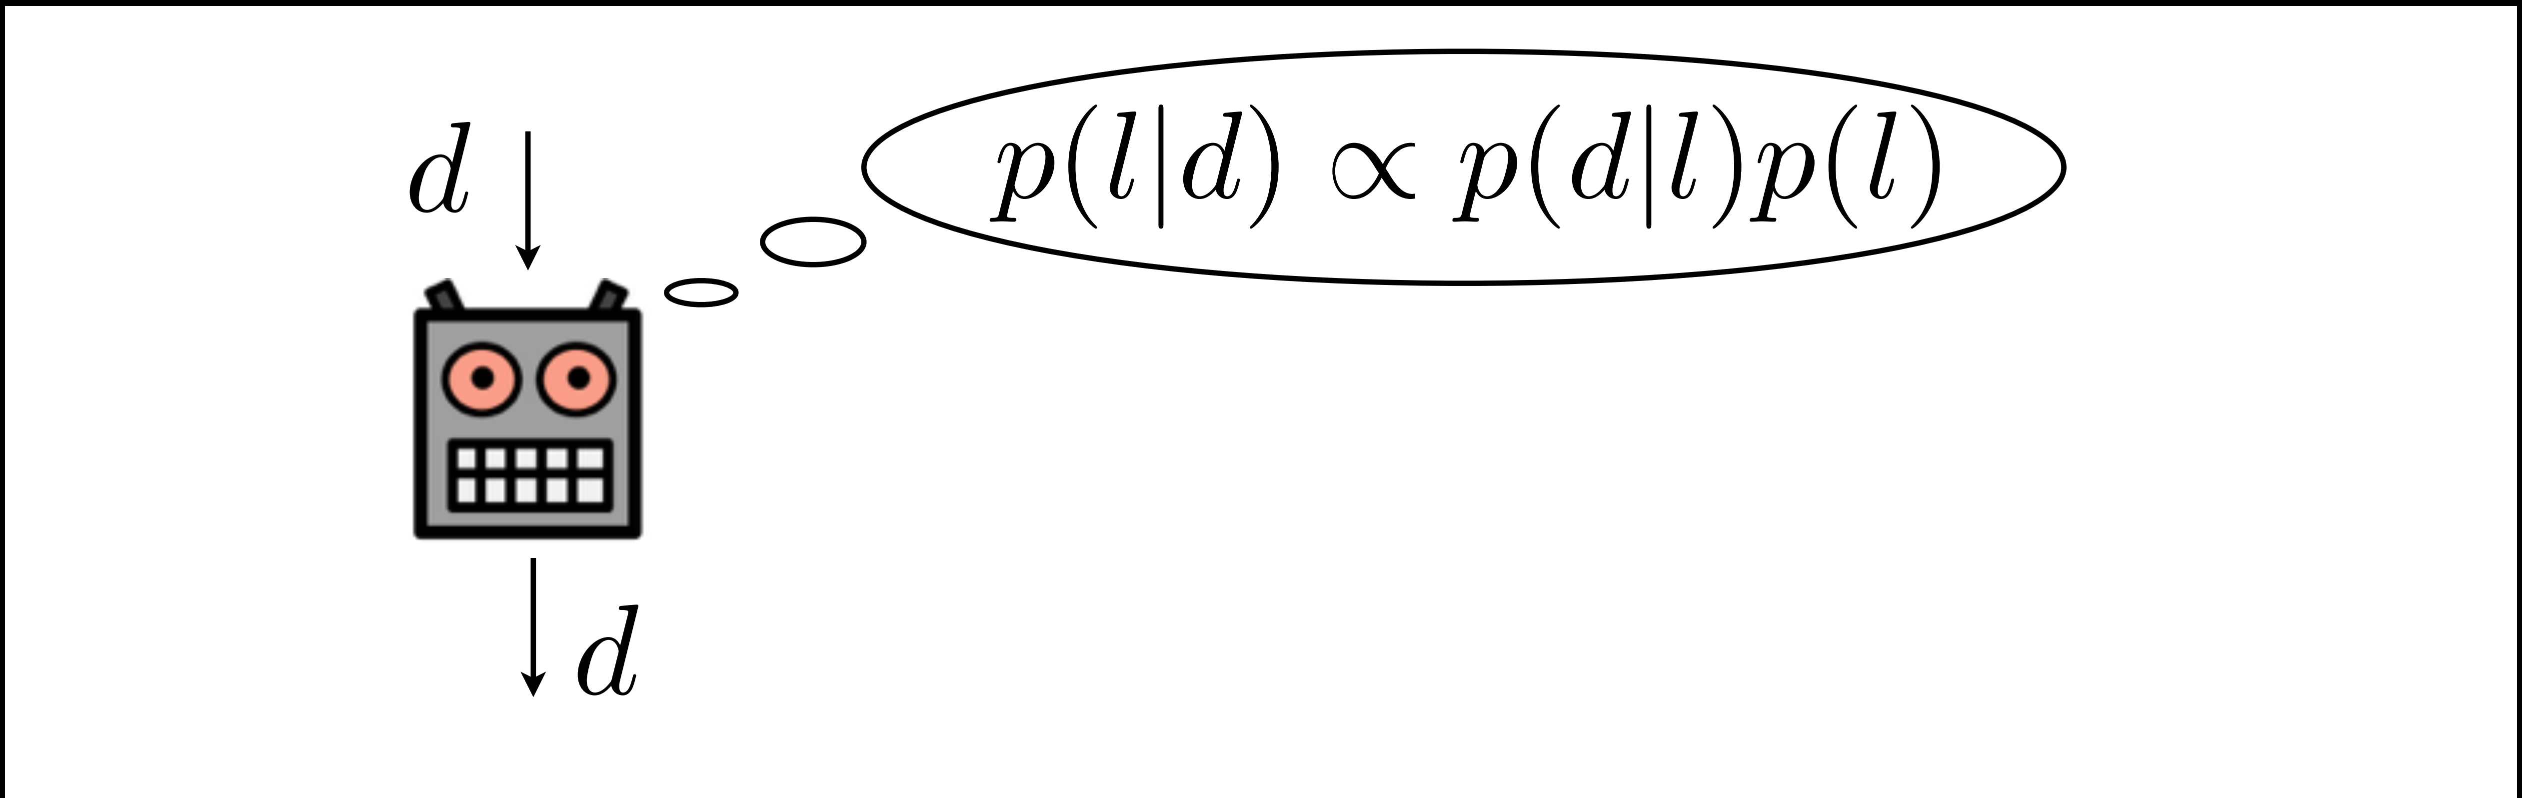

Data: Meaning-signal pairs

Languages: specify signal for each meaning

Likelihood p(d|l): avoid ambiguous utterances

Prior p(l): prefer compressible languages (shorter coding length)

Simulation of Exp 1

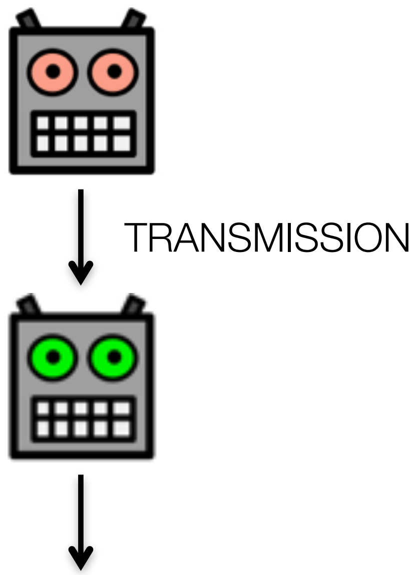

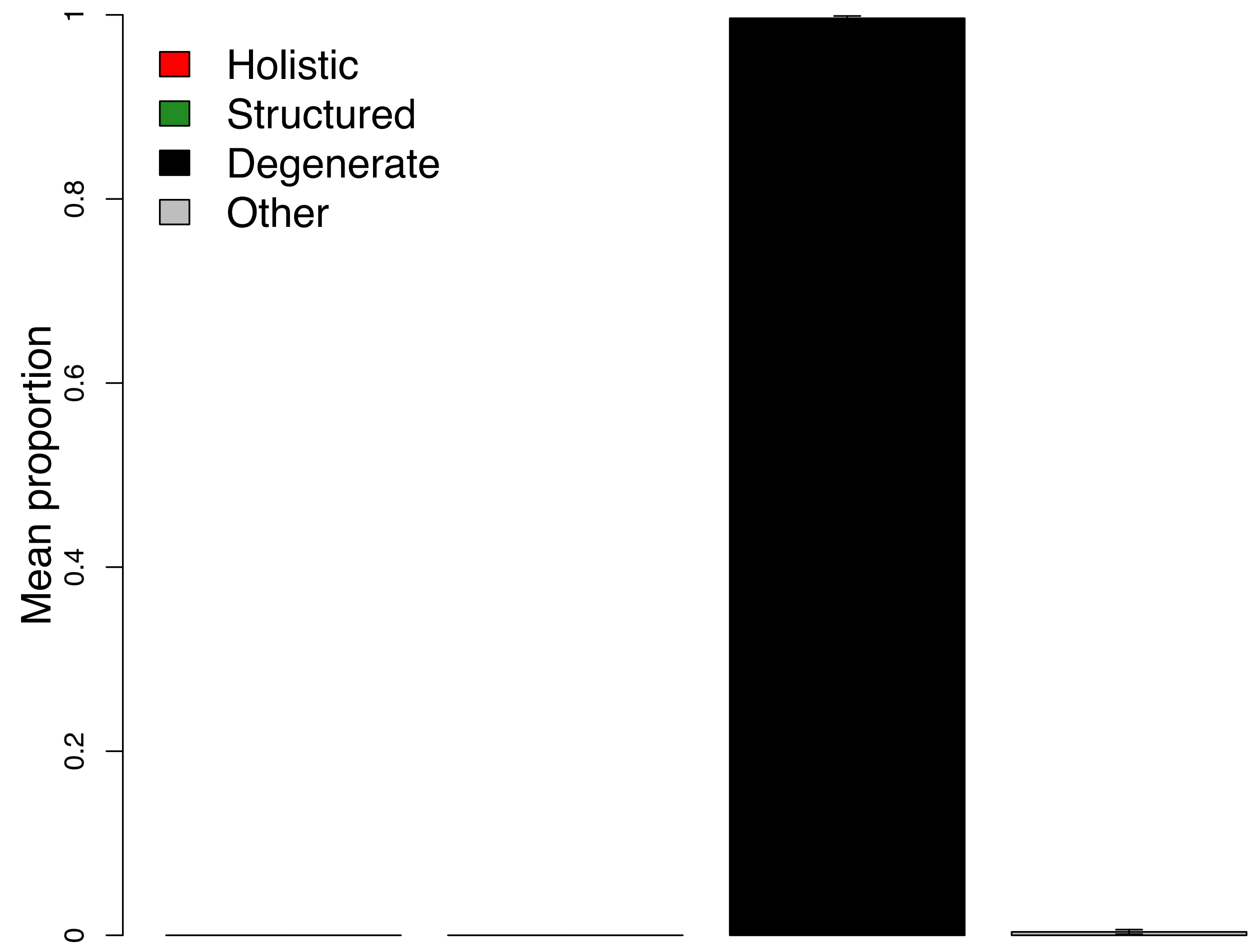

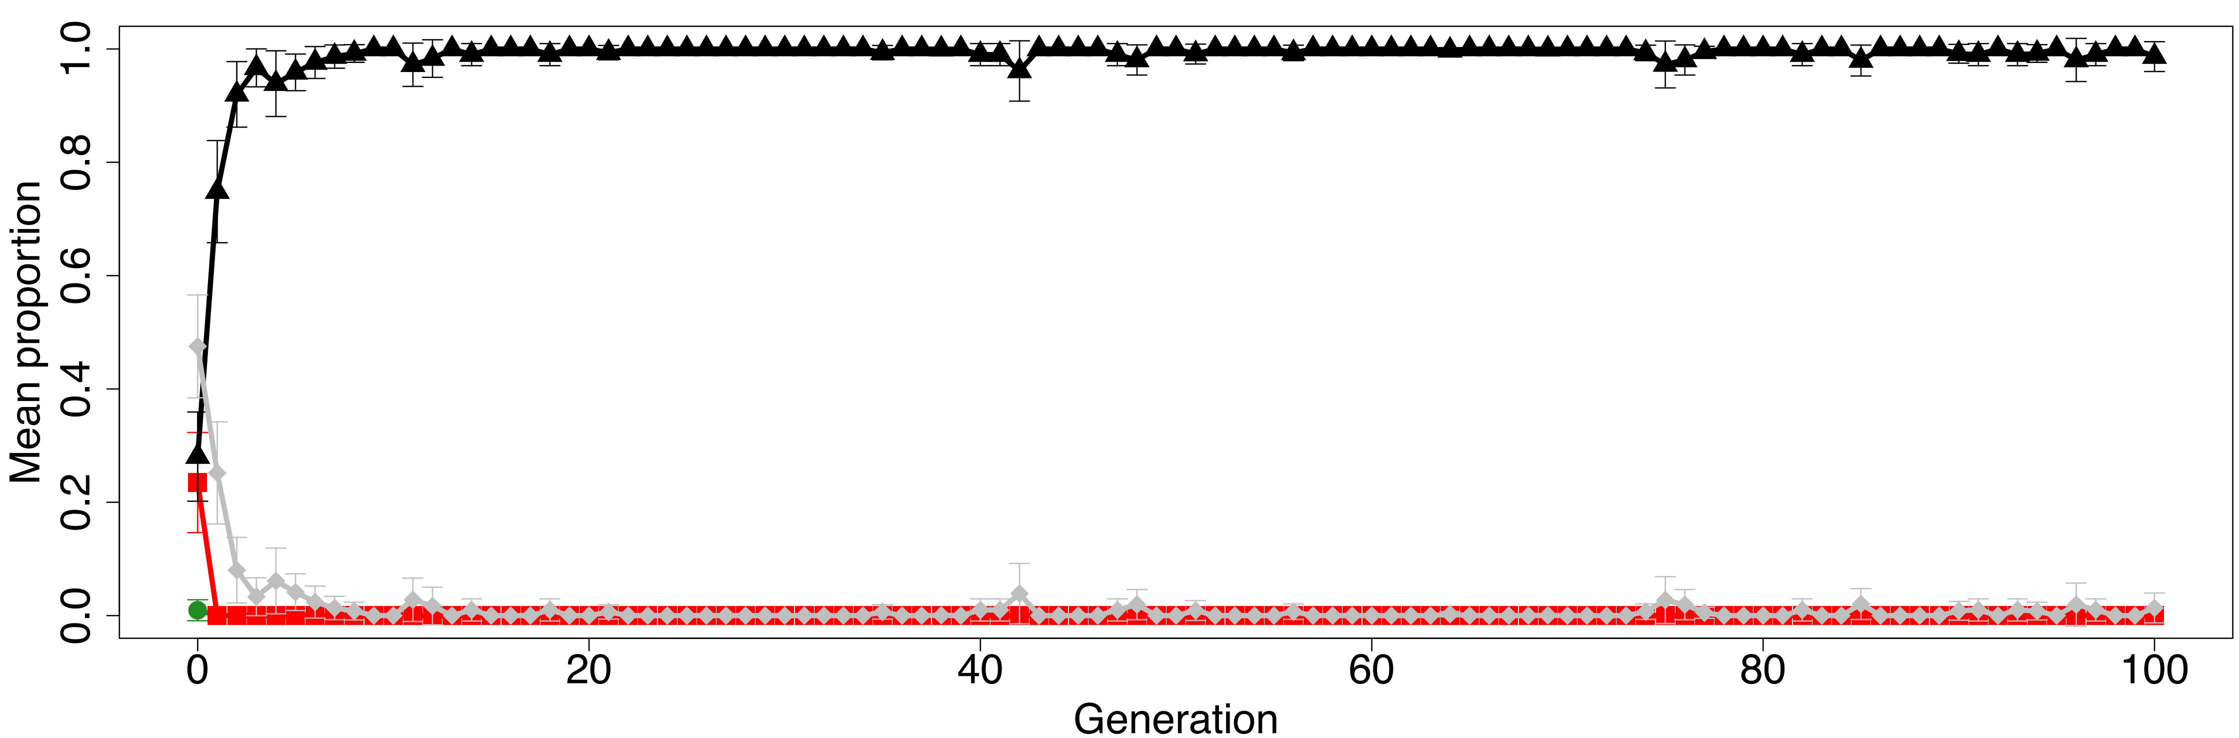

Simulation of Exp 2

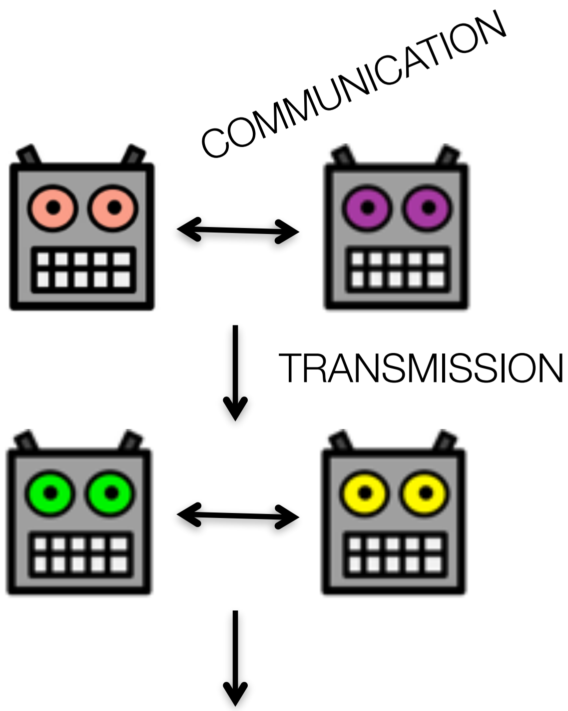

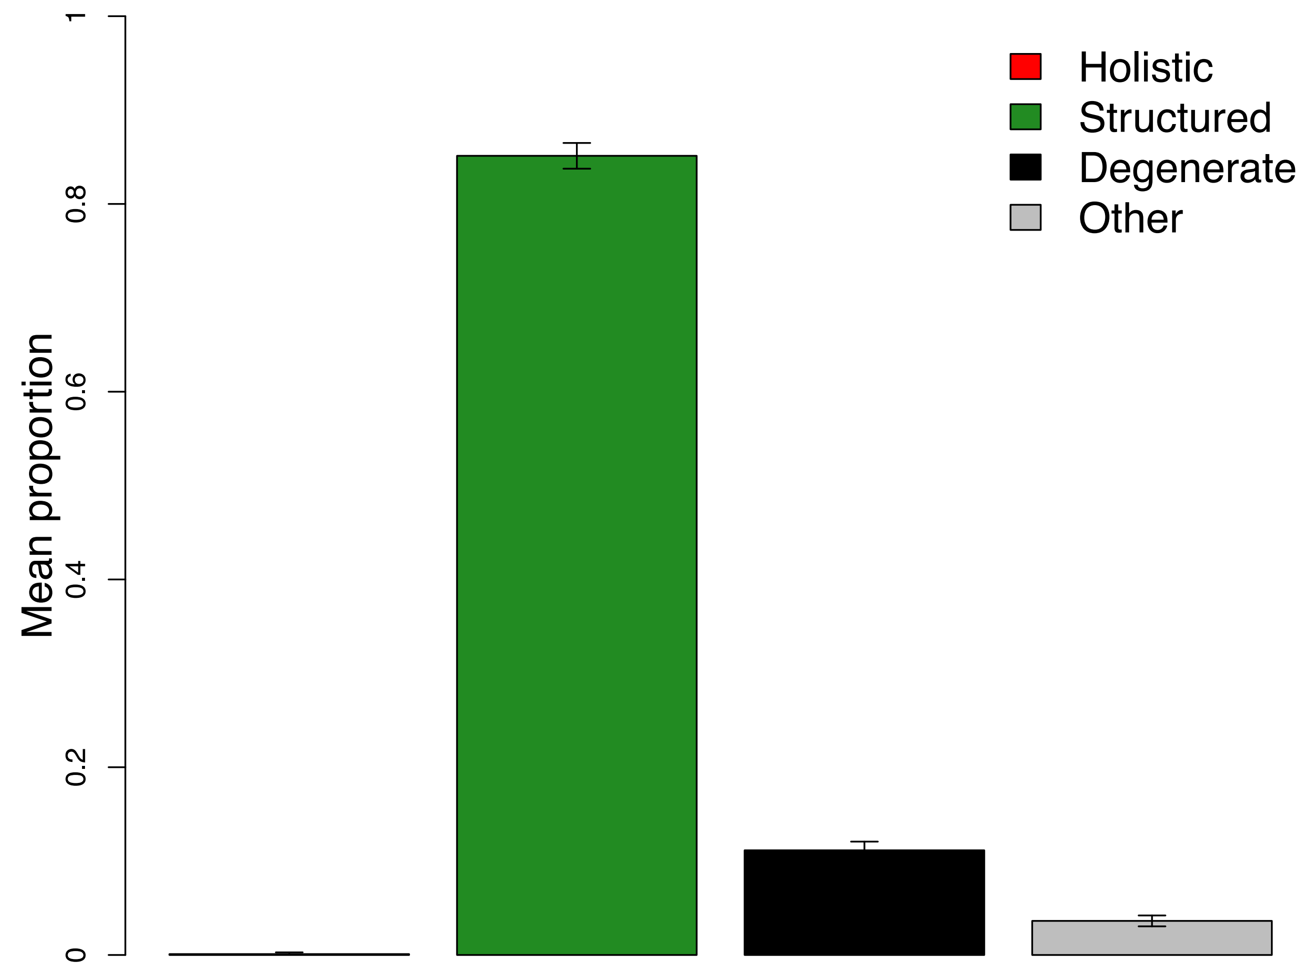

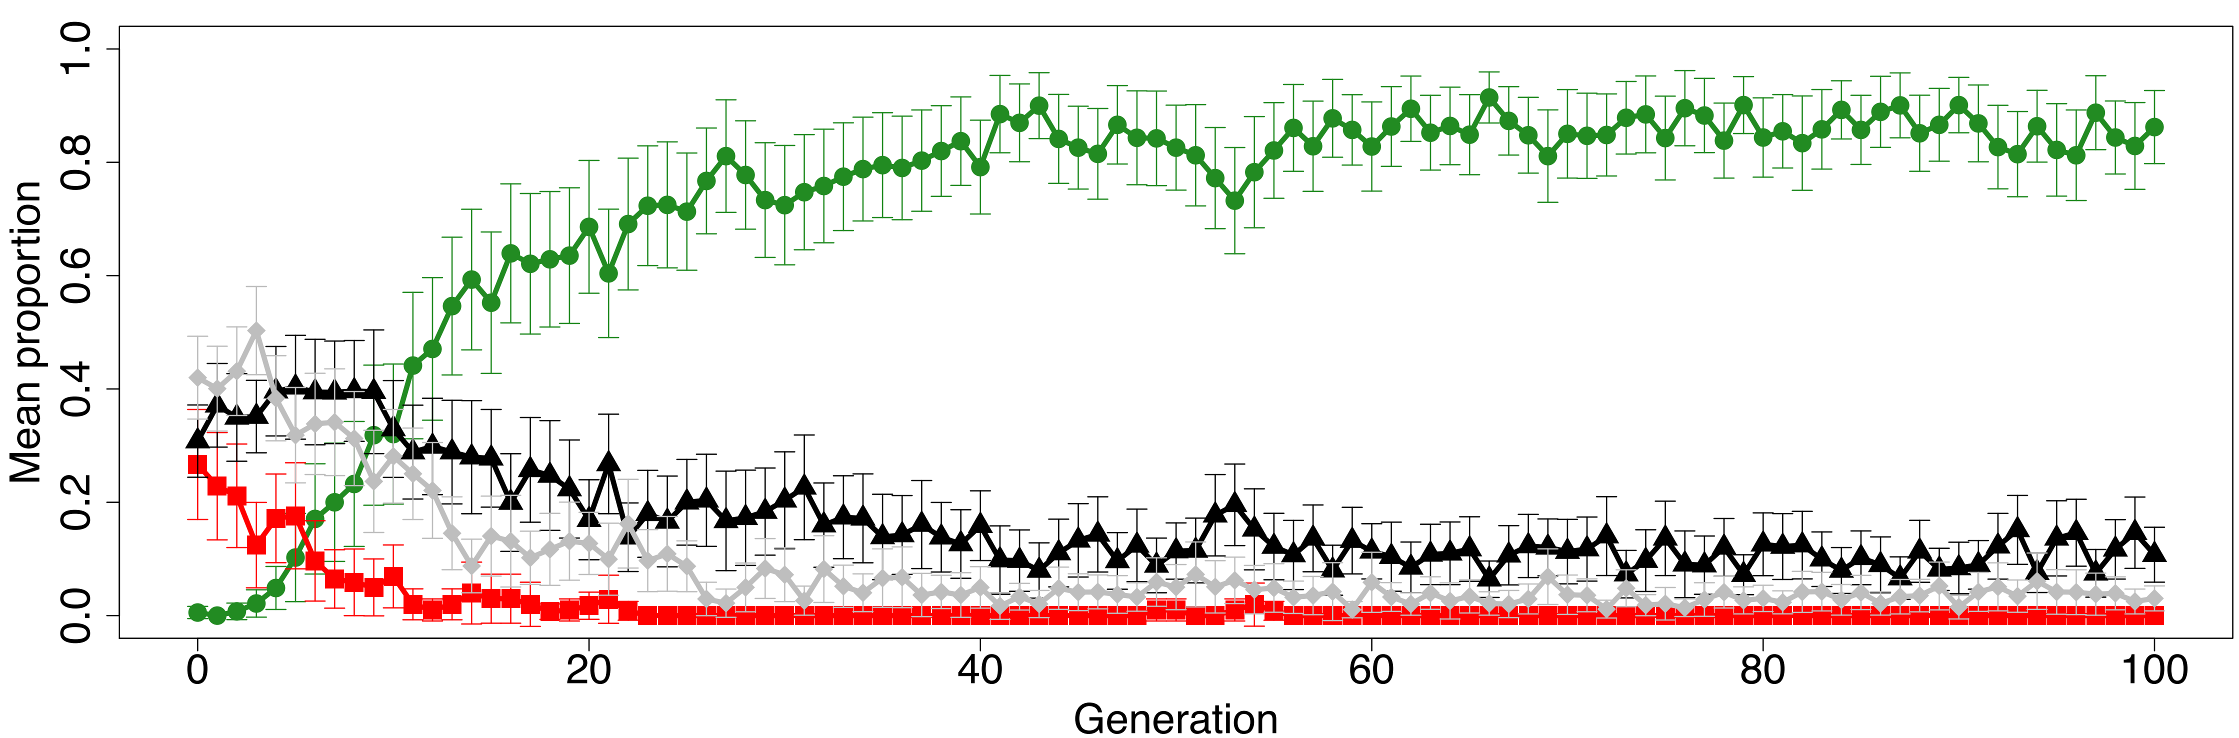

Simulation of Exp 3

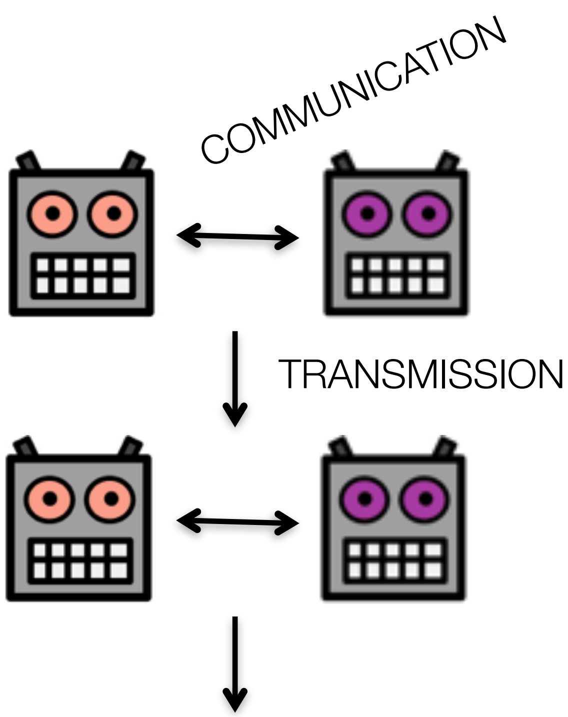

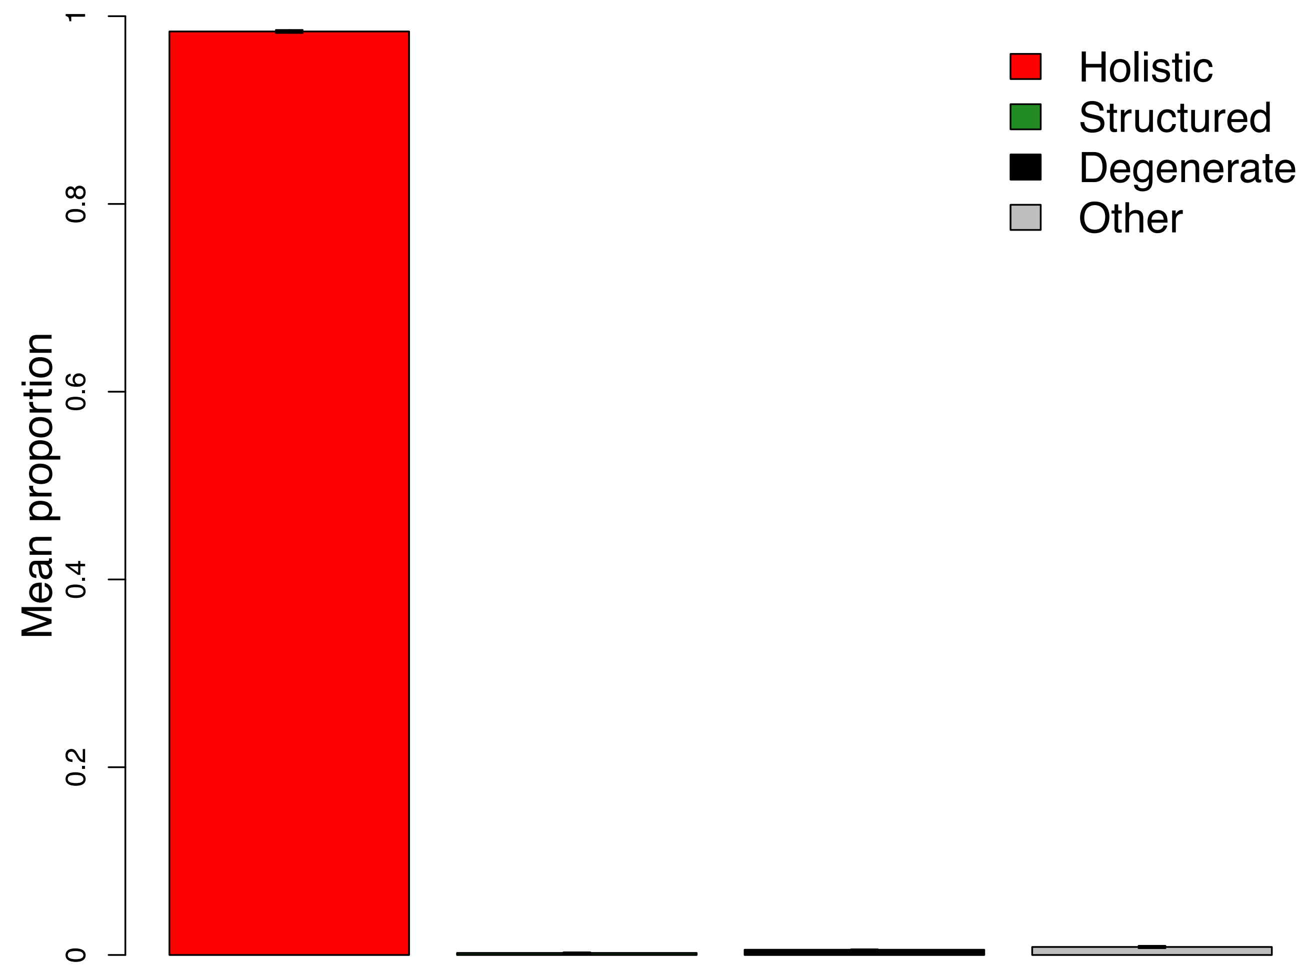

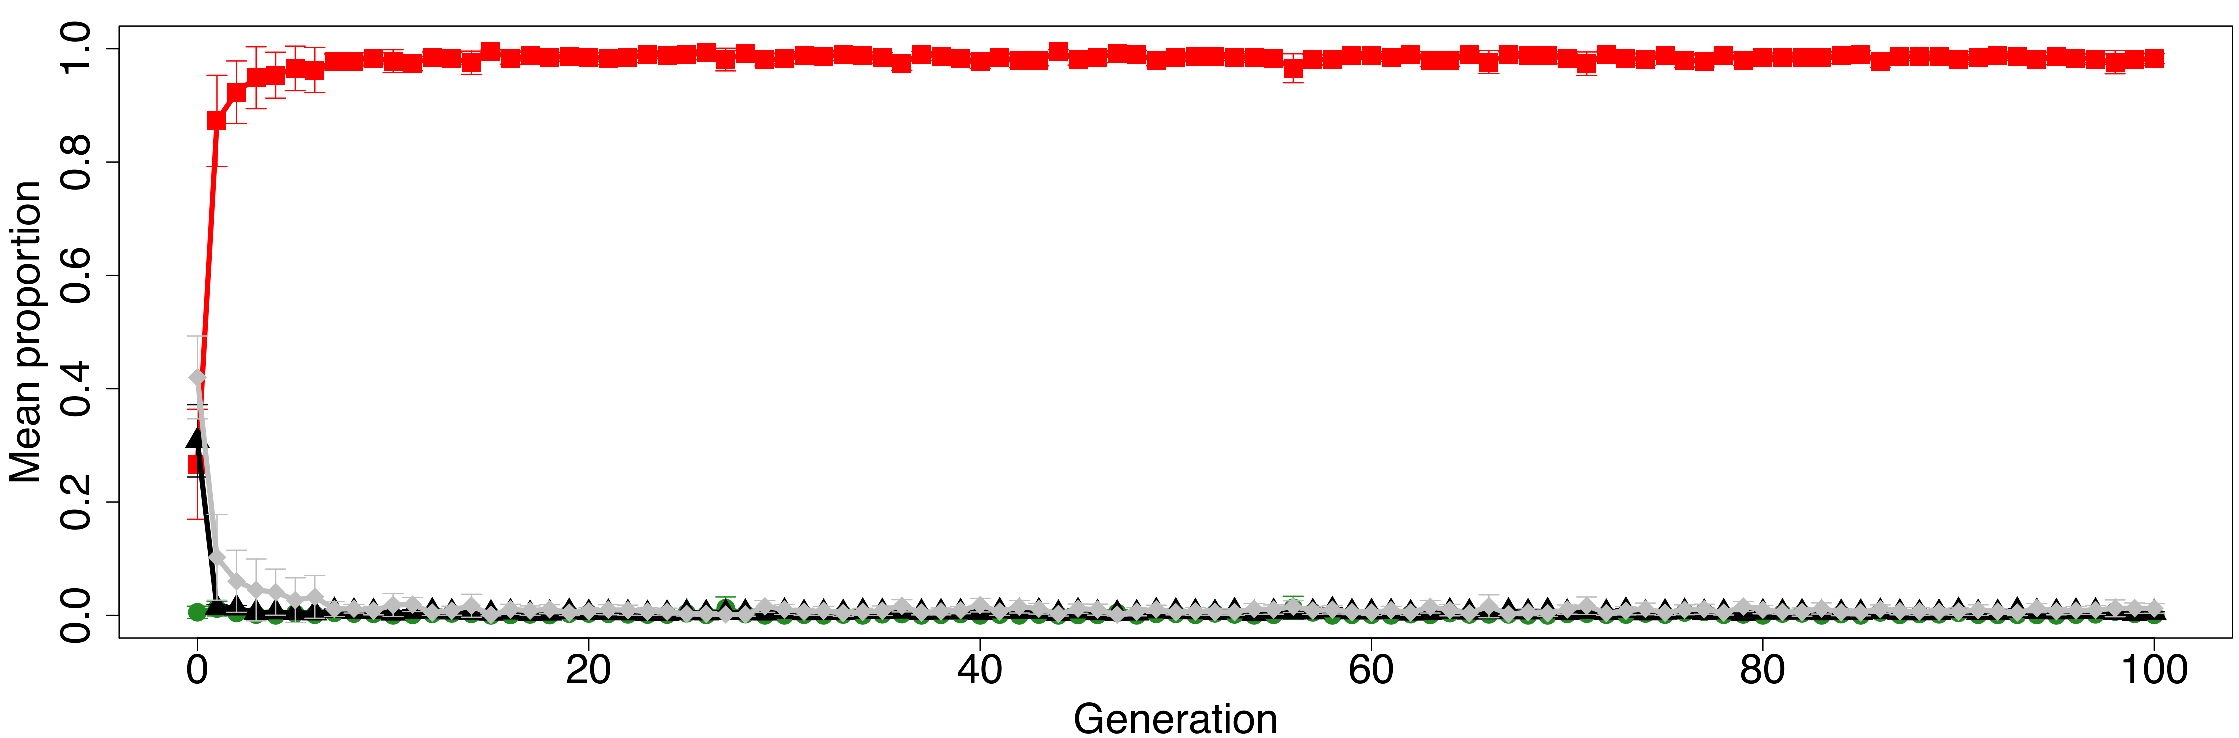

SIMULATIONS REPLICATE EXPERIMENT RESULTS

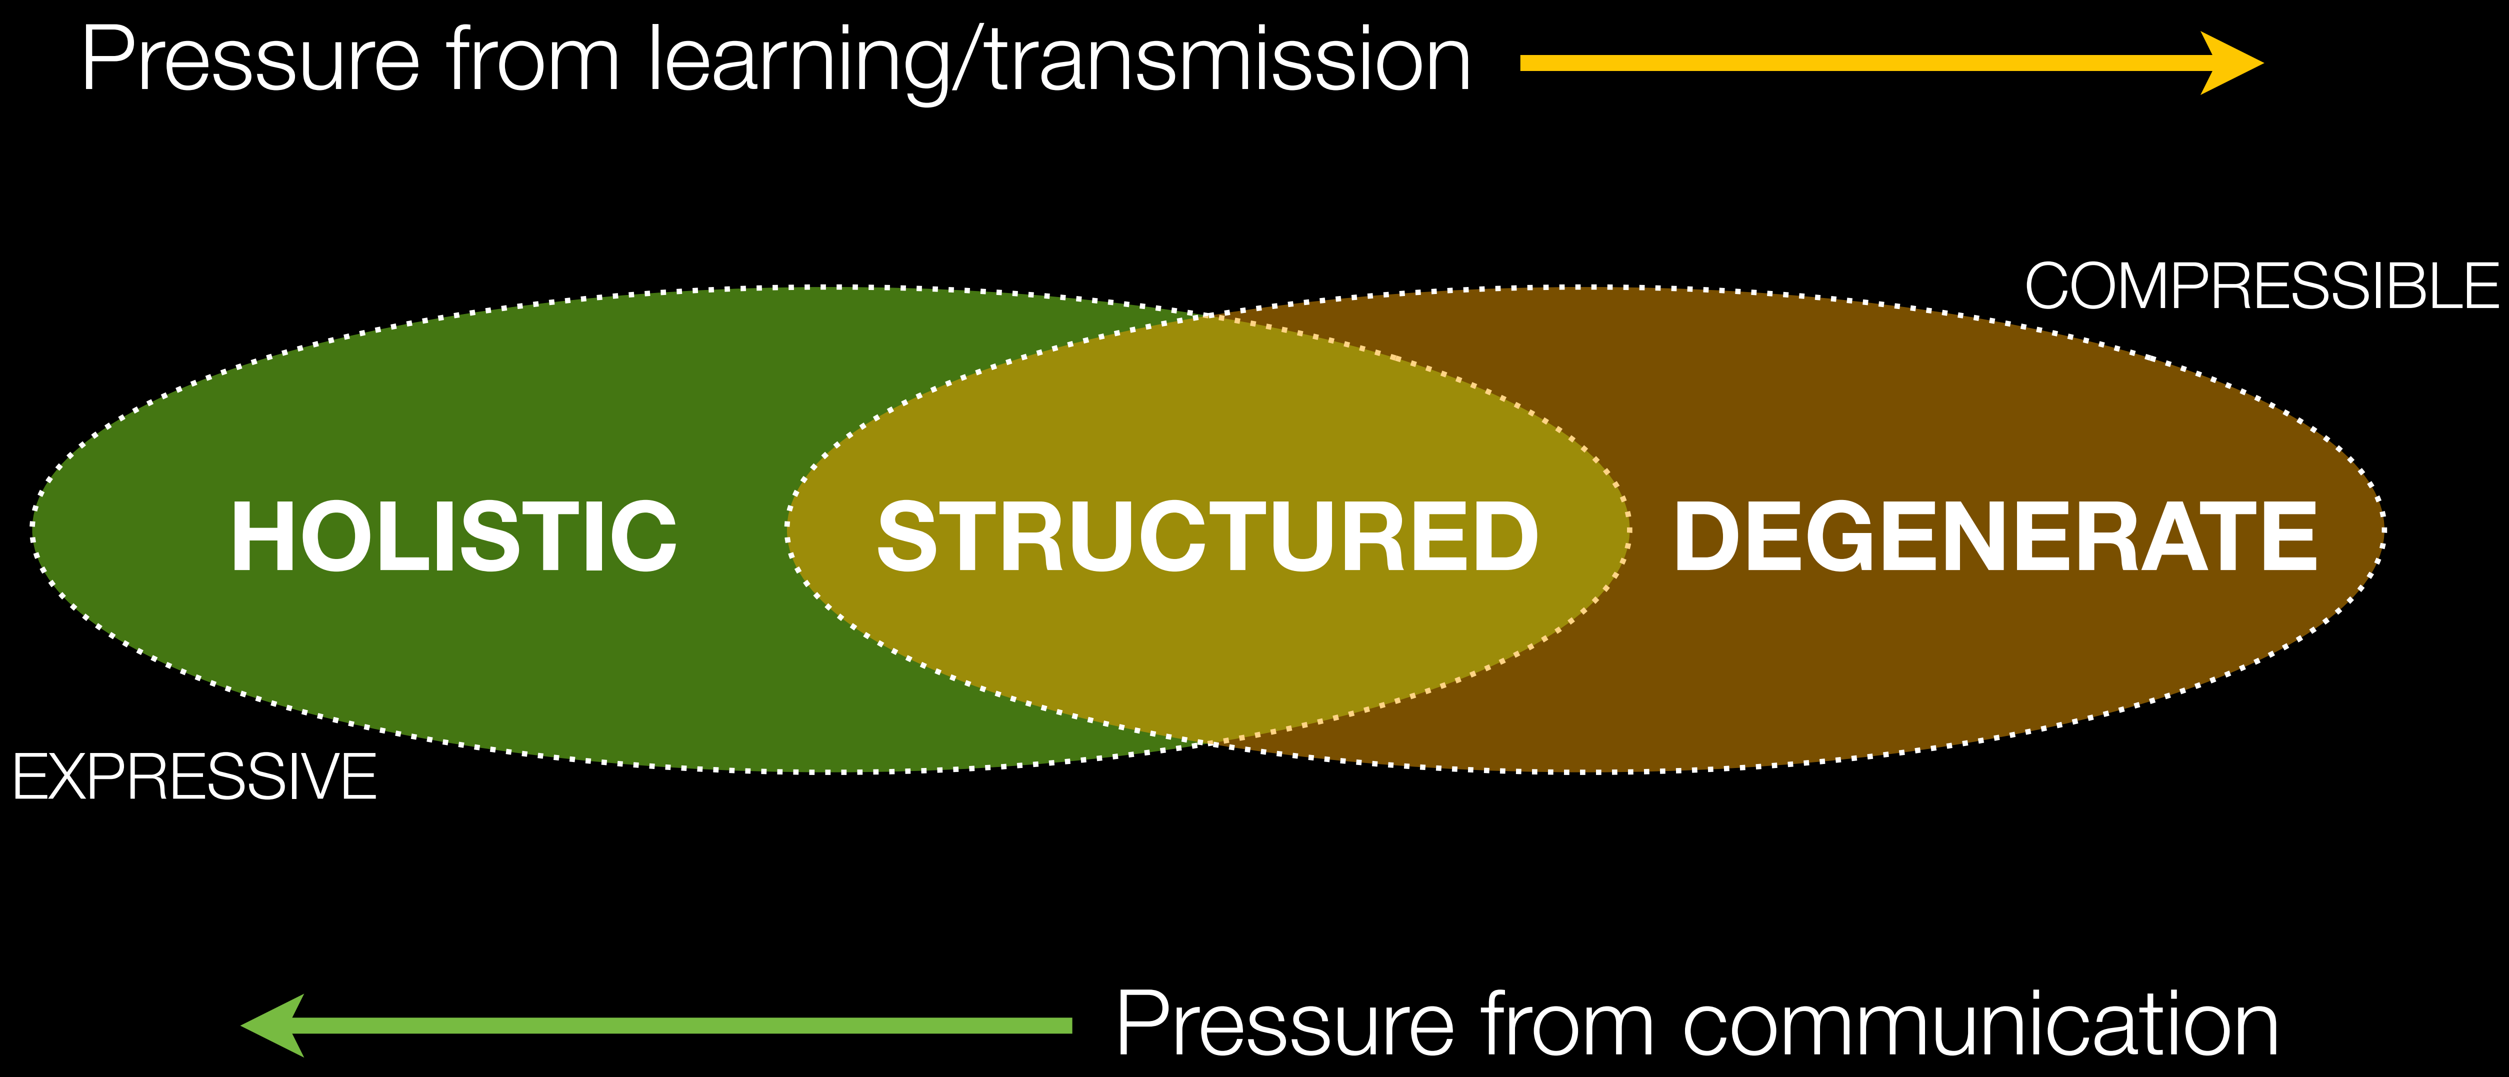

PART 3:
FROM EXPLICIT NAMING TO “EMBODIED” INTERACTION

THE EMERGENCE OF CULTURALLY TRANSMITTED COMMUNICATION

- The previous paradigm assumed that individuals:
    - want to communicate
    - know what to communicate about
    - have a dedicated “channel” for communication
    - want to share their communication system
- In other words, they are already *symbolic learners*
    - Can we explore the genuine emergence of symbols in the lab?
    - Experiment inspired by study in evolutionary robotics (Quinn 2001)

STUDY 2:
THE EMBODIED COMMUNICATION GAME

- Two-player cooperative computer game where the other player is in a different room
- Steer a character round a room with different coloured floor tiles and try to finish up on the same colour as the other player
- Similar to a study by Galantucci (2005) but *without a communication channel*

Scott-Phillips, Kirby & Ritchie (2009)

Player 1 sees:

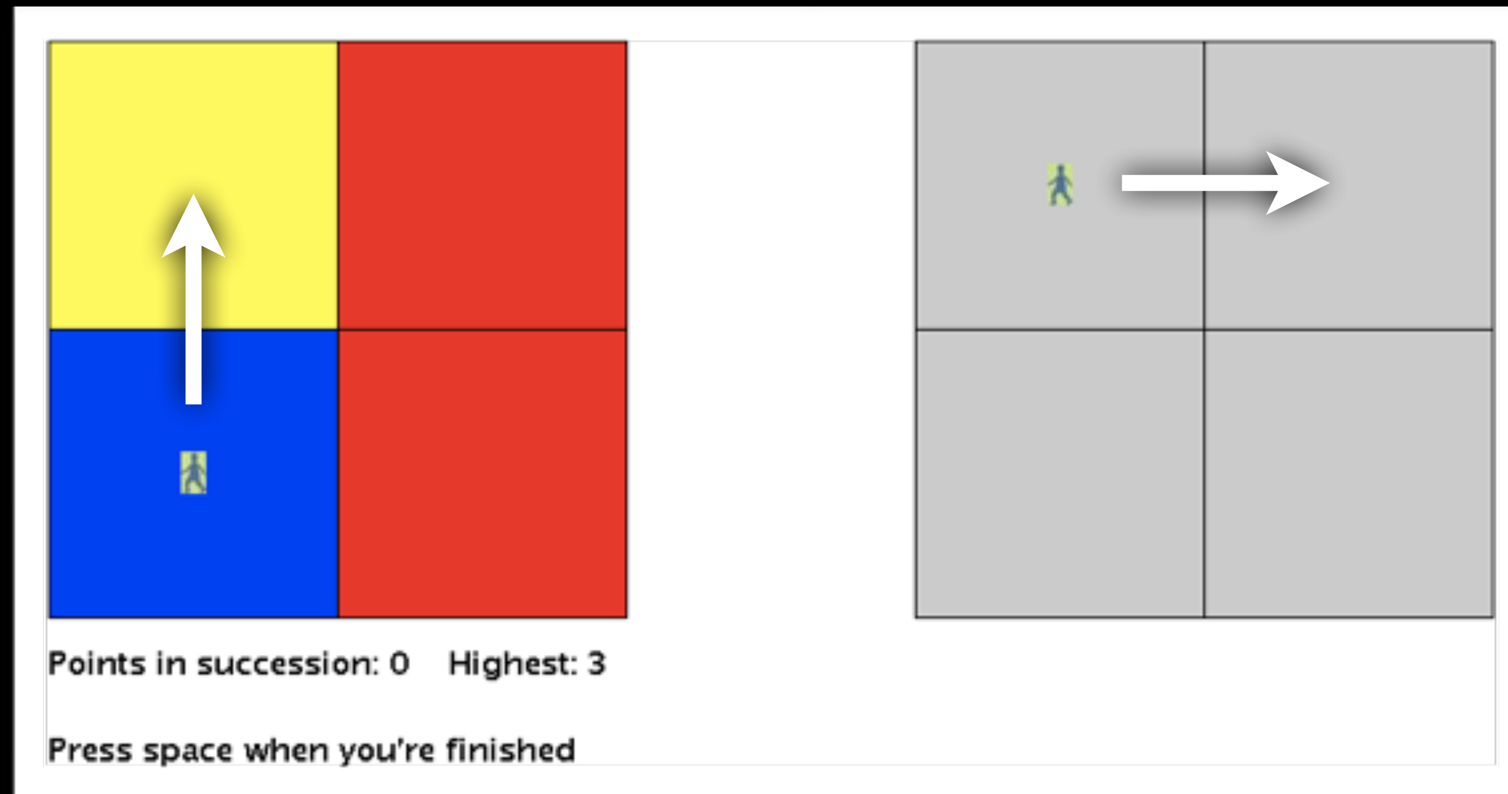

Player 2 sees:

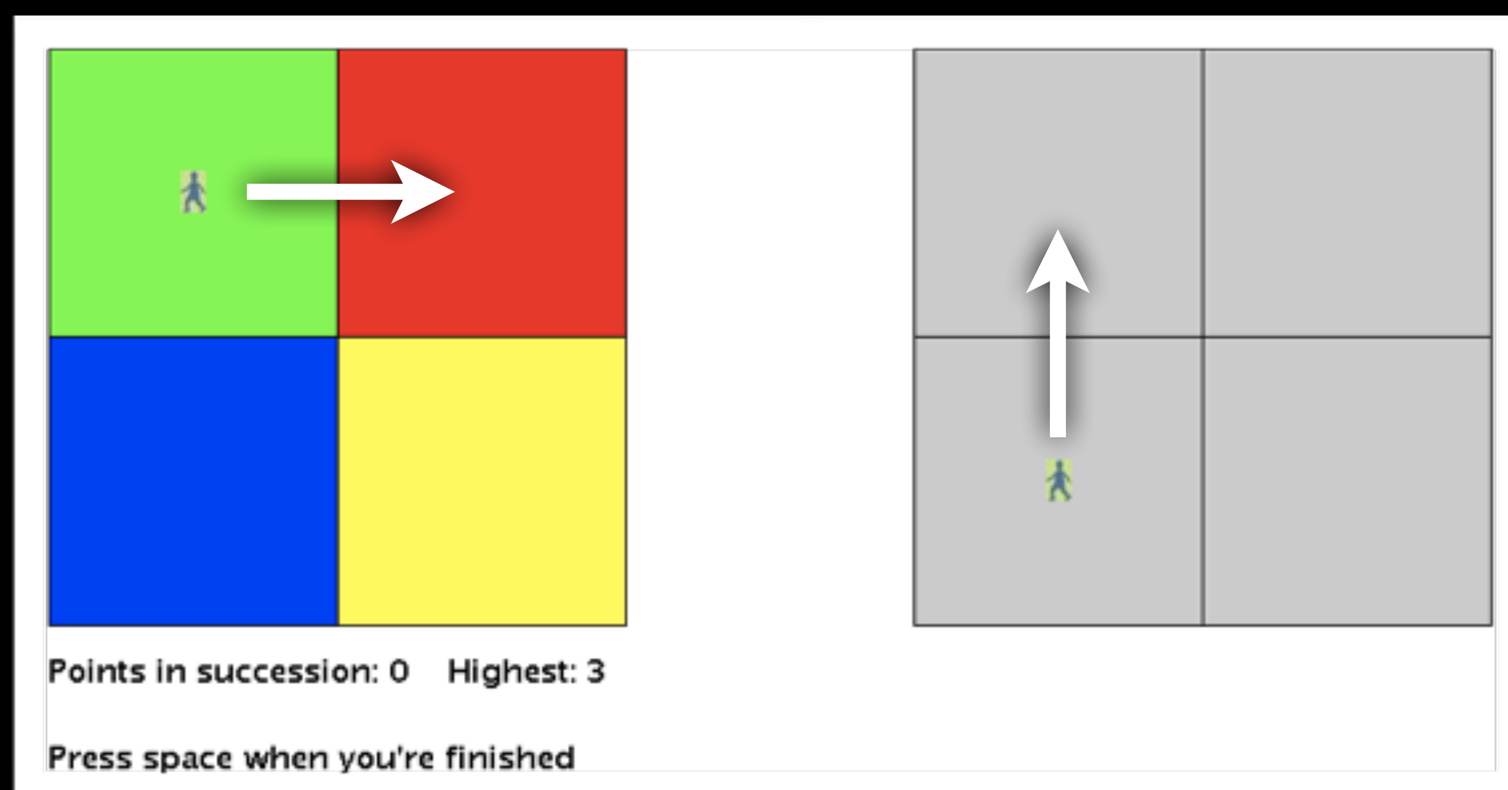

RULES

1. Score if on same colour after both press finish

2. Always at least one colour that's in both rooms (but equally there may be colours that are unique to room)

3. Colour assignment is completely random after each turn

4. After turn, other player's colours are revealed

It *is* possible to find a strategy for winning on every turn!

TYPICAL EARLY BEHAVIOUR

Player 1 sees:

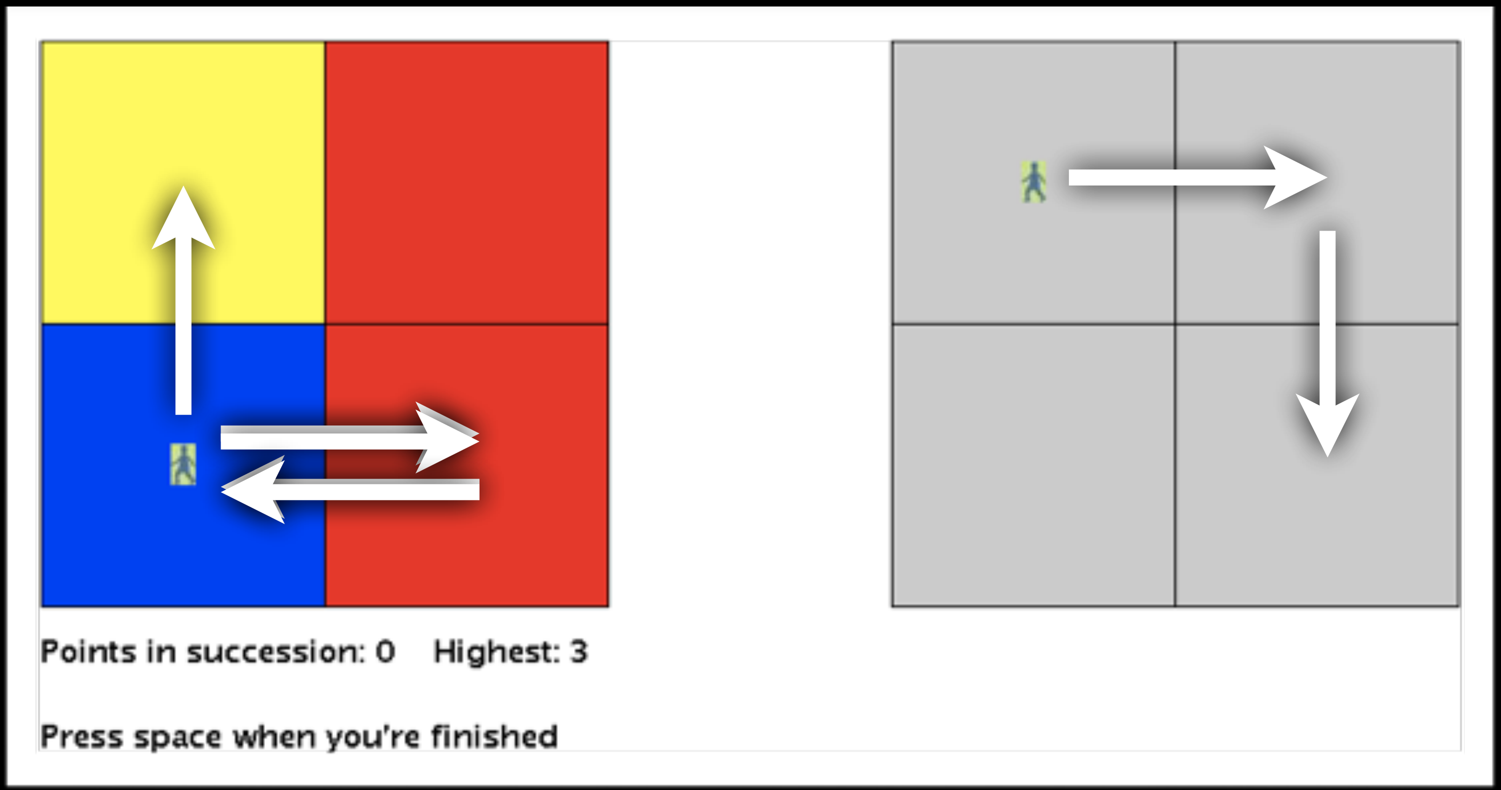

Player 2 sees:

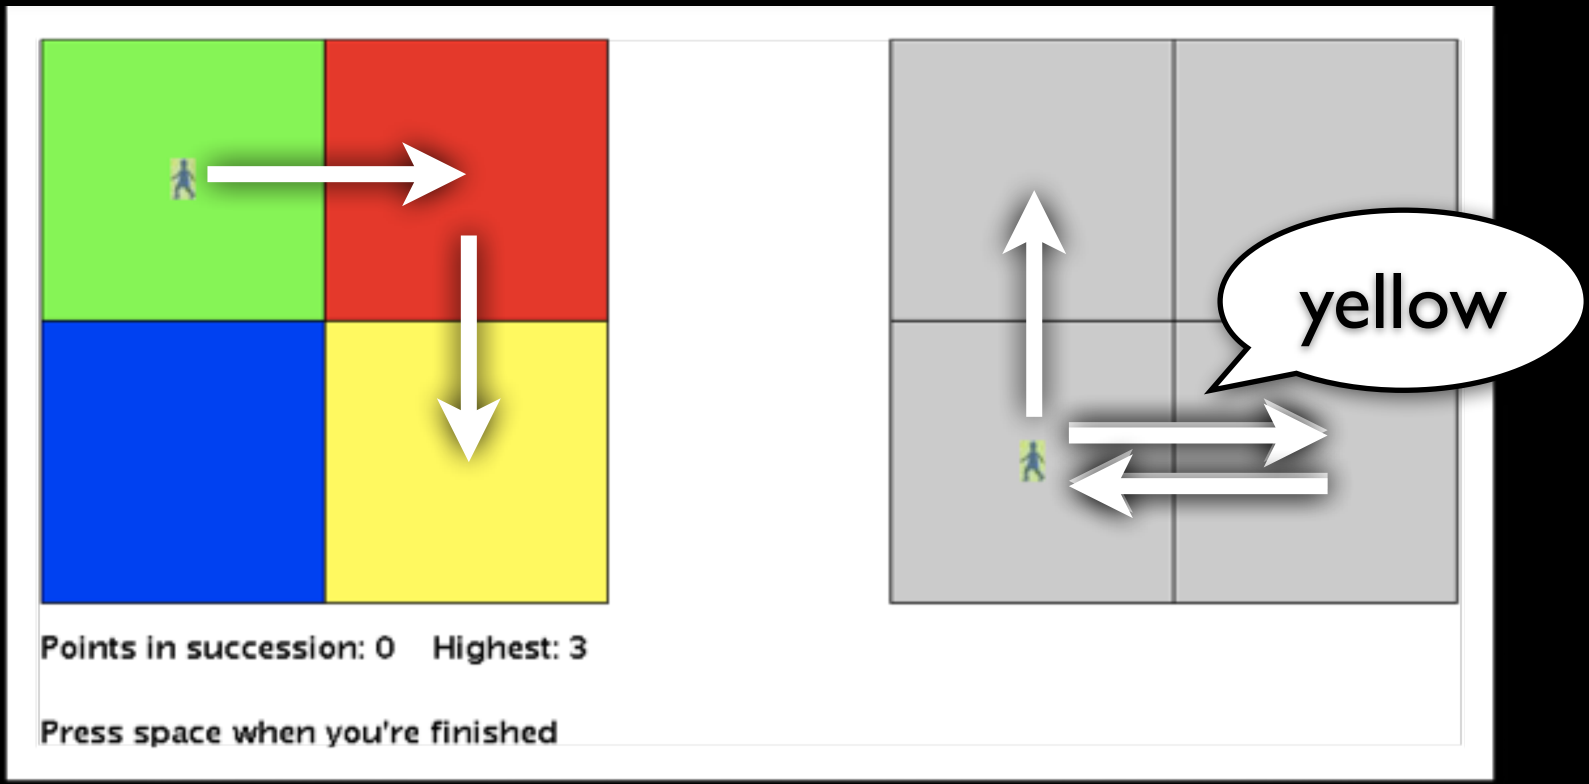

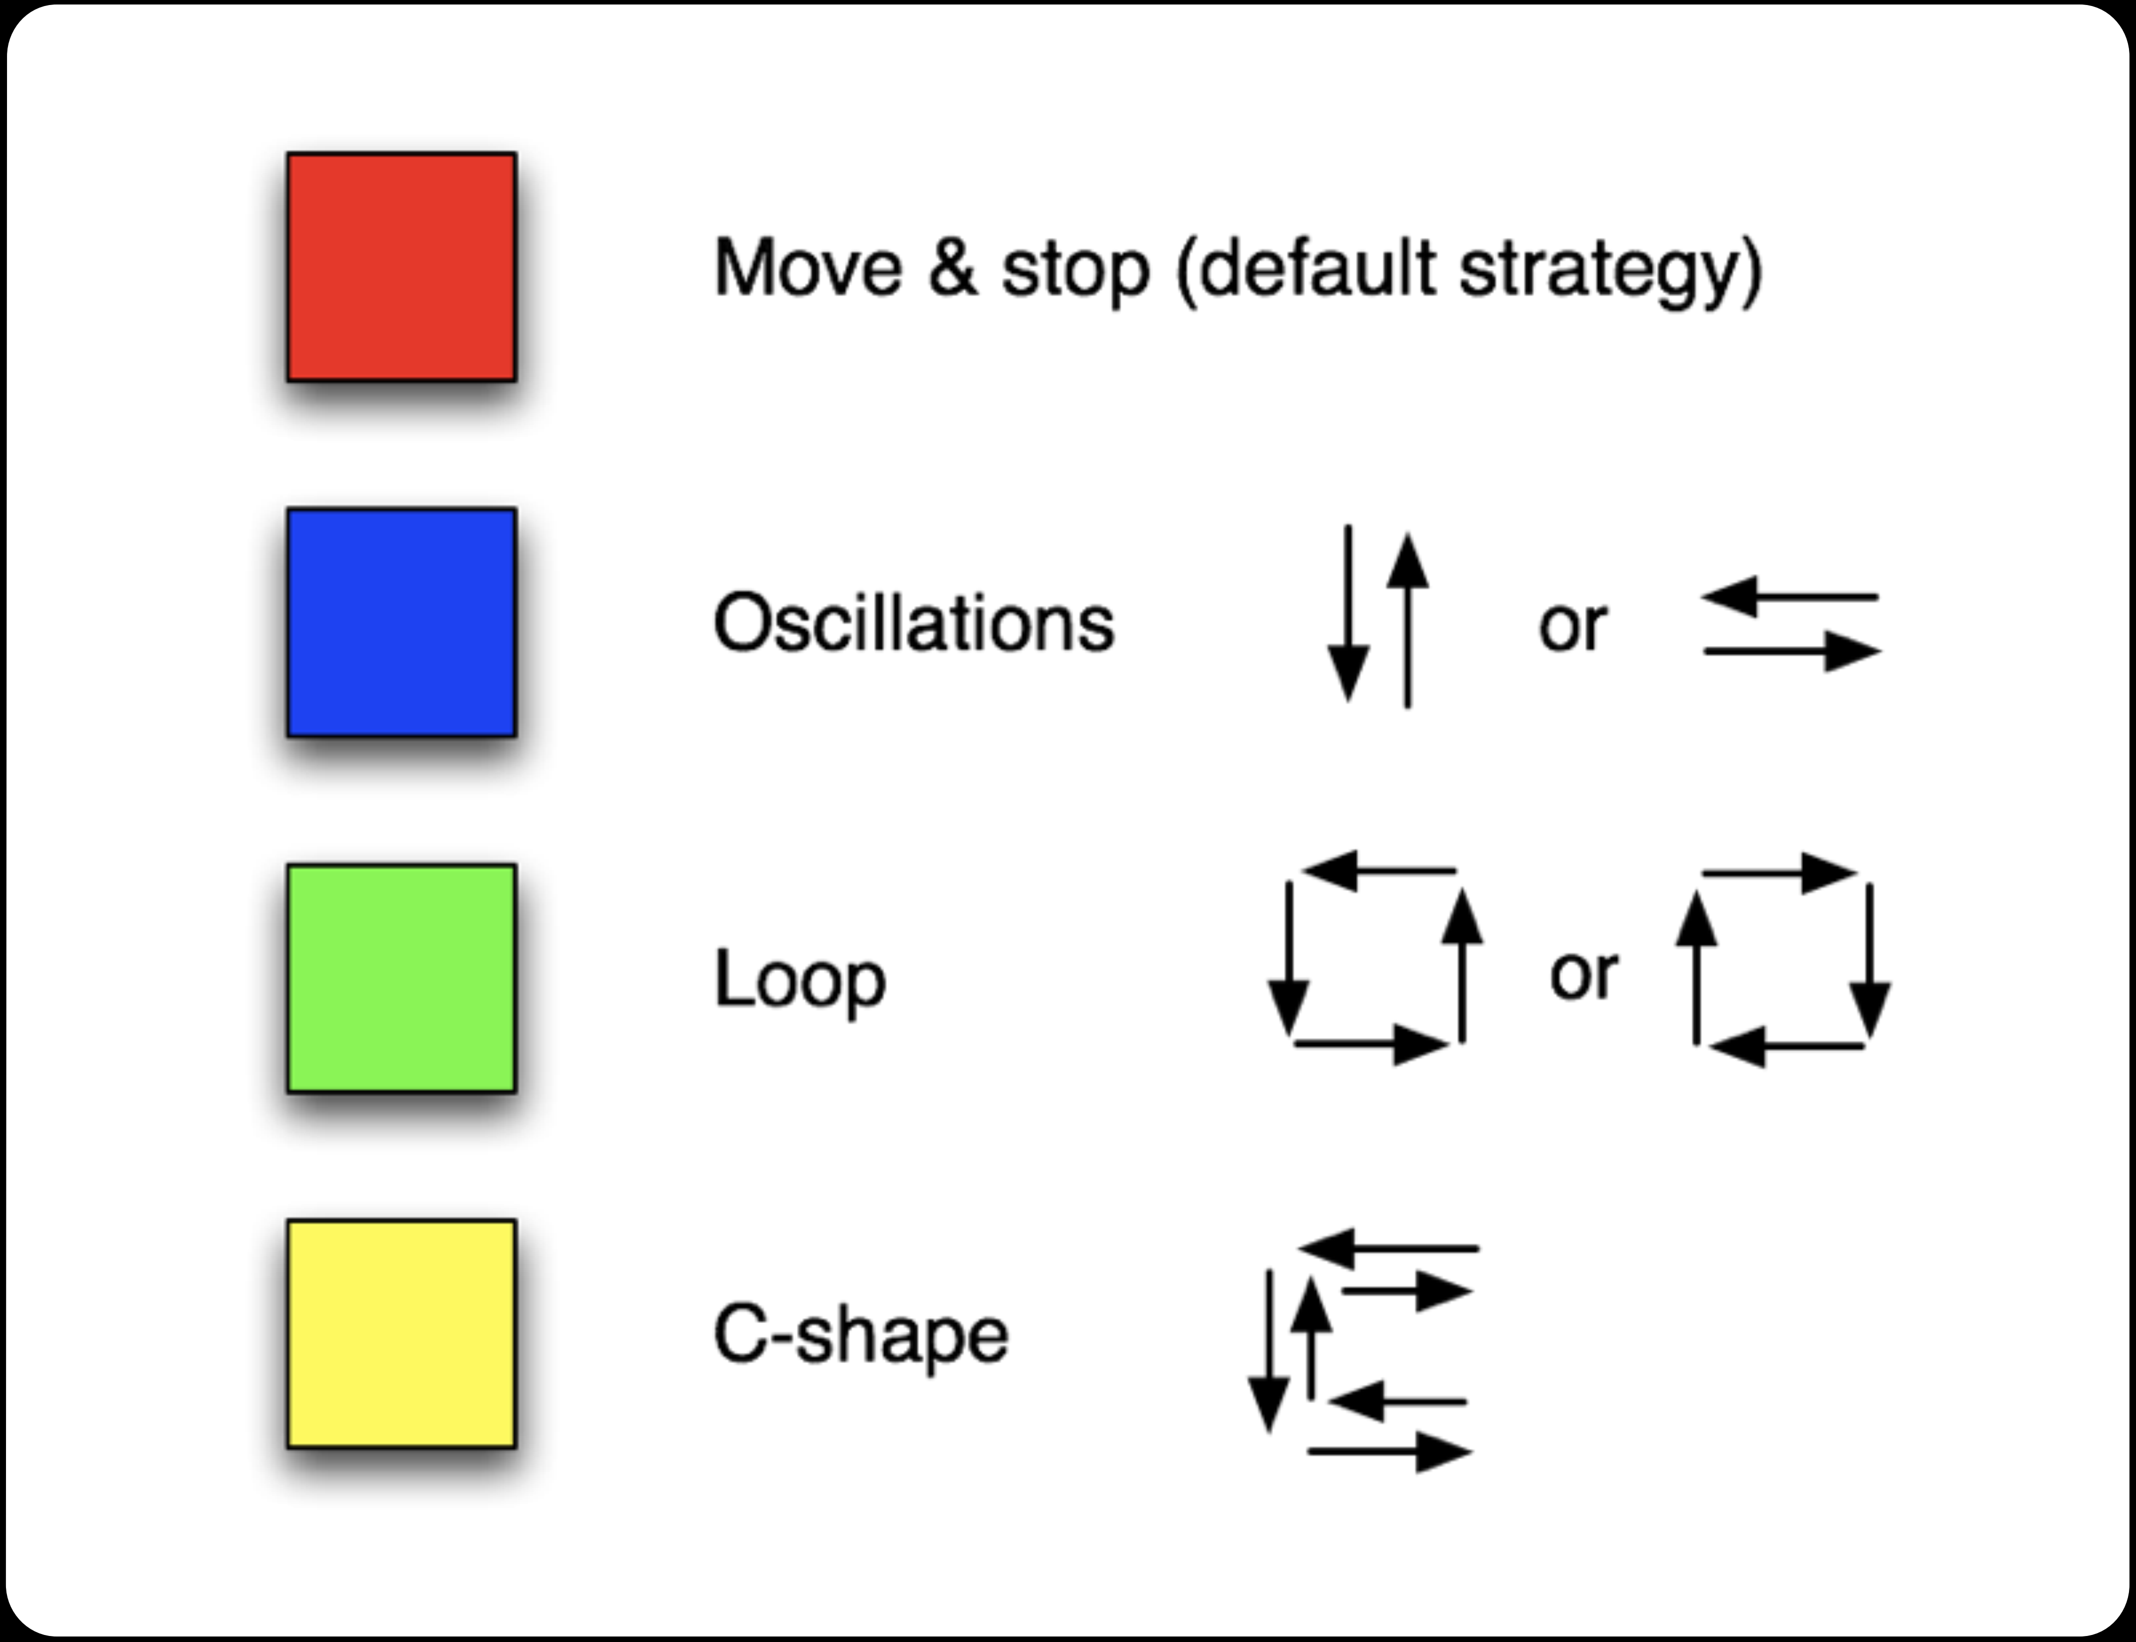
Move & stop (default strategy)
Oscillations
or
Loop
or
C-shape

AN EXAMPLE OF DIALOGUE

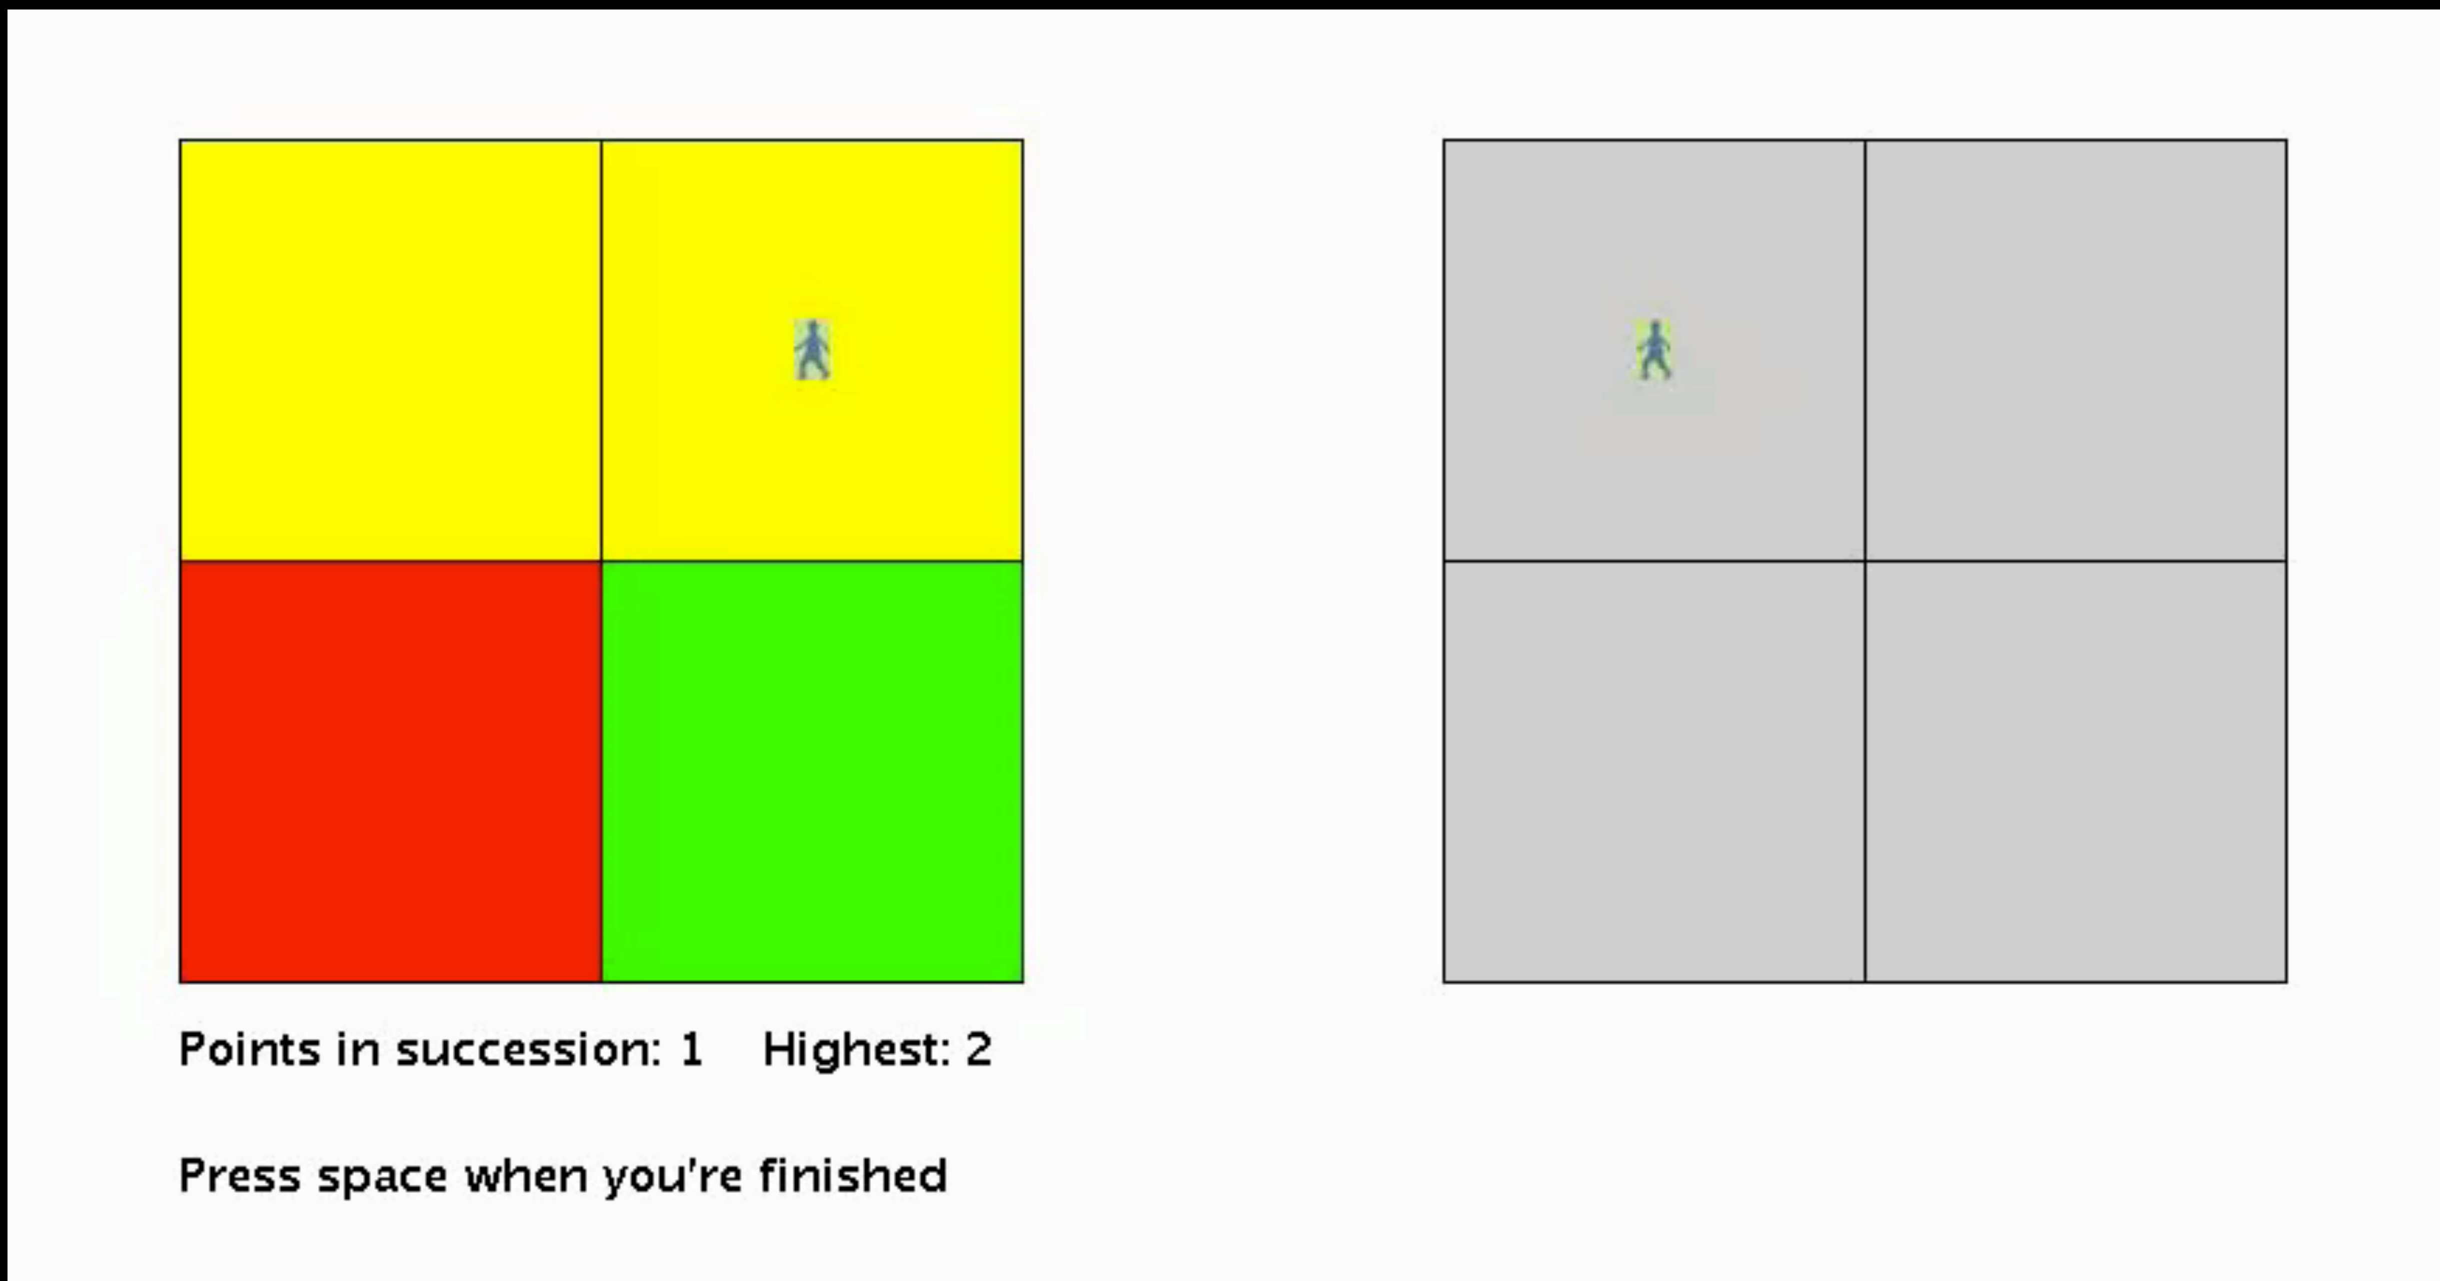

A TYPICAL PATTERN OF EMERGENCE

1. First a "default" strategy emerges
2. Then a signal to mean "something's wrong!"
3. Ritualised to mean a particular colour
4. Extended to the other colours

- Demonstrates again the fundamental importance of the socio/cultural process.
- Shows how embodied behaviours get exploited to carry meaning.

CULTURAL EVOLUTION

- Cultural evolution is just as important (if not more so) than biological evolution in understanding human language
  - This means we need to abandon some of the idealisations of the orthodox individual-based approach

    Language structure does not spring directly from cognitive constraints/biases

- We can study cultural evolution in the lab by building experiments based on simulation

- **Results suggest that we might think of culture itself as a computational system.**

PART 4:
AN OVERVIEW OF LOTS OF WORK IN PROGRESS.
OR, WHAT DO WE DO ALL DAY IN MY LAB?

STUDY 3: ARTIFICIAL SIGN LANGUAGE

- **Training**: learn gestures that communicate manner and path of a moving ball (12 out of 16 meanings)
- **Testing**: try to communicate all 16 meanings using gesture
- **Initial input**: improvised gestures of 16 different participants

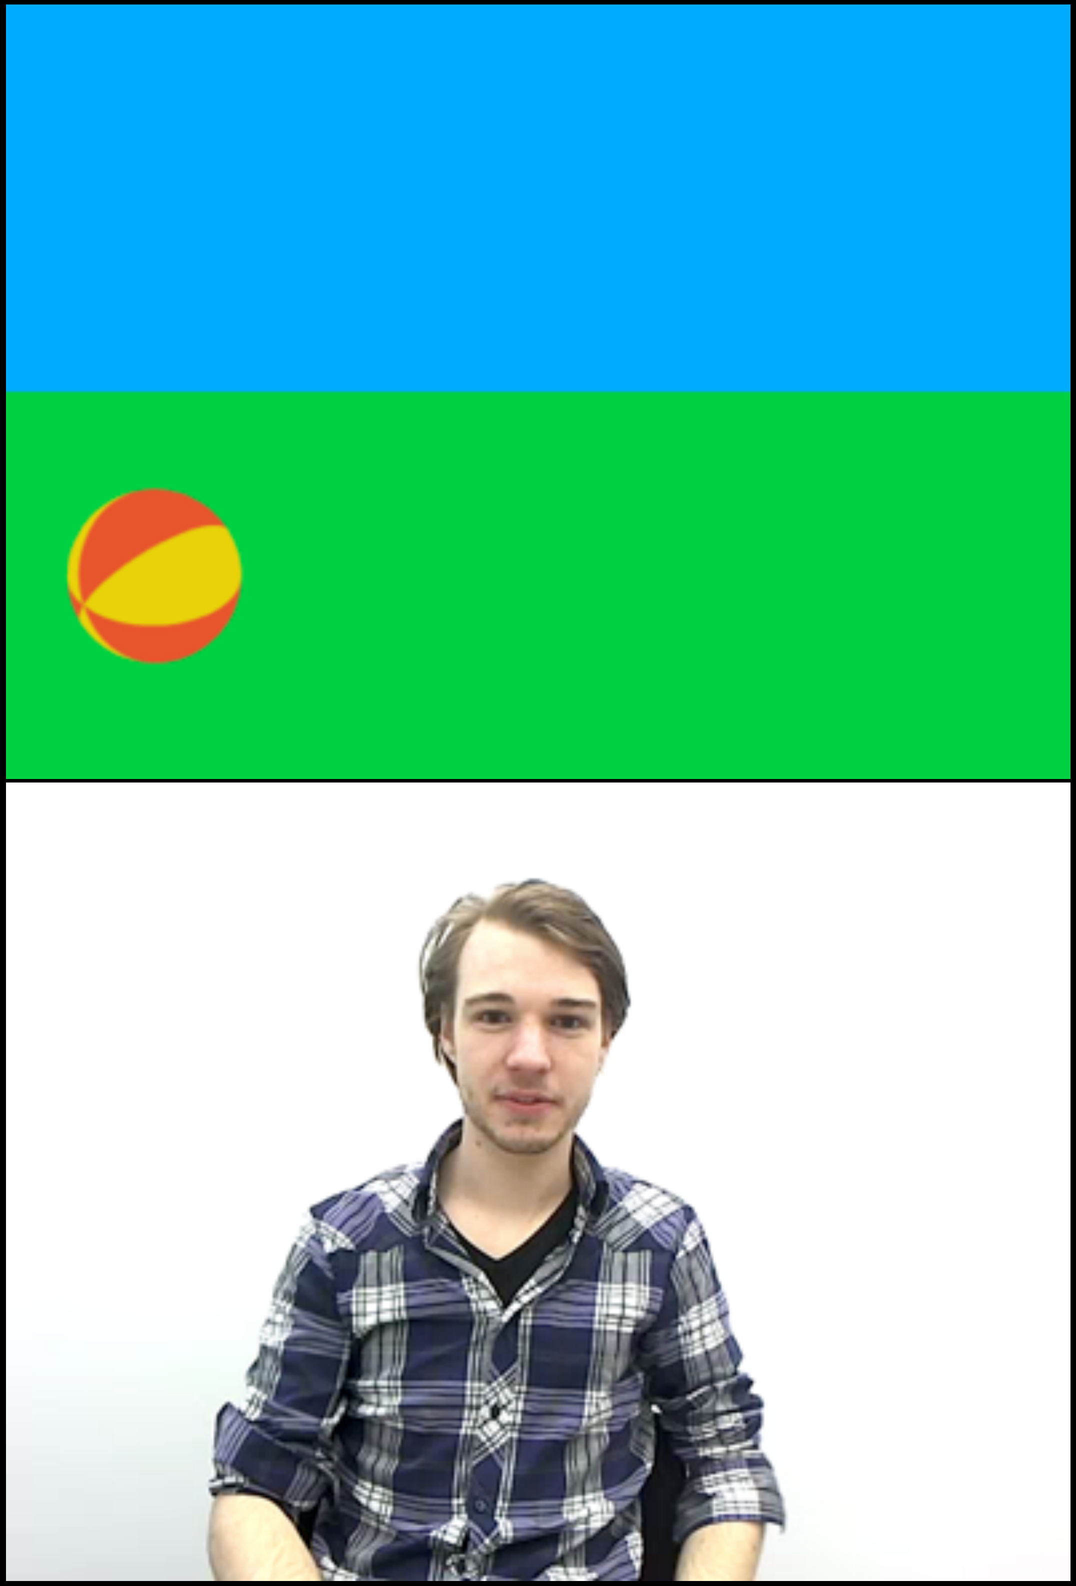

Smith, Abramova, Cartmill & Kirby (in prep)

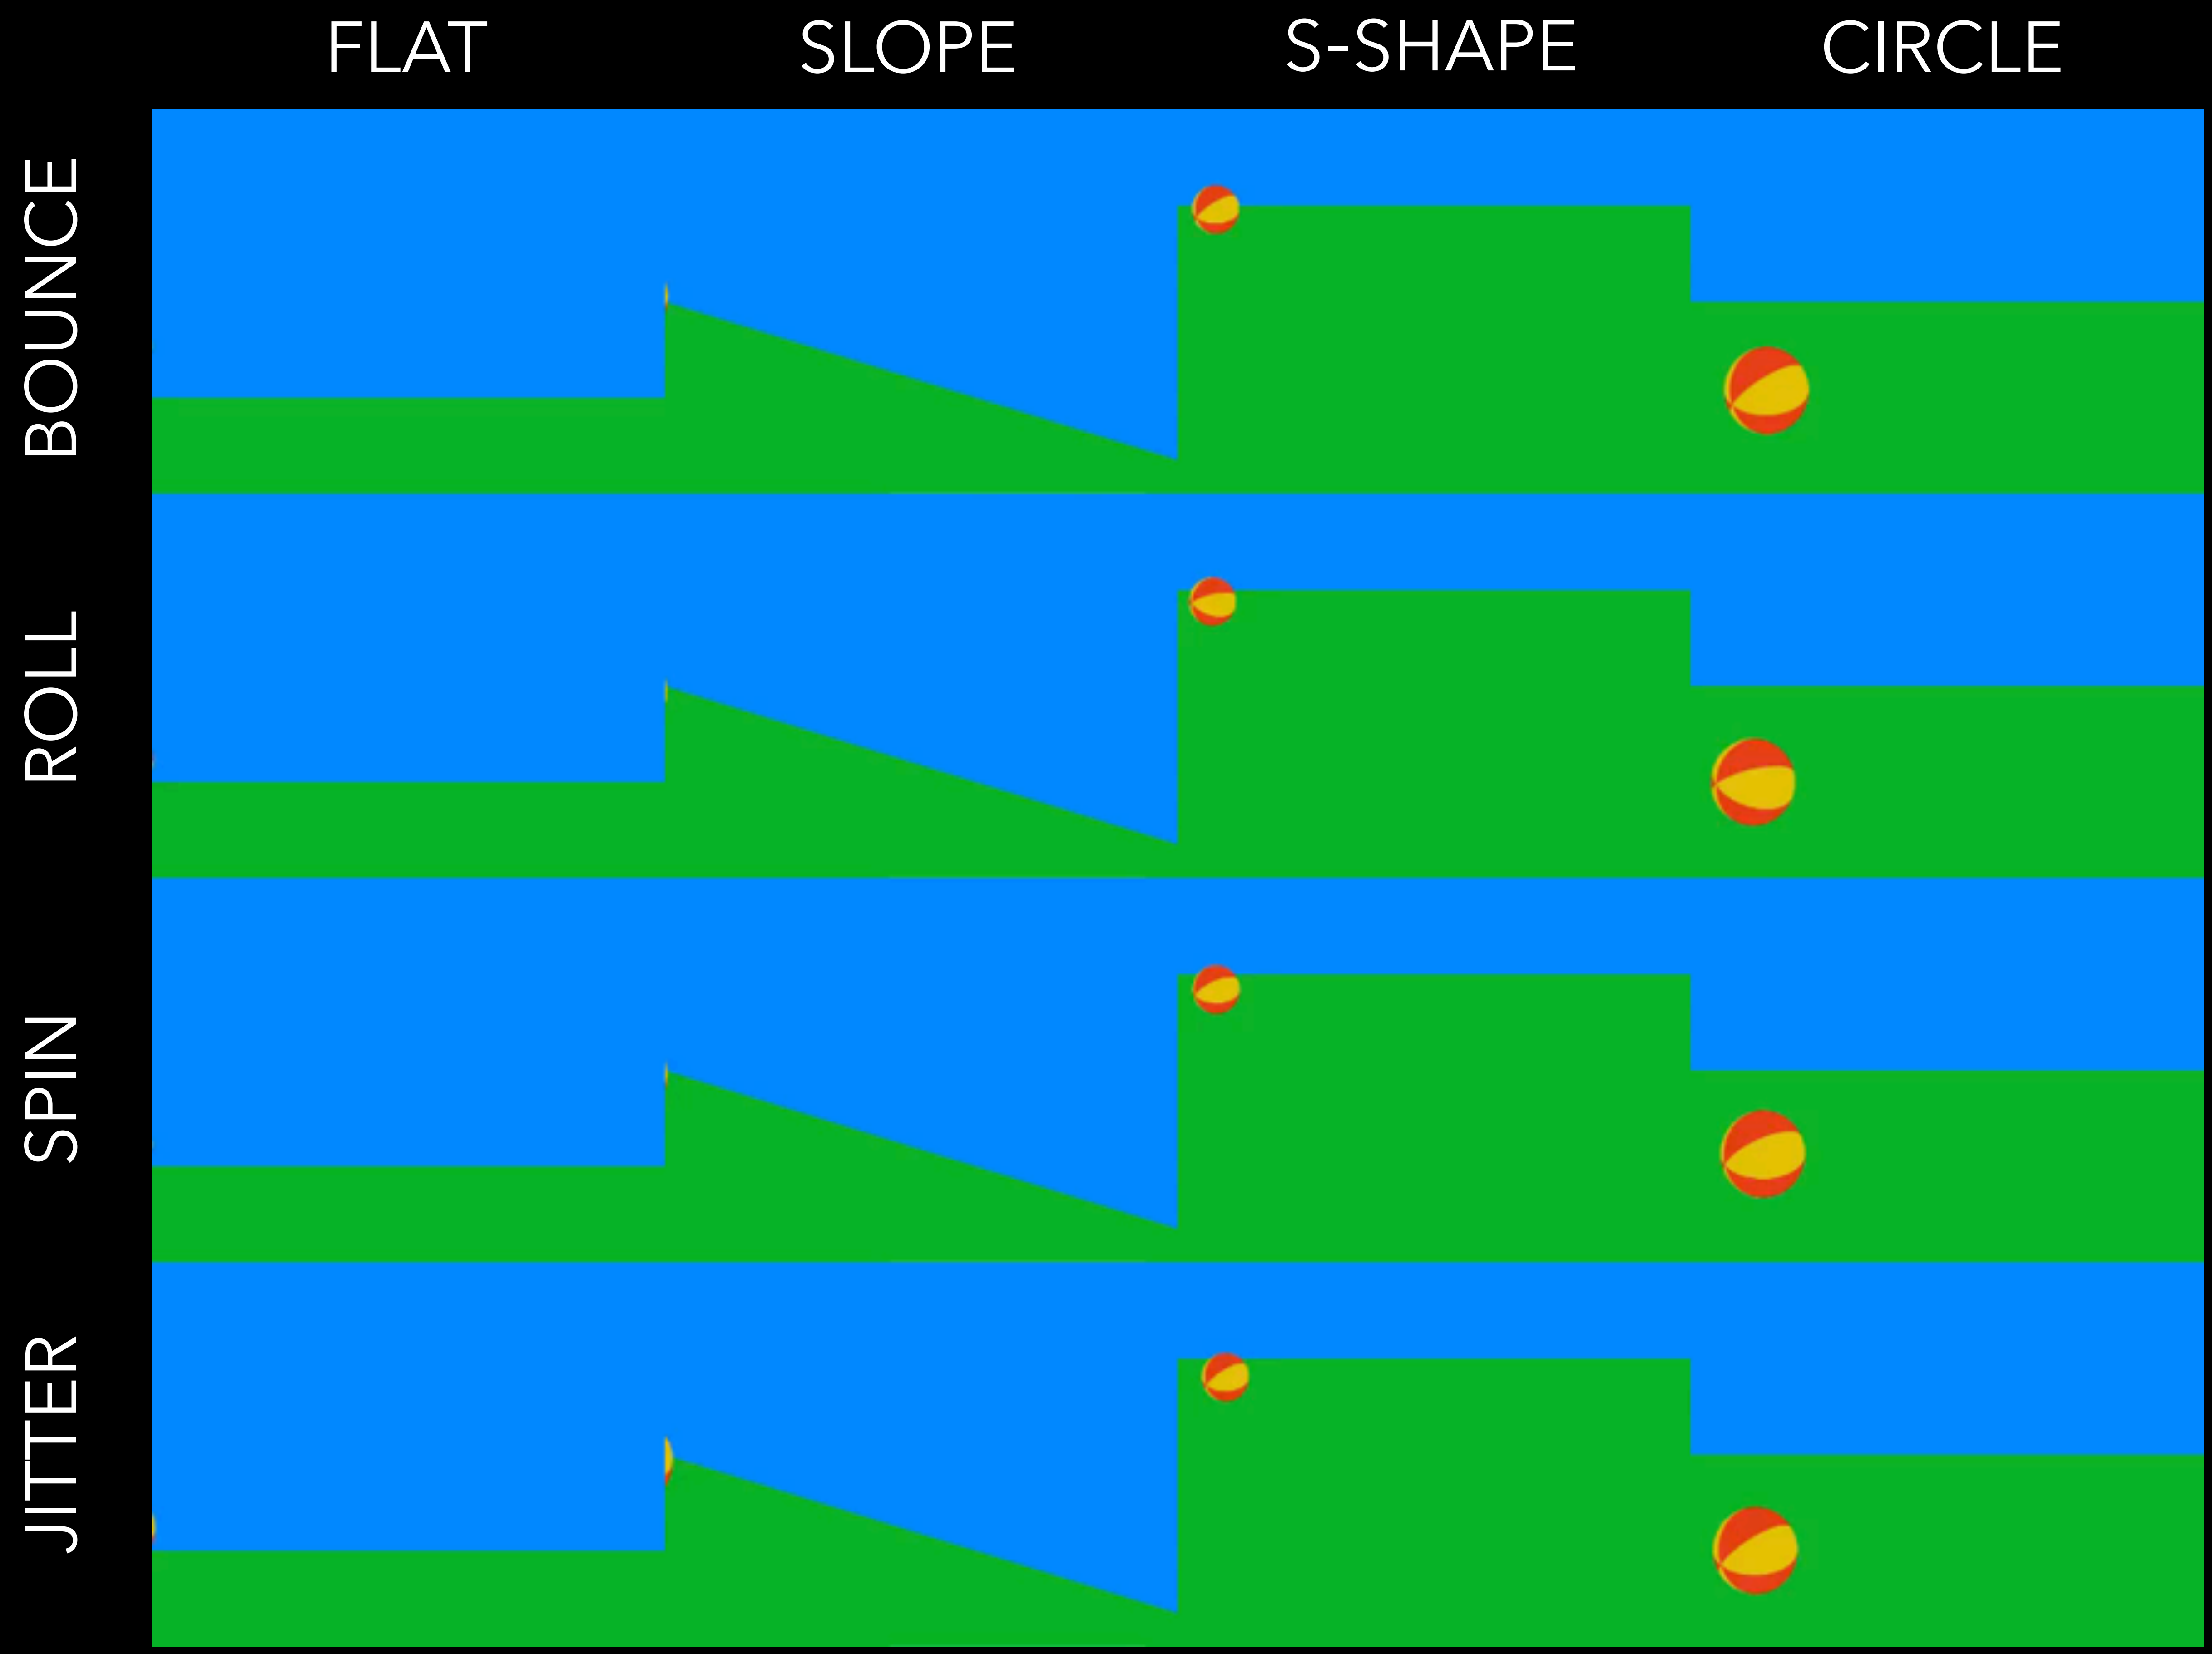
FLAT
SLOPE
S-SHAPE
CIRCLE
BOUNCE
ROLL
SPIN
JITTER


GRAND

STUDY 3: QUANTITATIVE RESULTS, LANGUAGE BECOMES MORE SYSTEMATIC

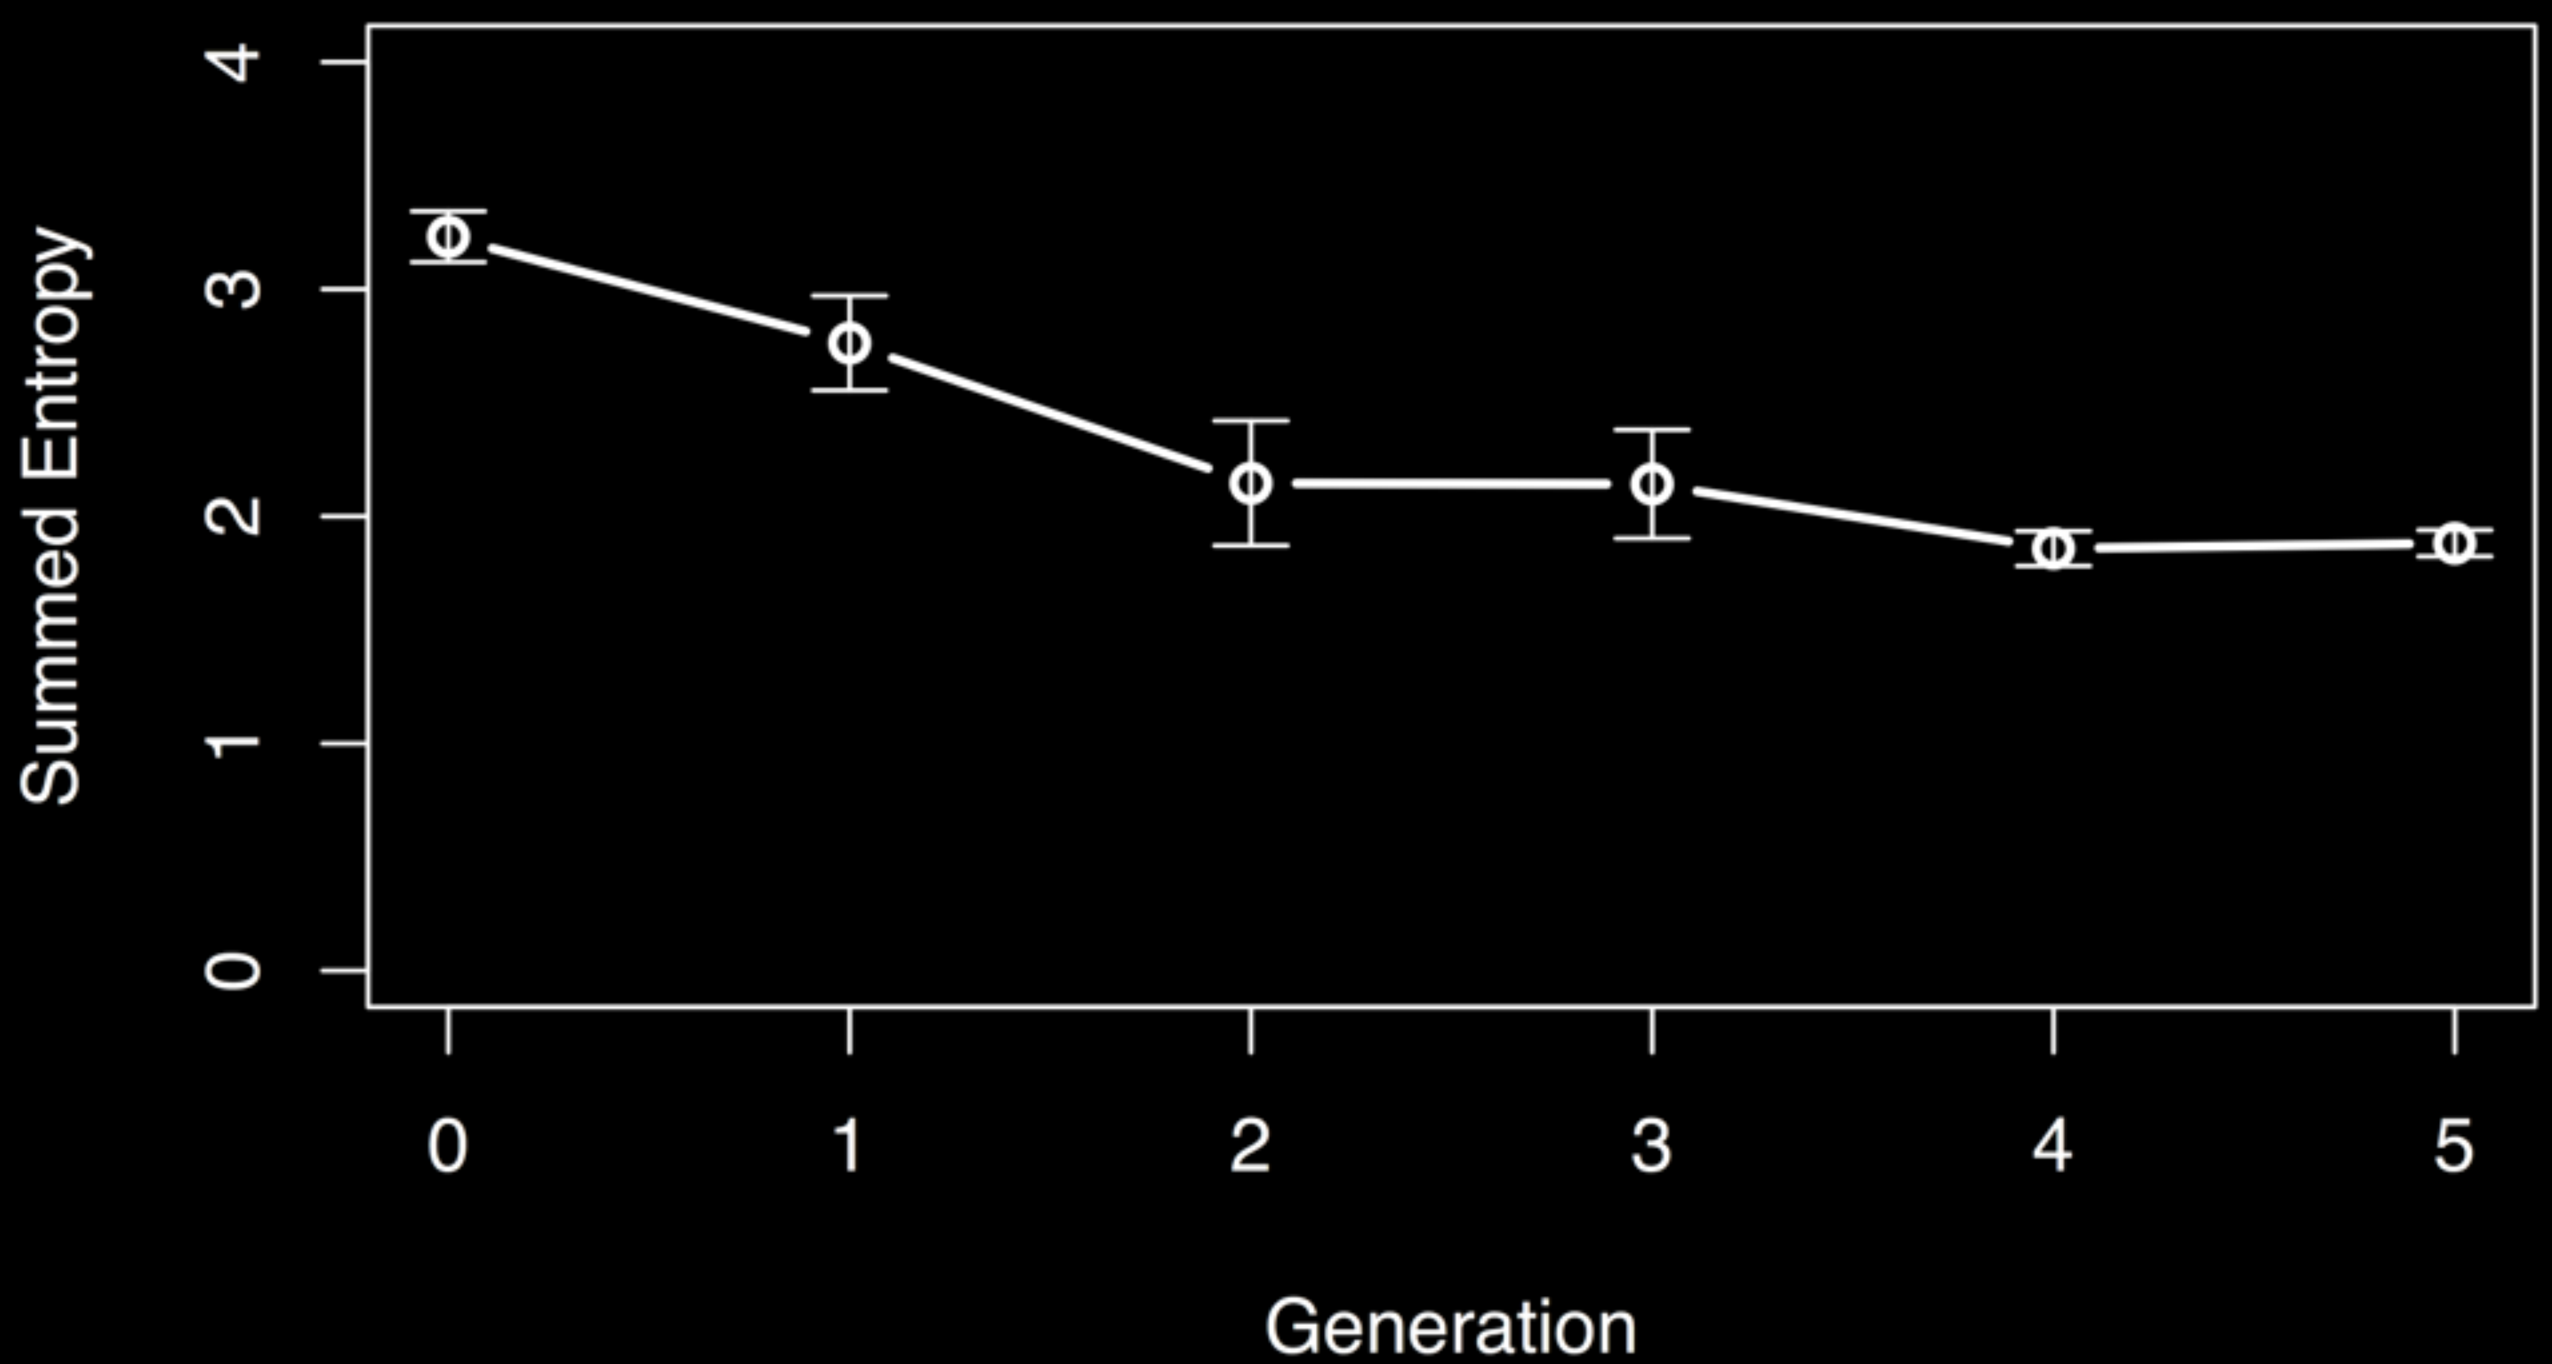

STUDY 4: SIMON GAME
THE EMERGENCE OF A SYSTEM

- **Training + testing**: observe sequences of coloured lights and immediately imitate them

  Each trial presented as an independent immediate imitation task.

- **Initial input**: 60 random sequences of 12 lights with equal probability of each light occurring

Cornish, Smith & Kirby (2013)


Picking
sequence...

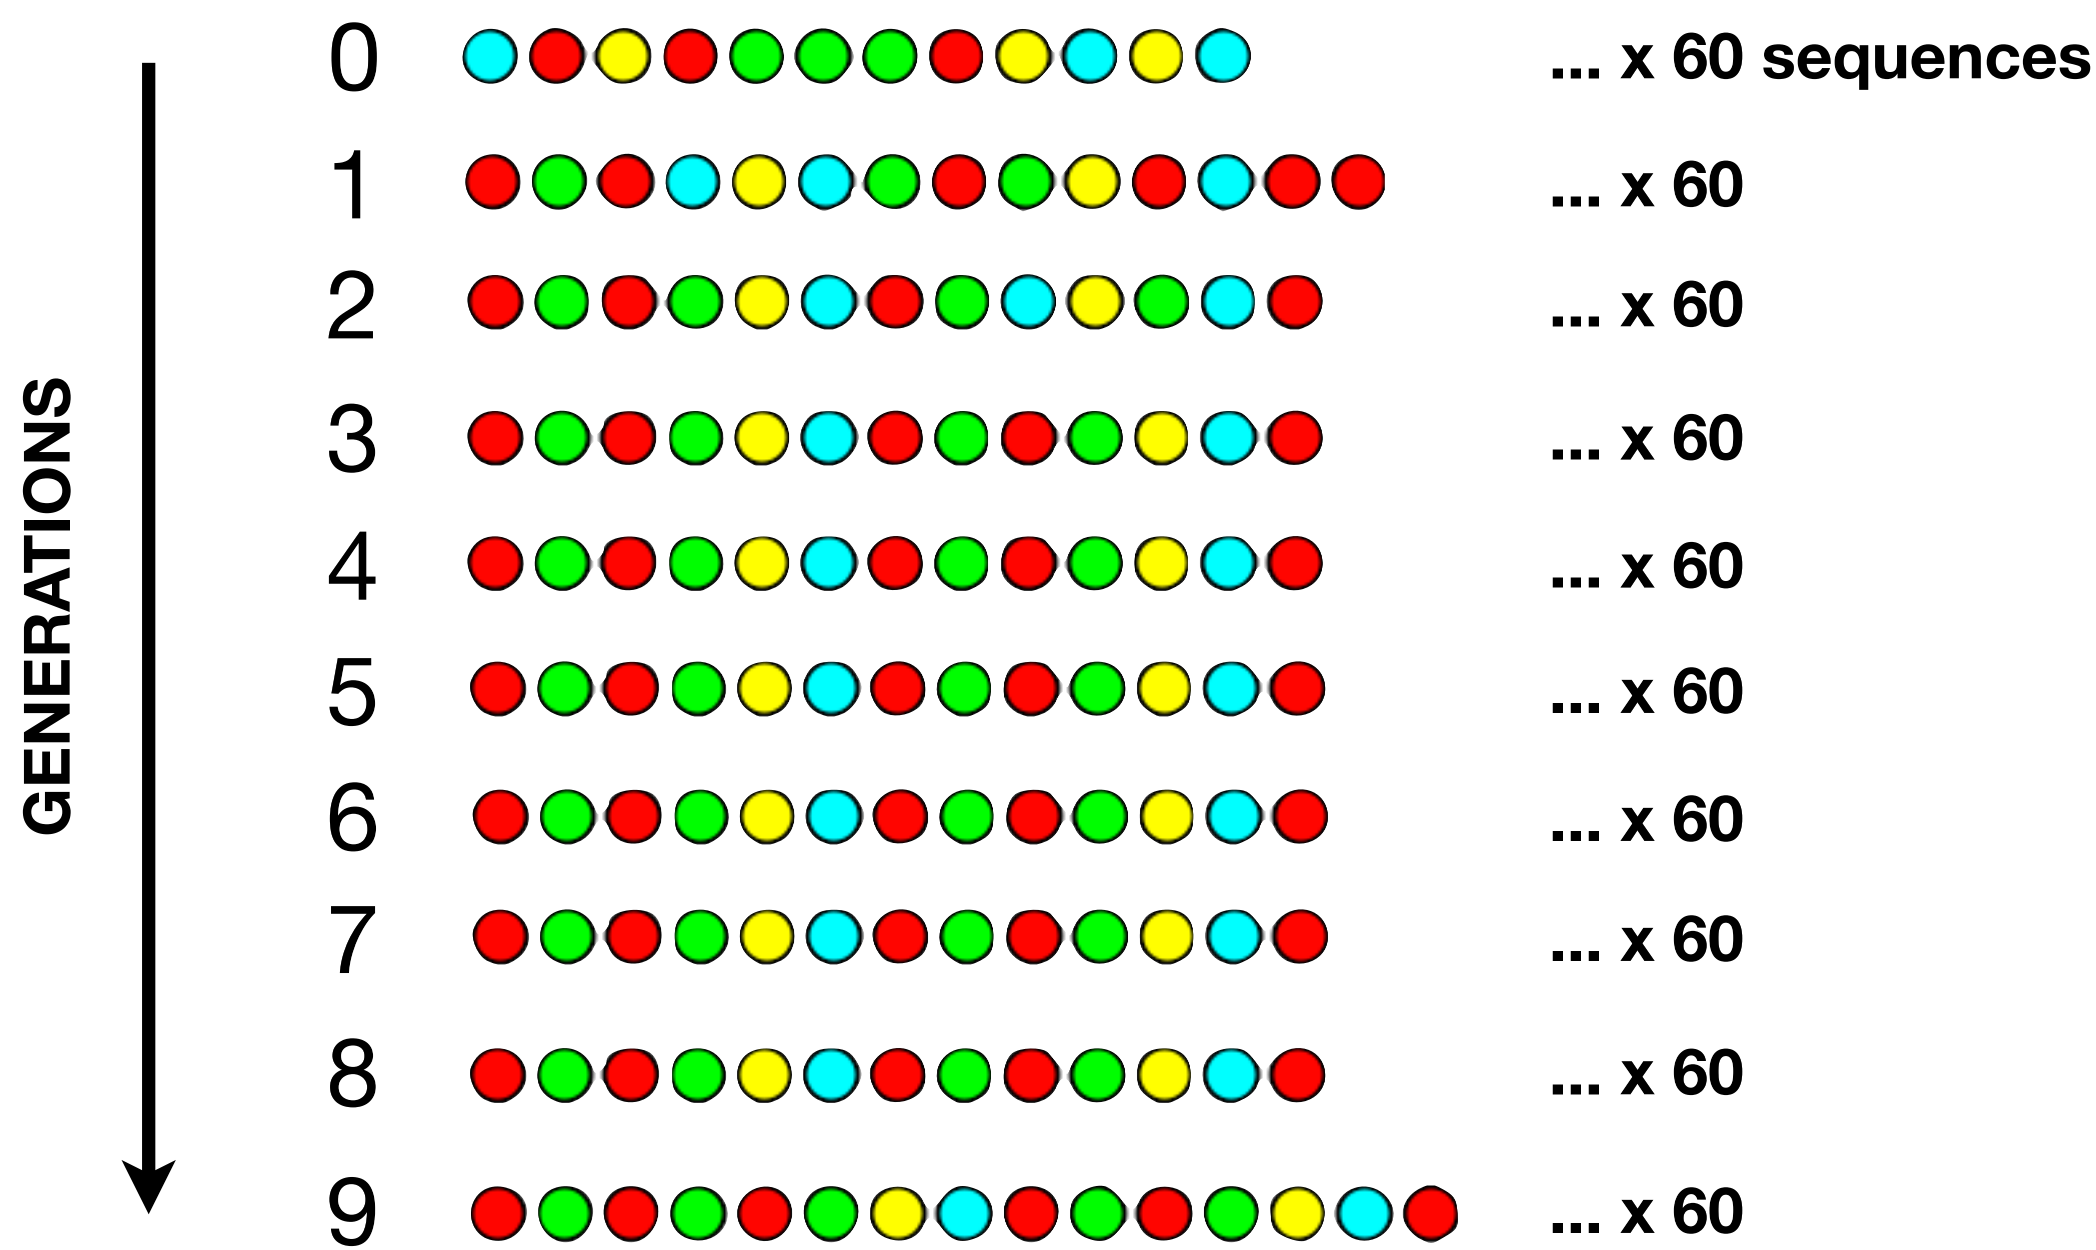

GENERATIONS
0
1
2
3
4
5
6
7
8
9
... x 60 sequences
... x 60
... x 60
... x 60
... x 60
... x 60
... x 60
... x 60
... x 60
... x 60

STUDY 4: RESULTS
TASK GETS EASIER OVER GENERATIONS

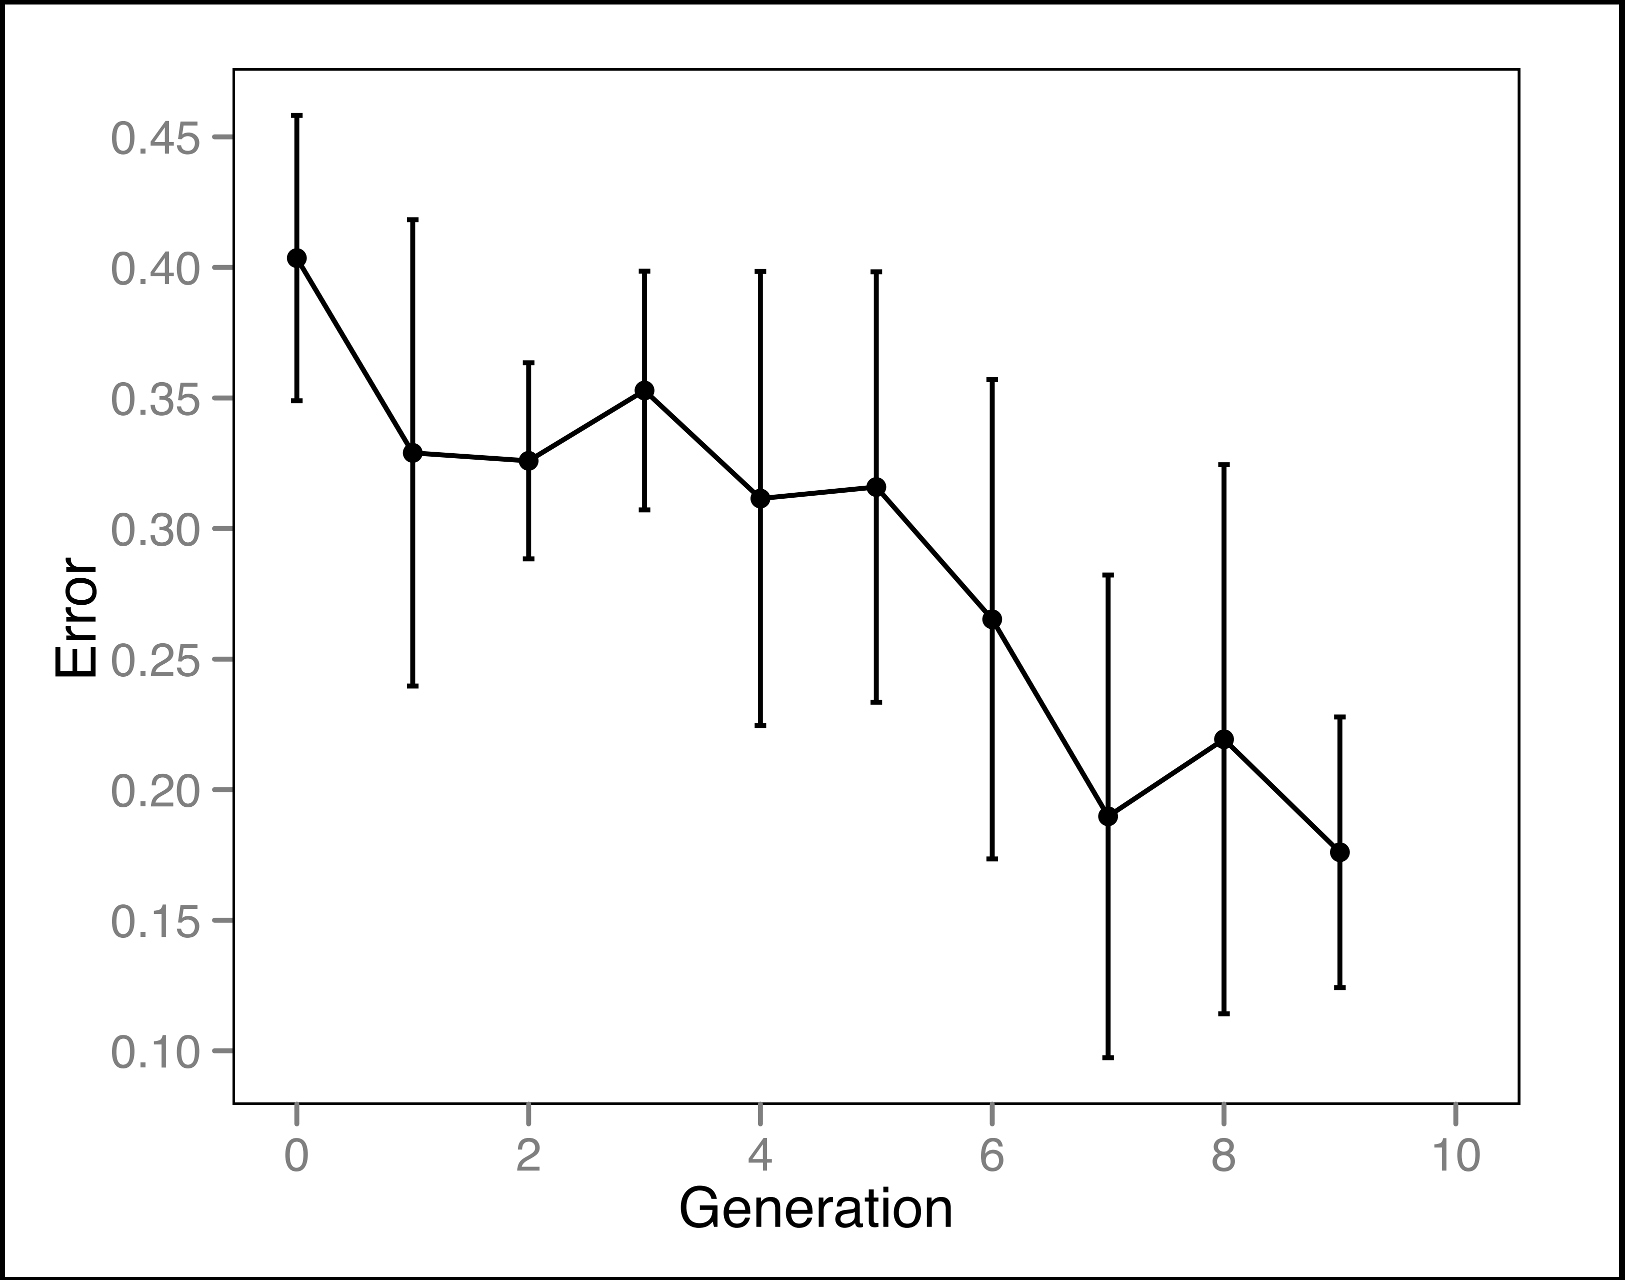

STUDY 4: RESULTS
SET BECOMES STRUCTURED

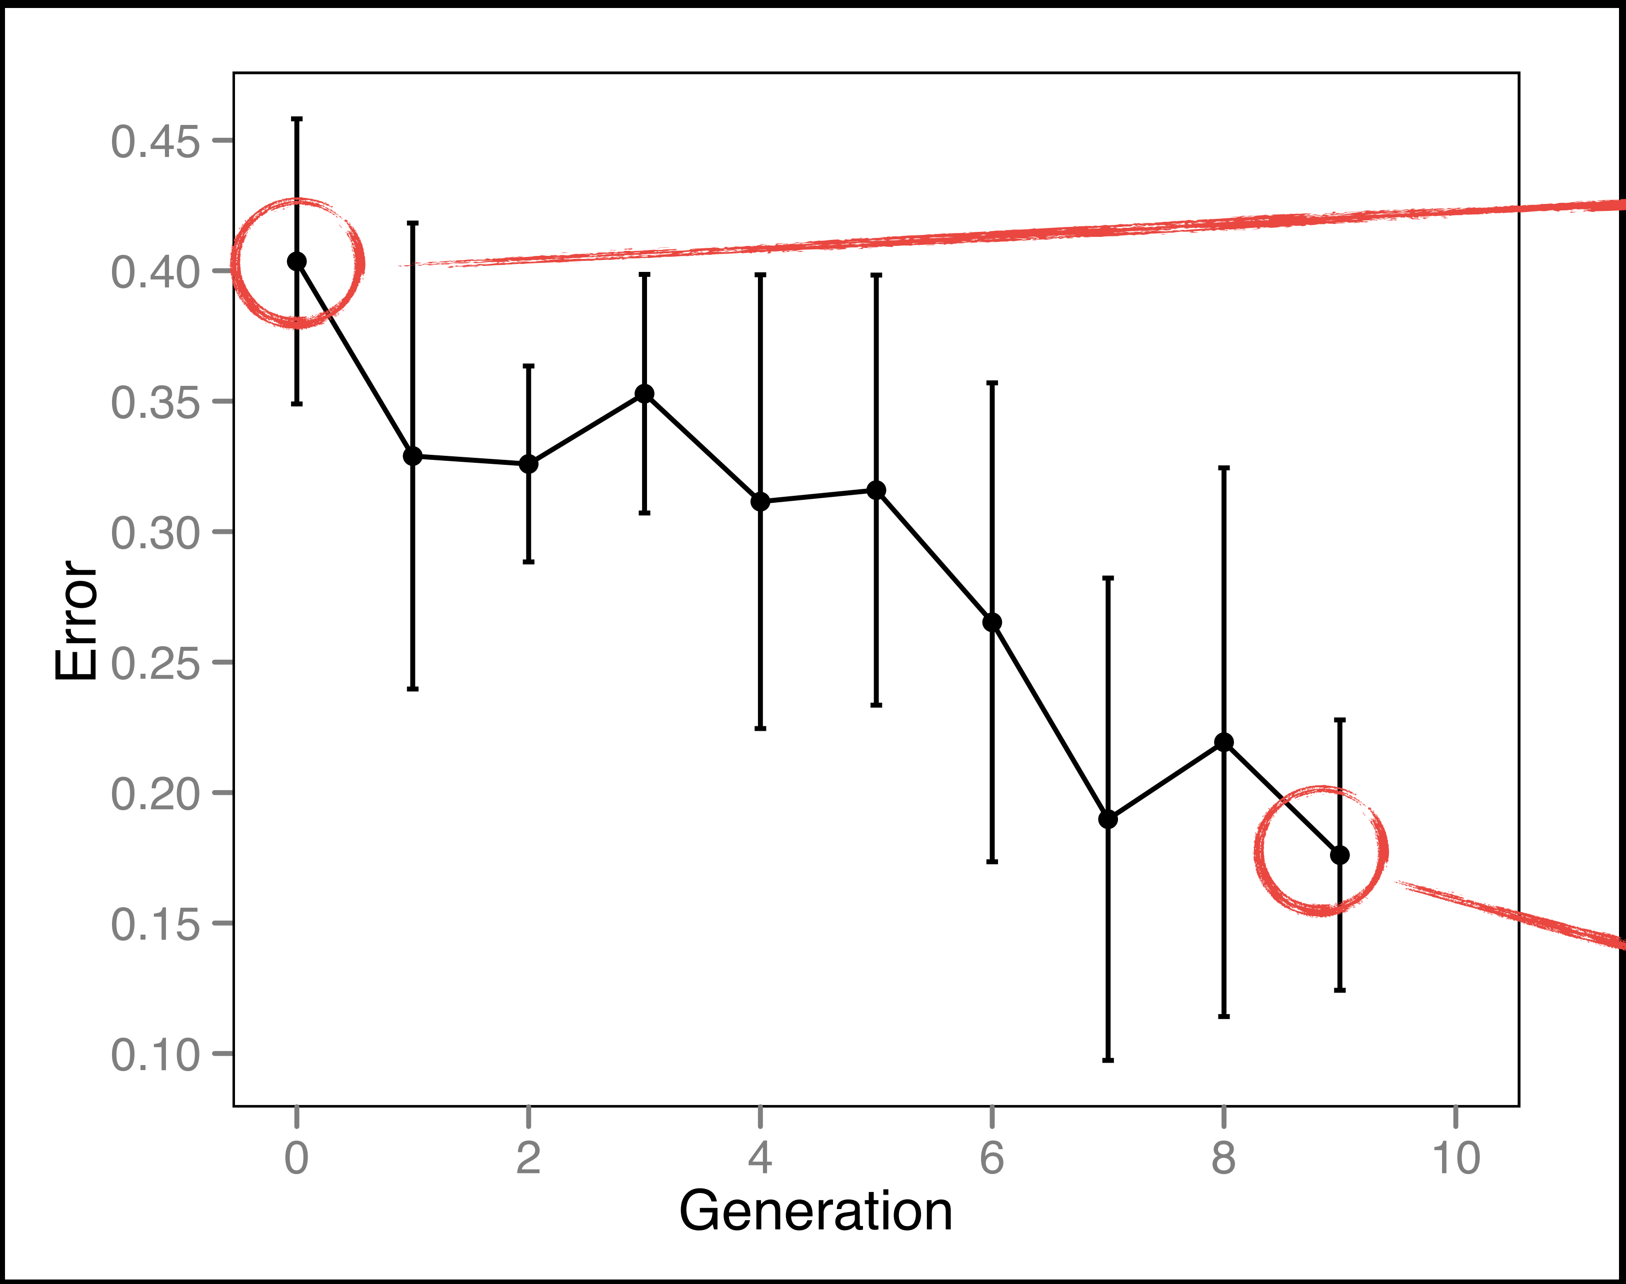

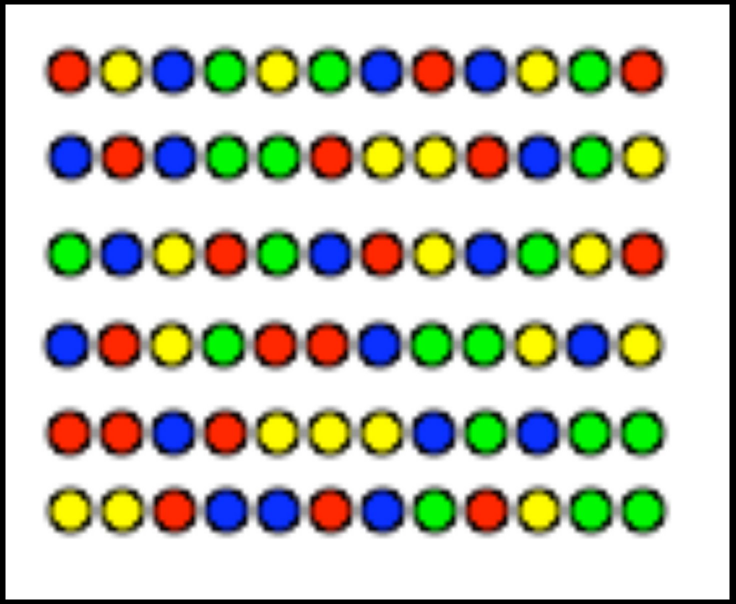

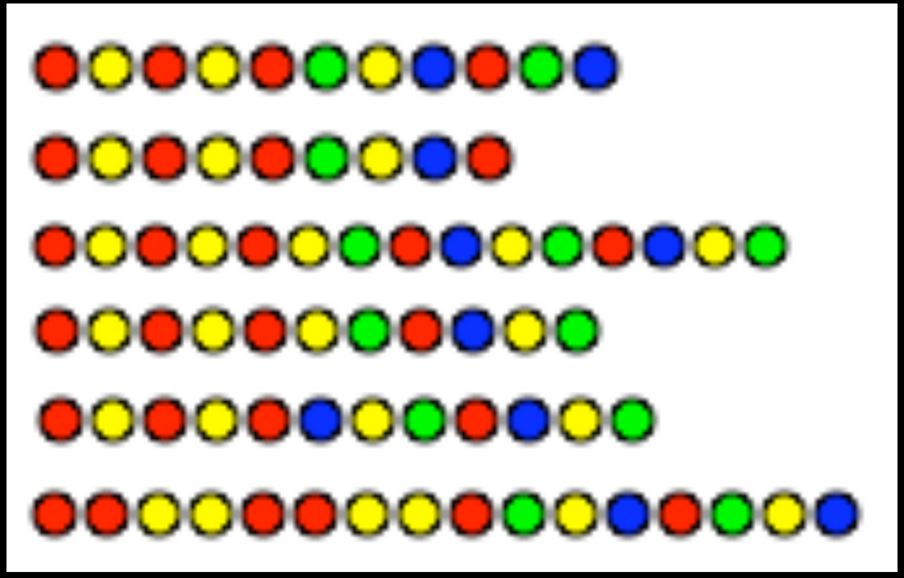

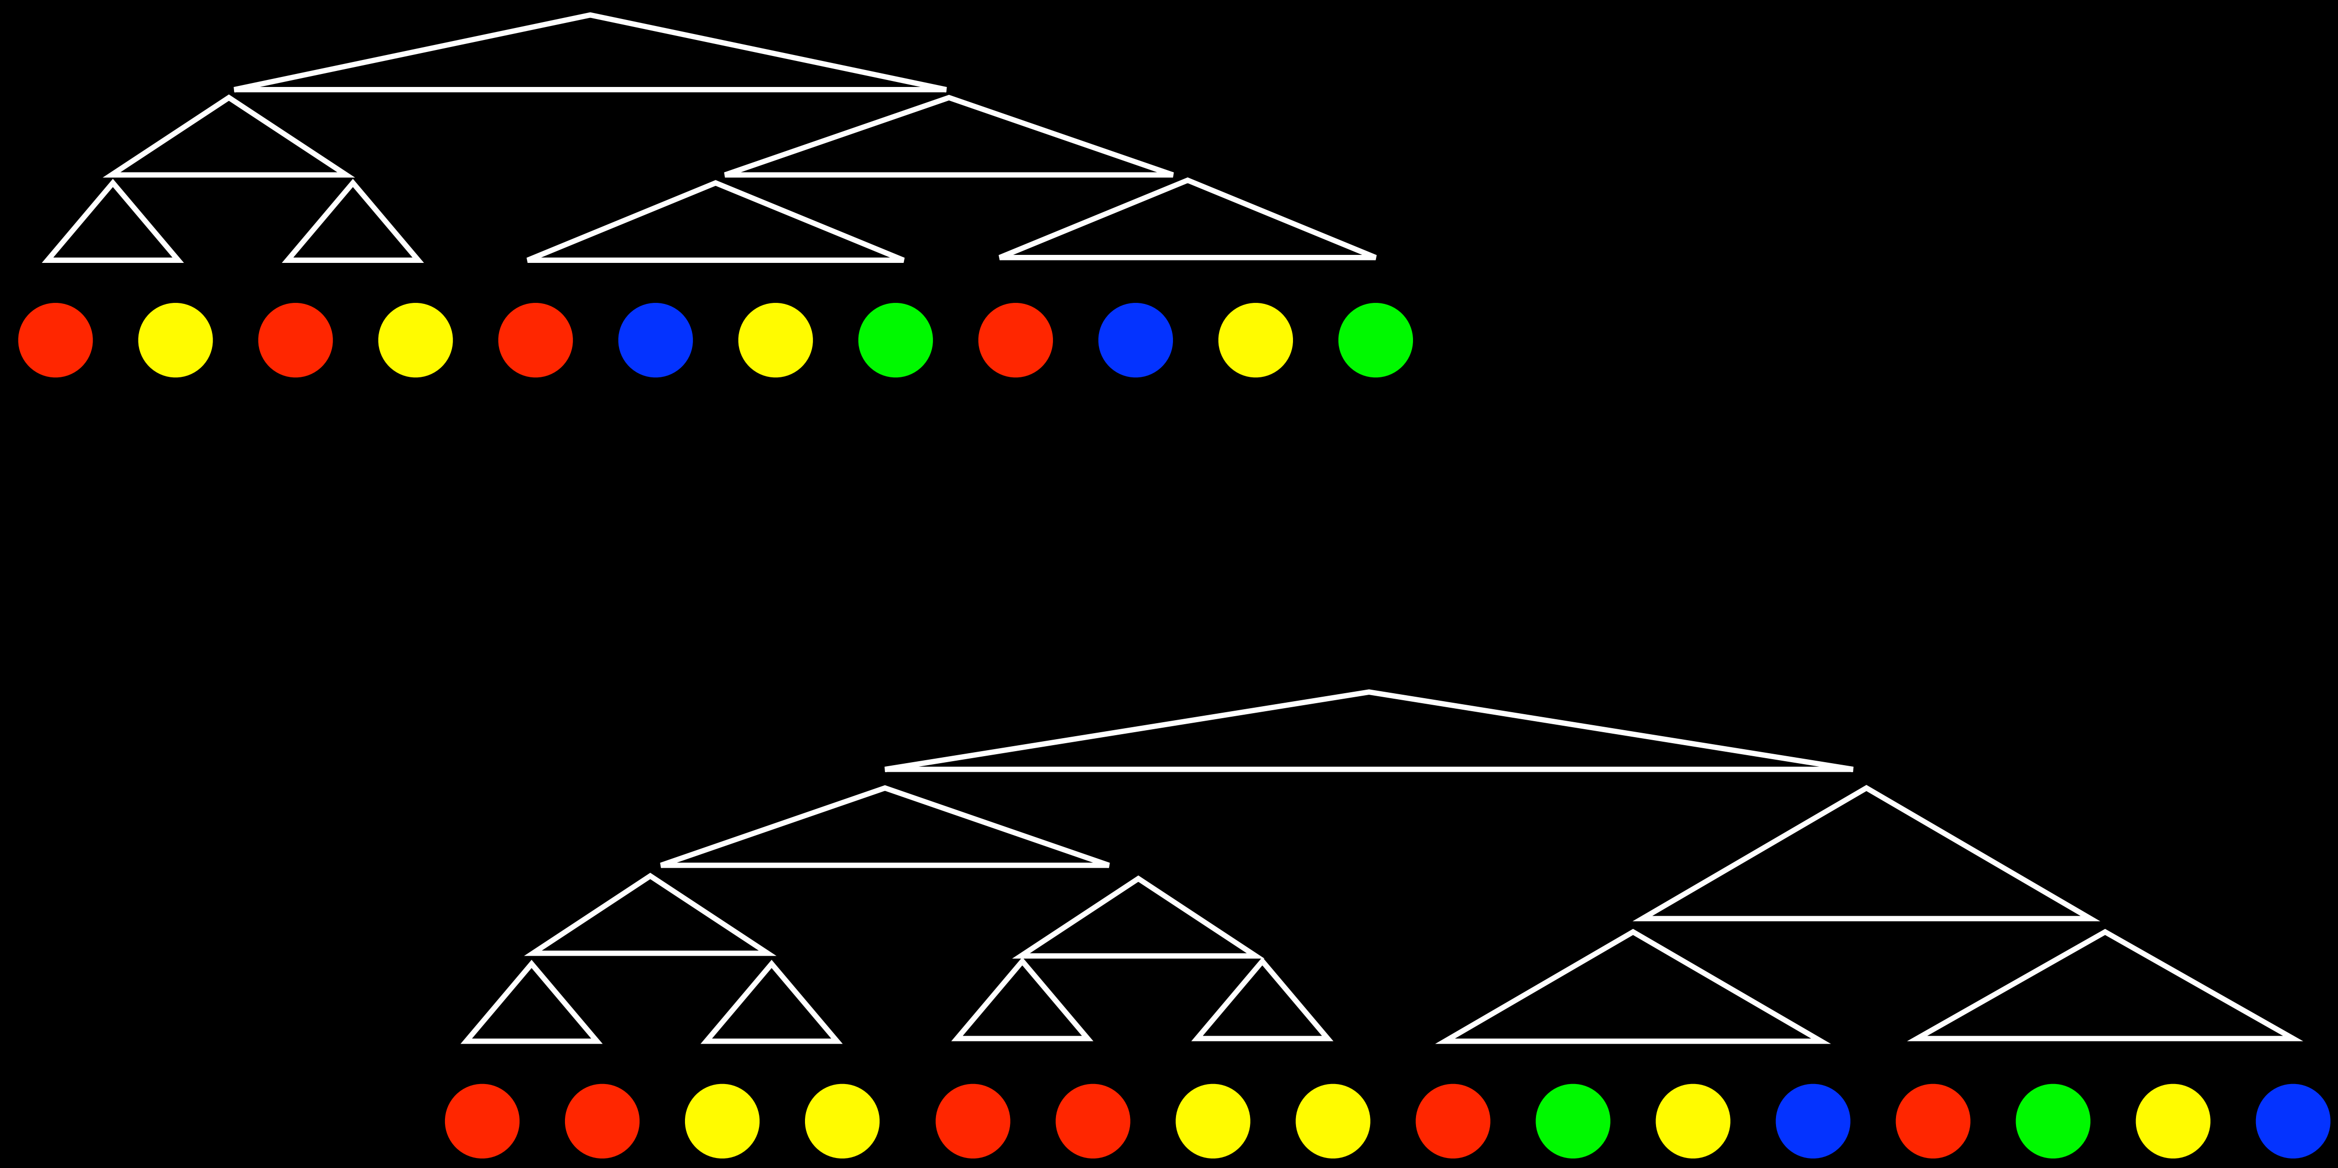

STUDY 4: RESULTS
LEARNABILITY OF SHUFFLED VS. NON-SHUFFLED SETS

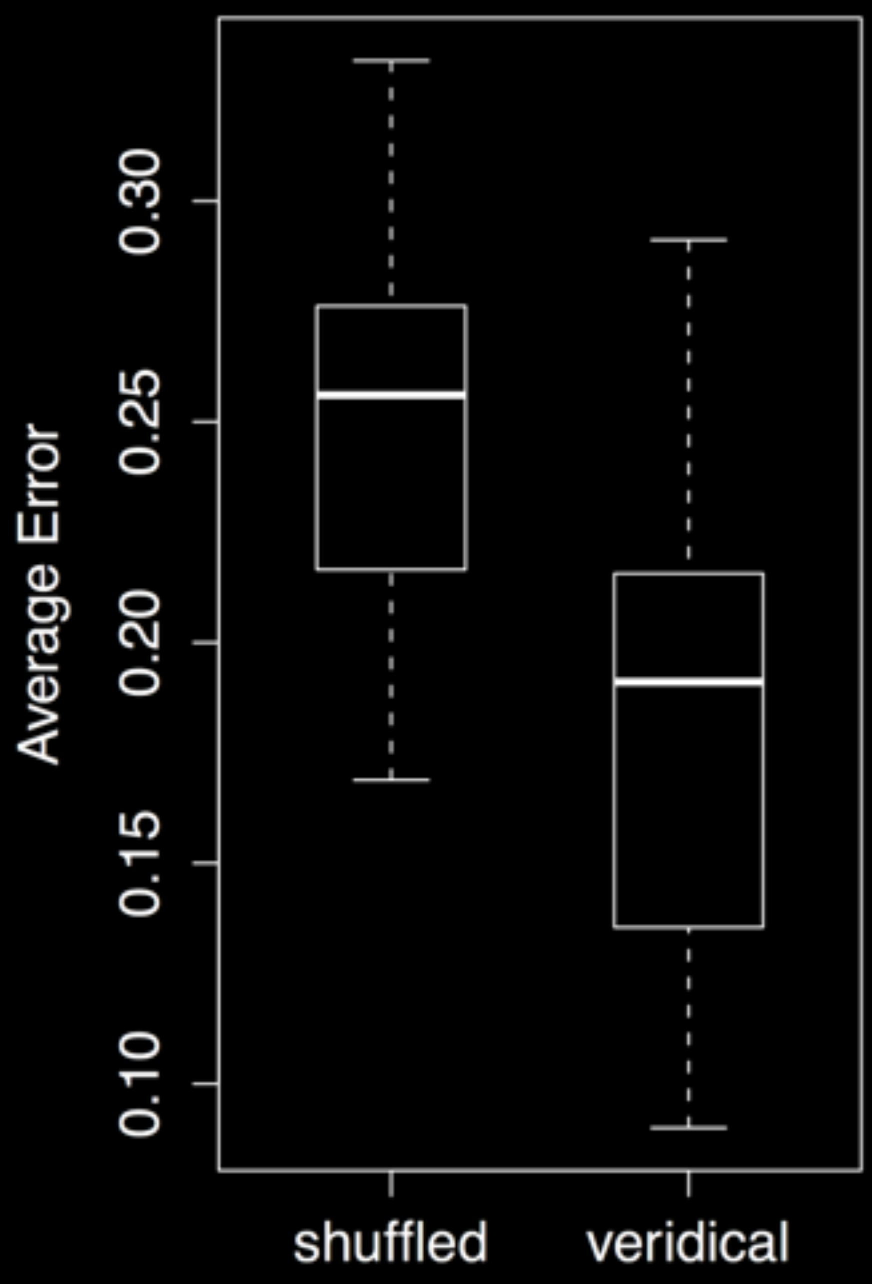

- Tested new participants on sequences from the end of our experiments
- Either exposed to sequences all from same chain, or from a mix of different chains
- Sequences are easier to copy if they are presented alongside others from the same chain
- Demonstrates that sequences now act as a system

STUDY 5: ITERATED LEARNING IN BABOONS

- **Training + testing**: observe pattern of four lights illuminated on touch screen and immediately recall the location of lights.
- **Initial input**: 50 random grid patterns

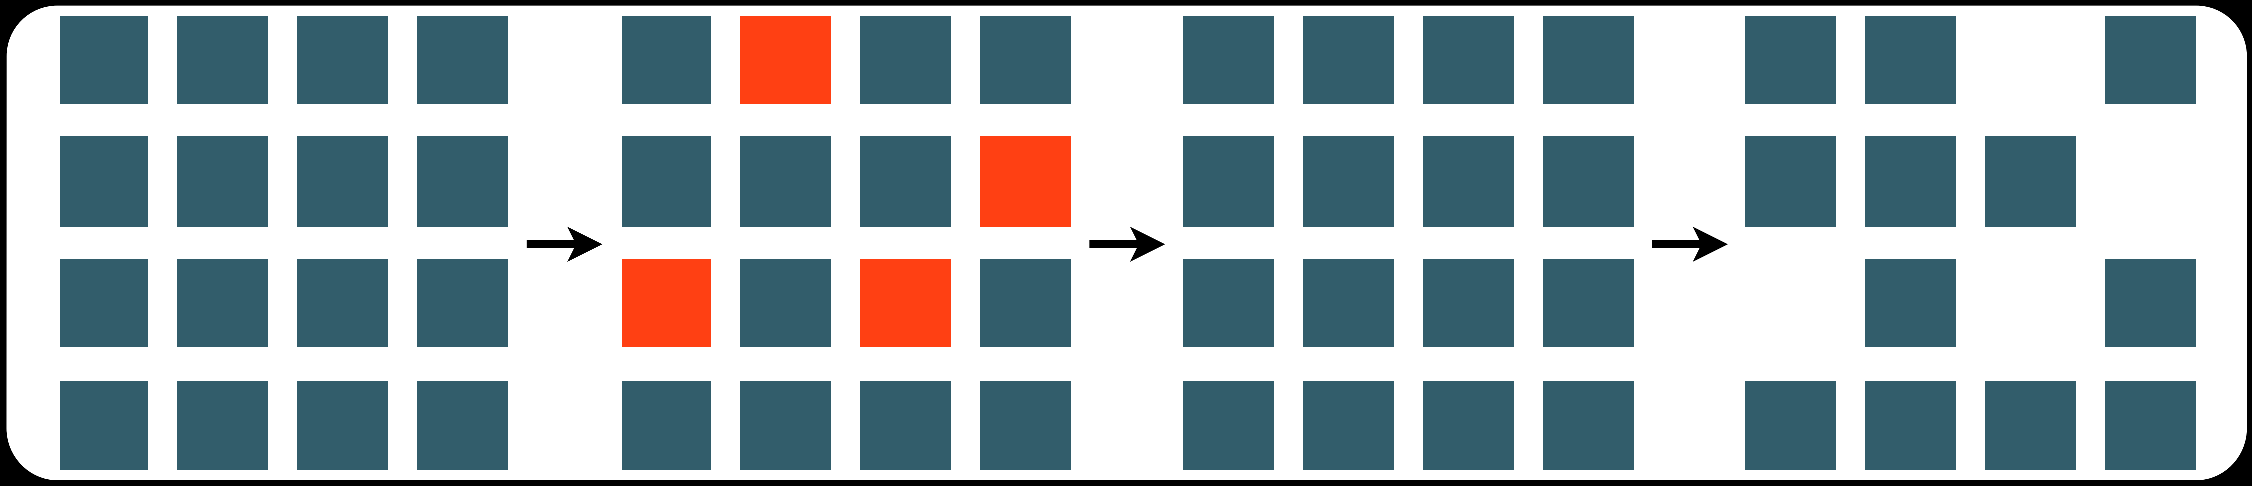

Claidiere, Smith, Kirby & Fagot (in press)

STUDY 5: QUANTITATIVE RESULTS
TASK GETS EASIER OVER GENERATIONS

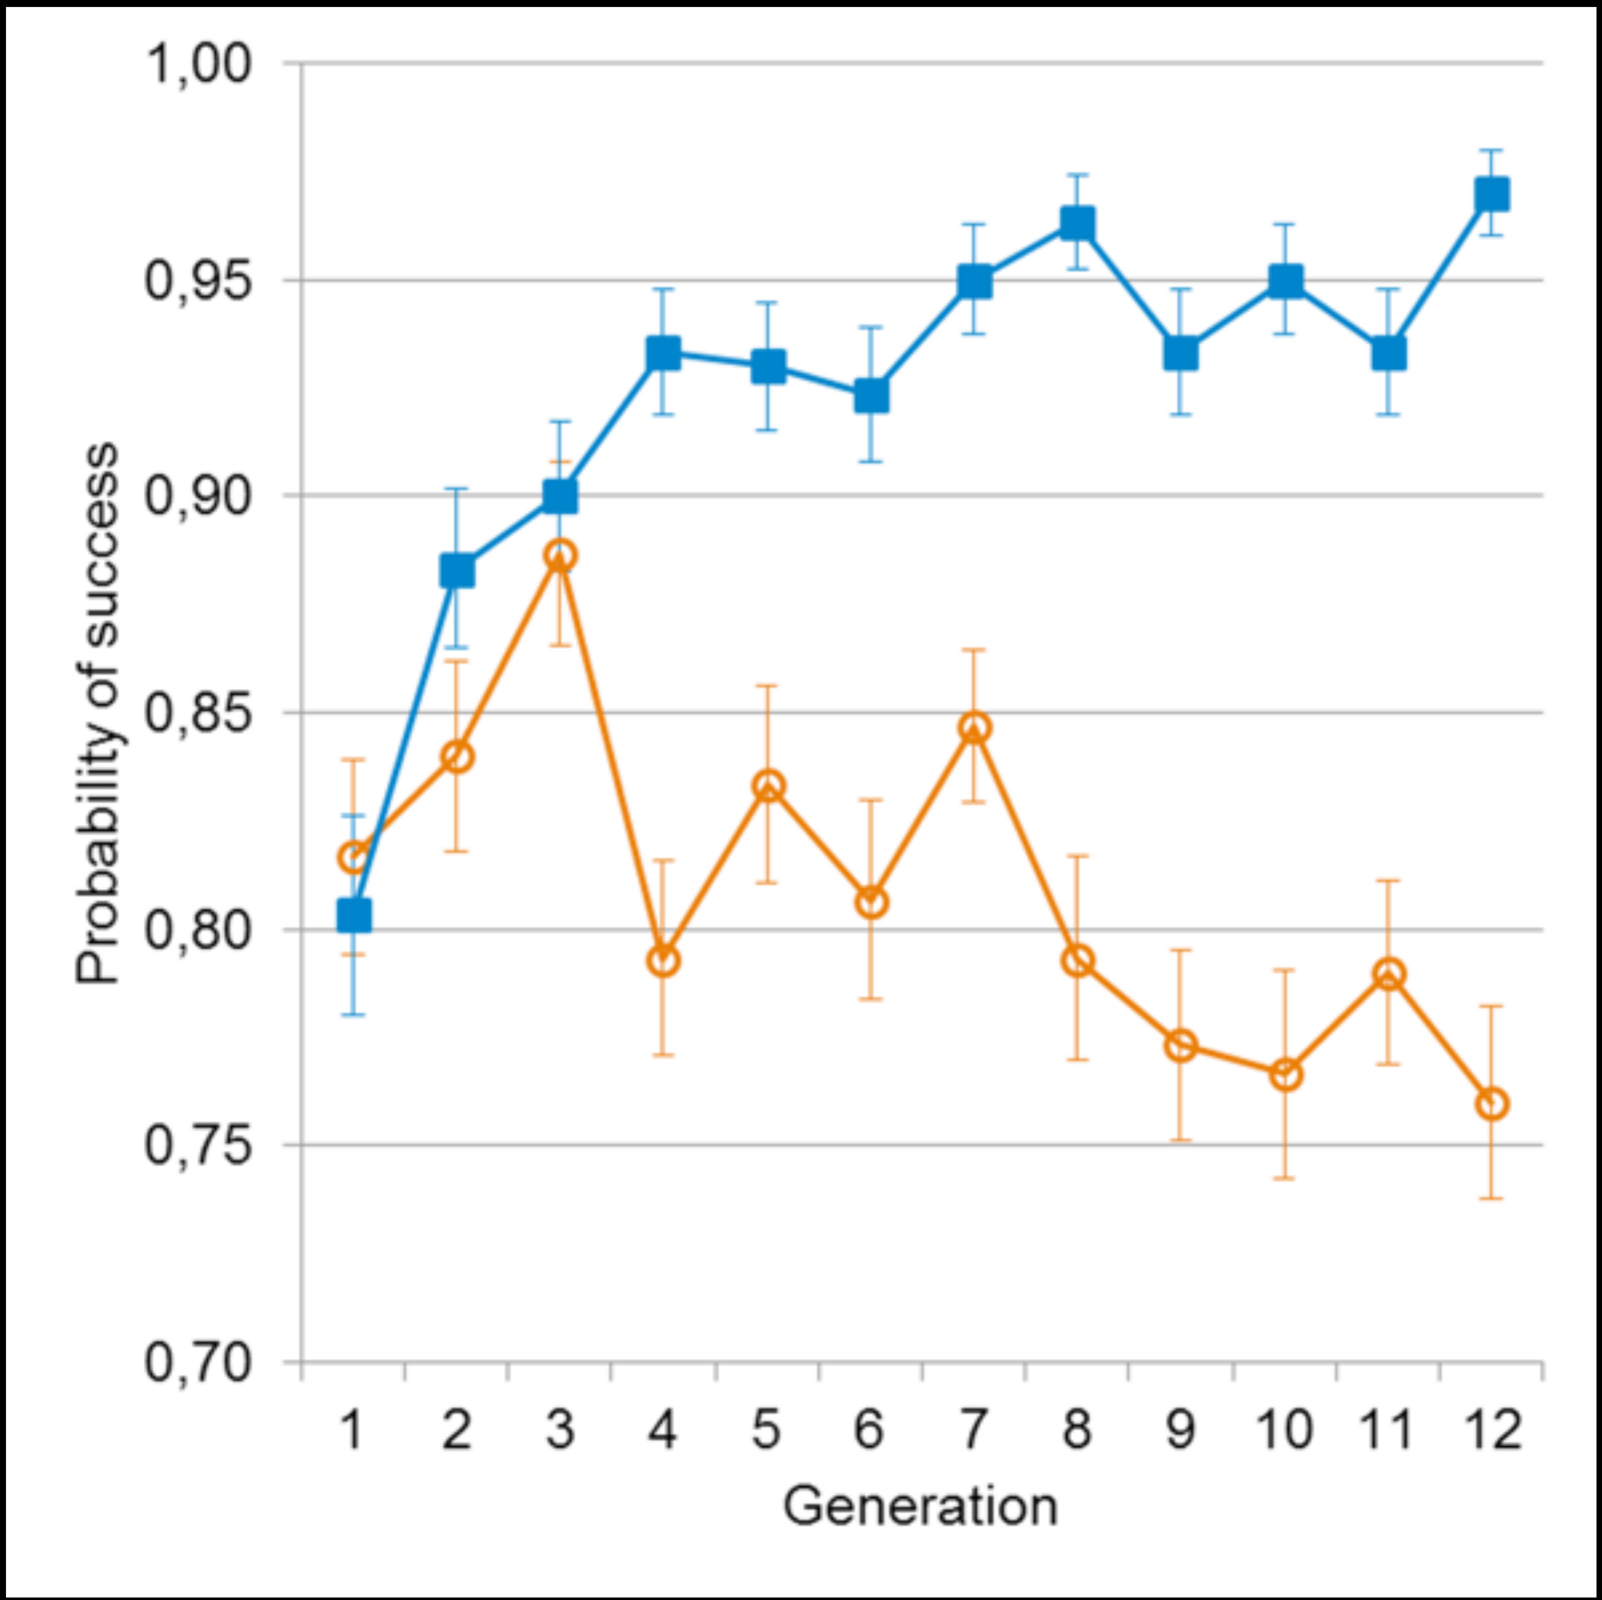

Transmission trials

Matched random trials

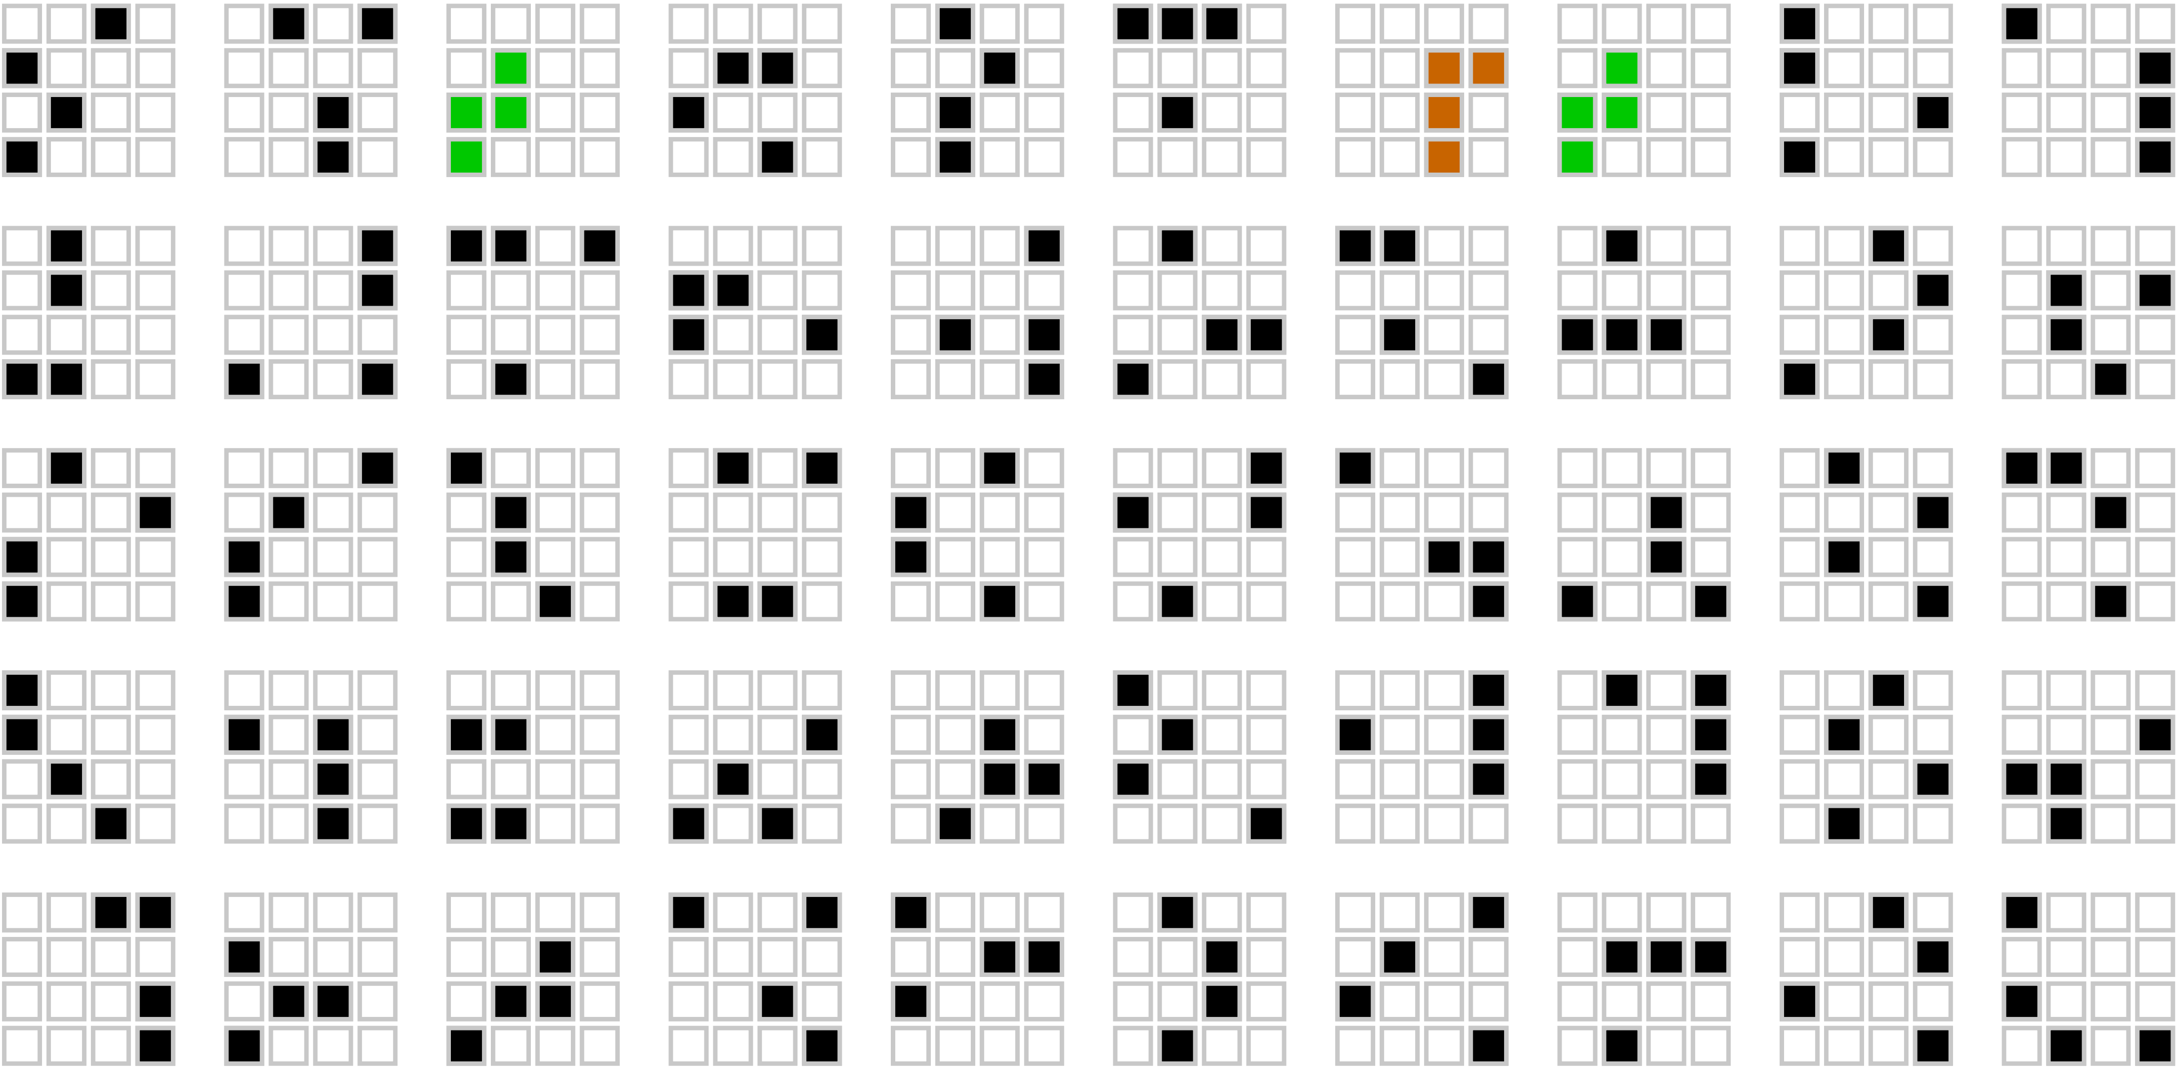

4:0

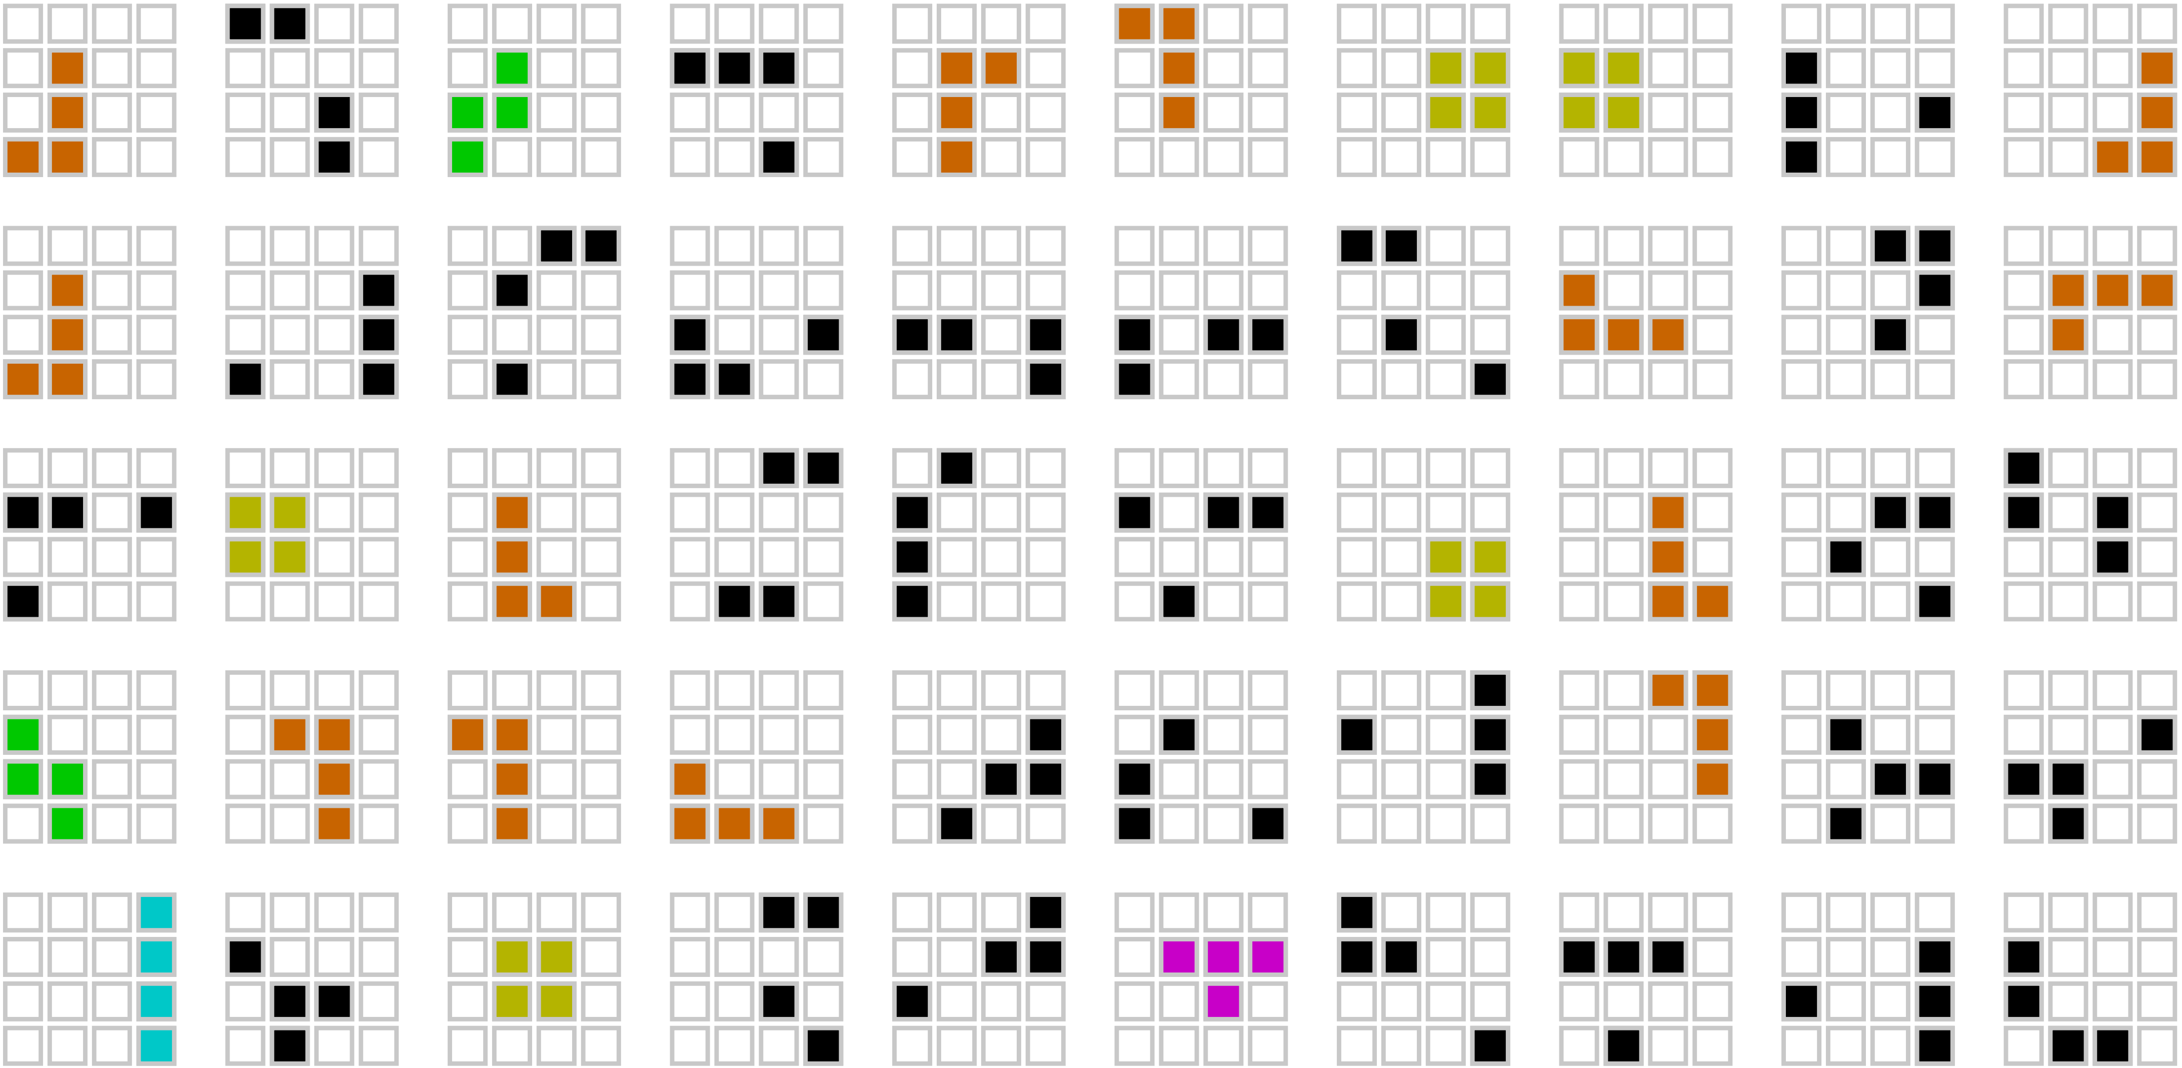

4:1

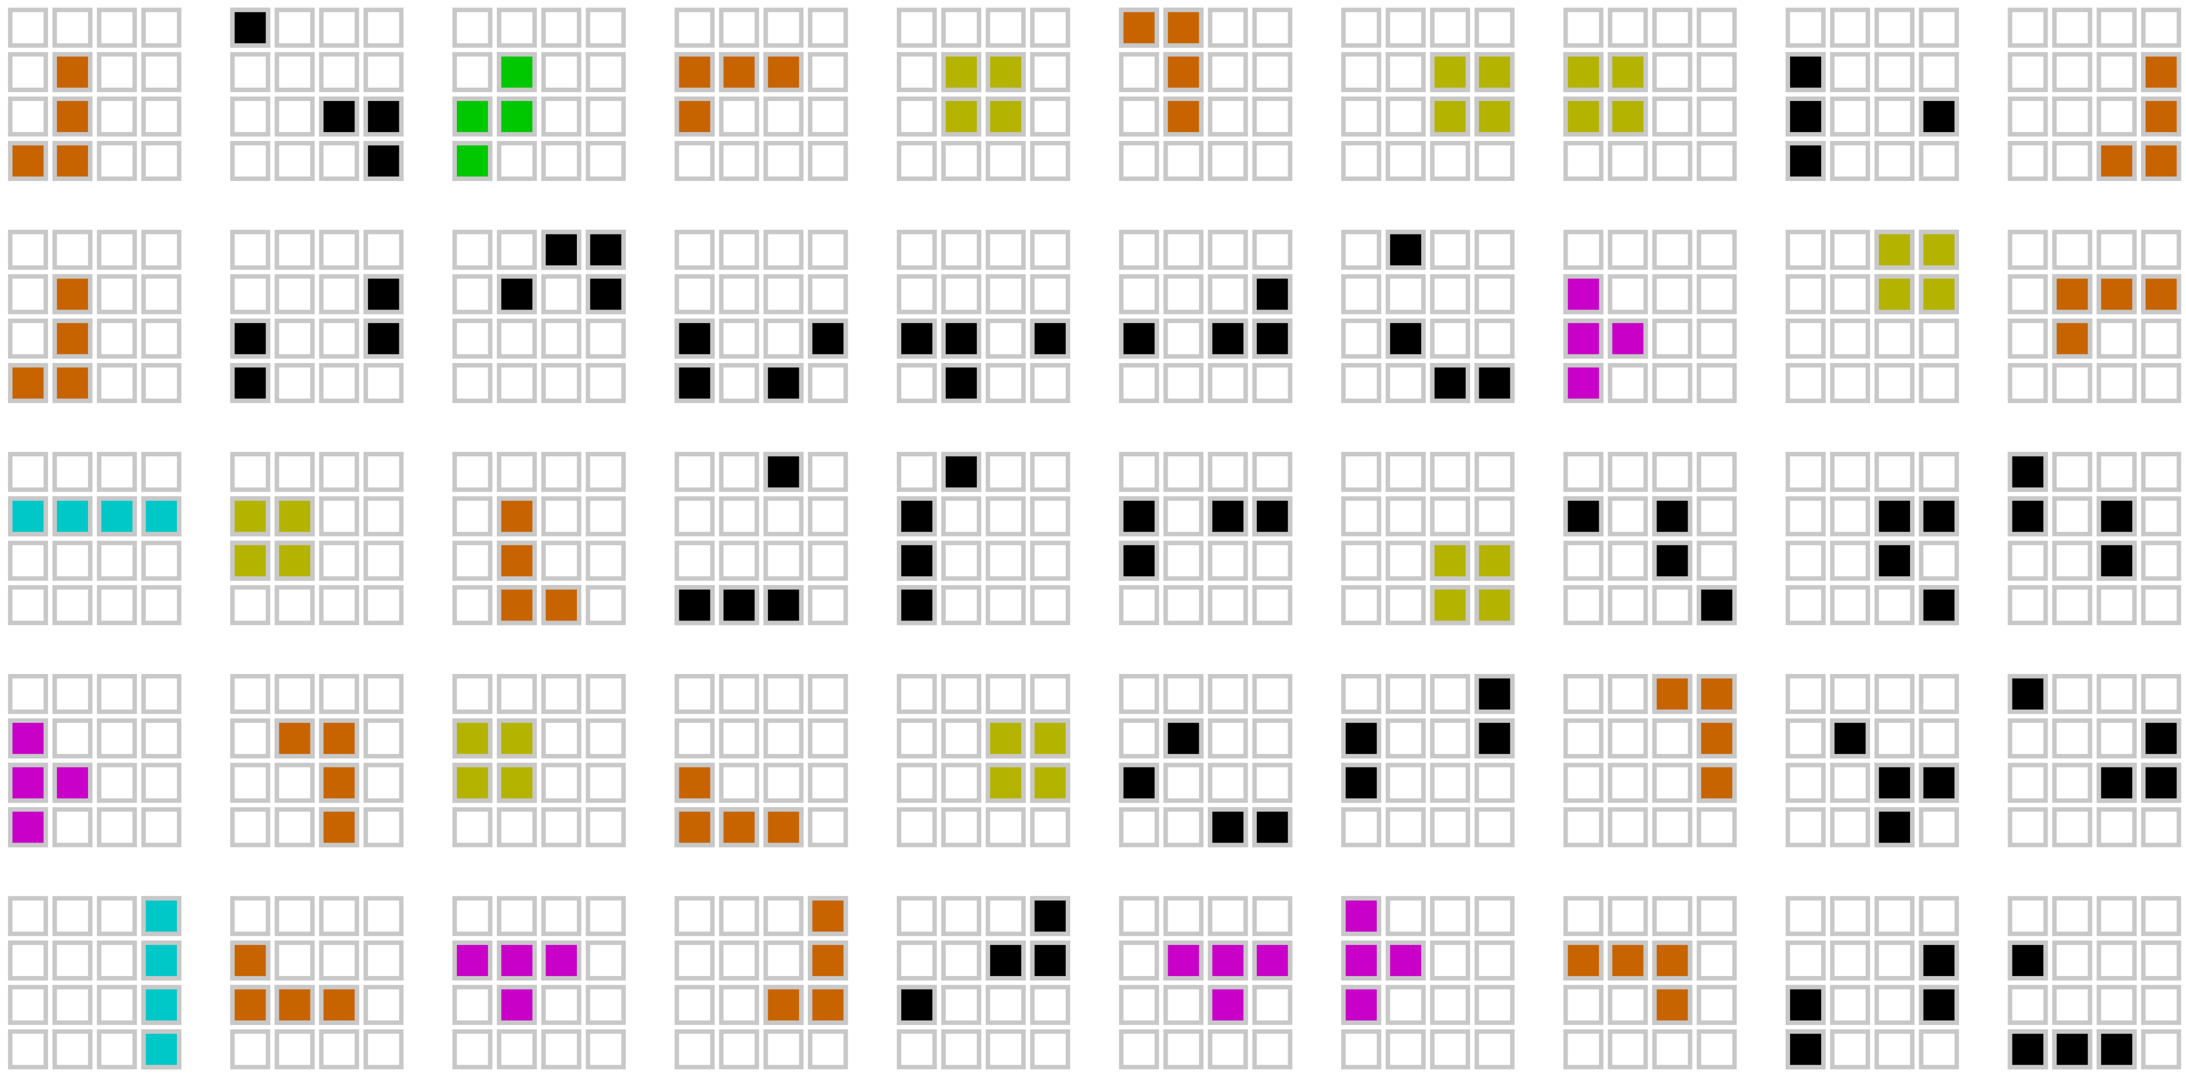

4:2

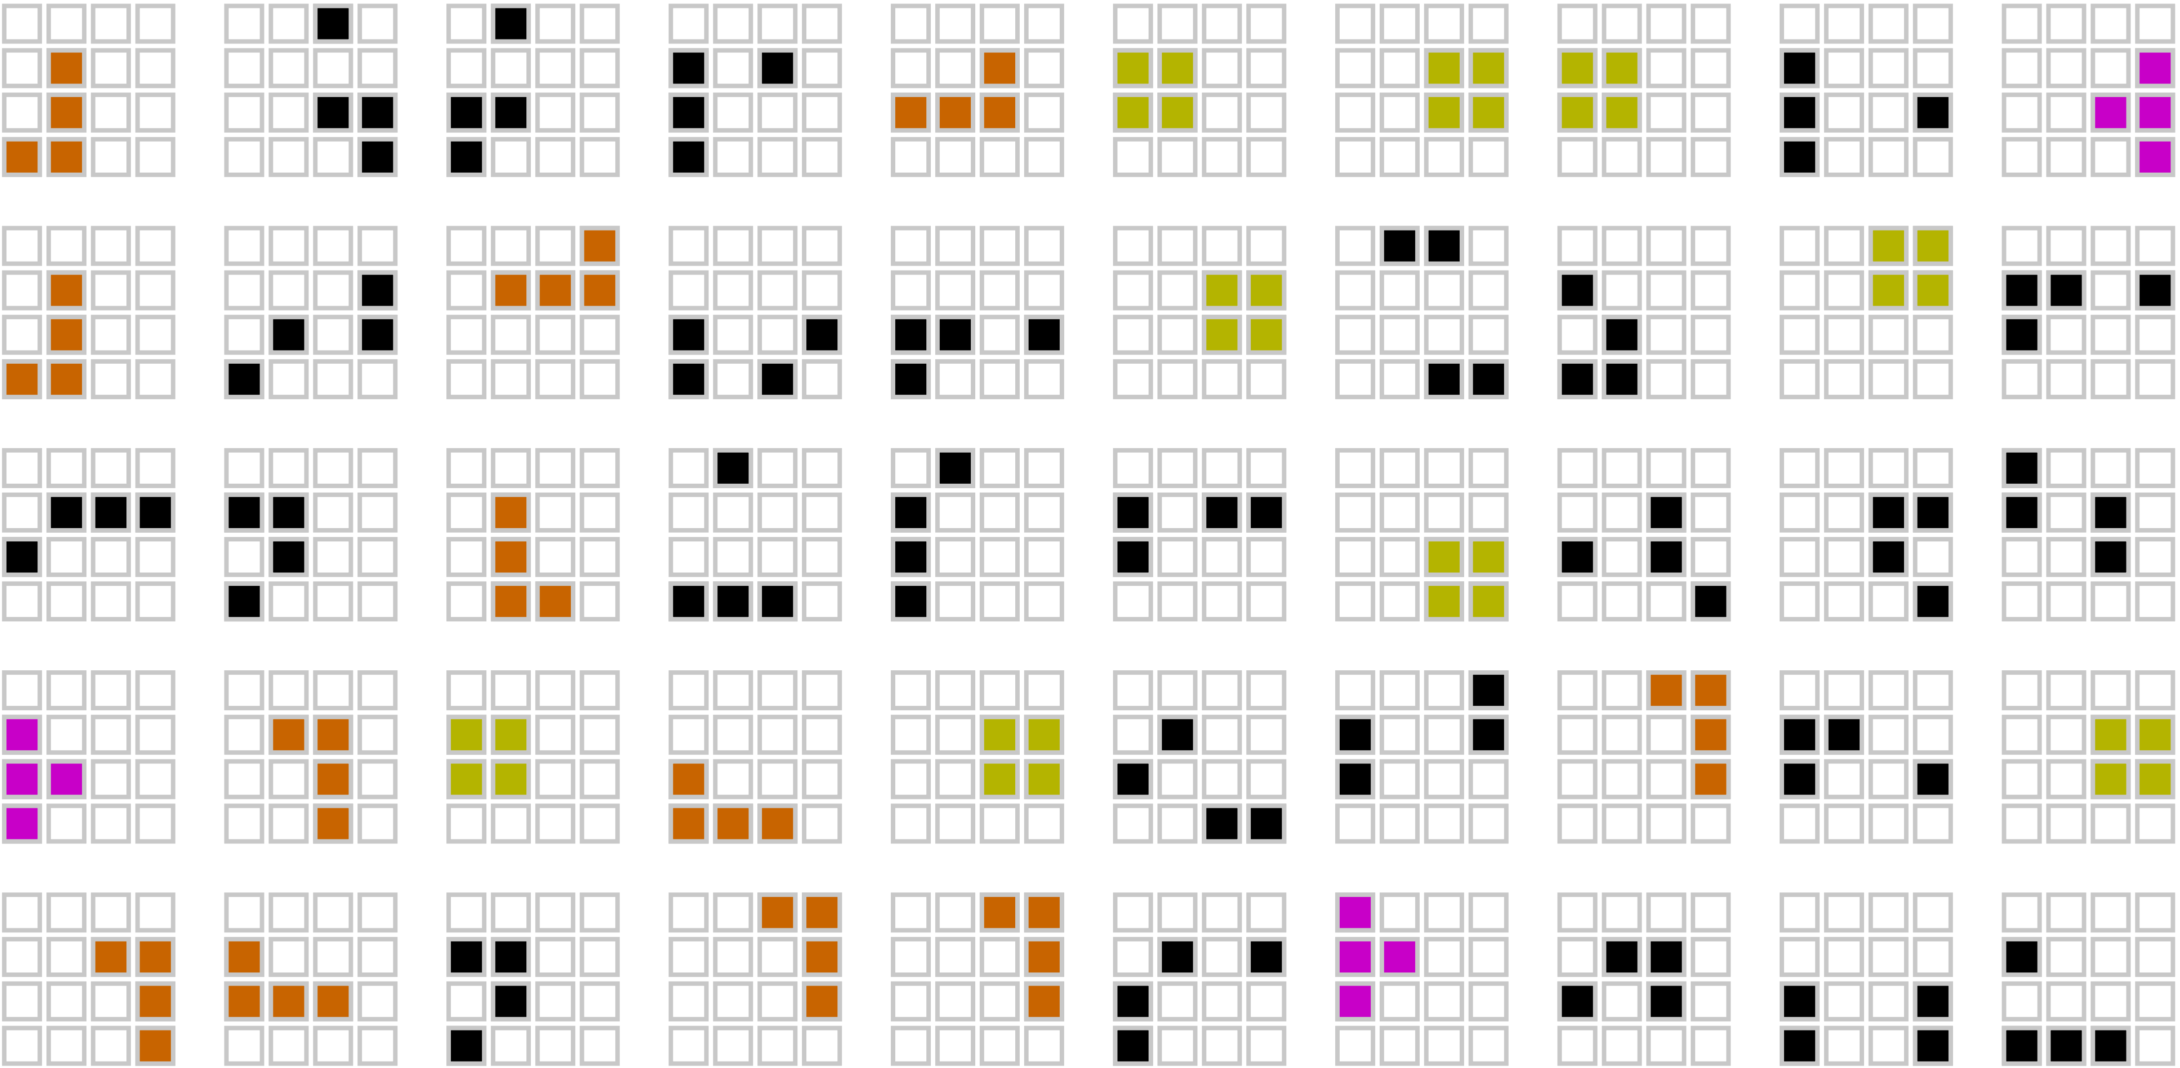

4:3

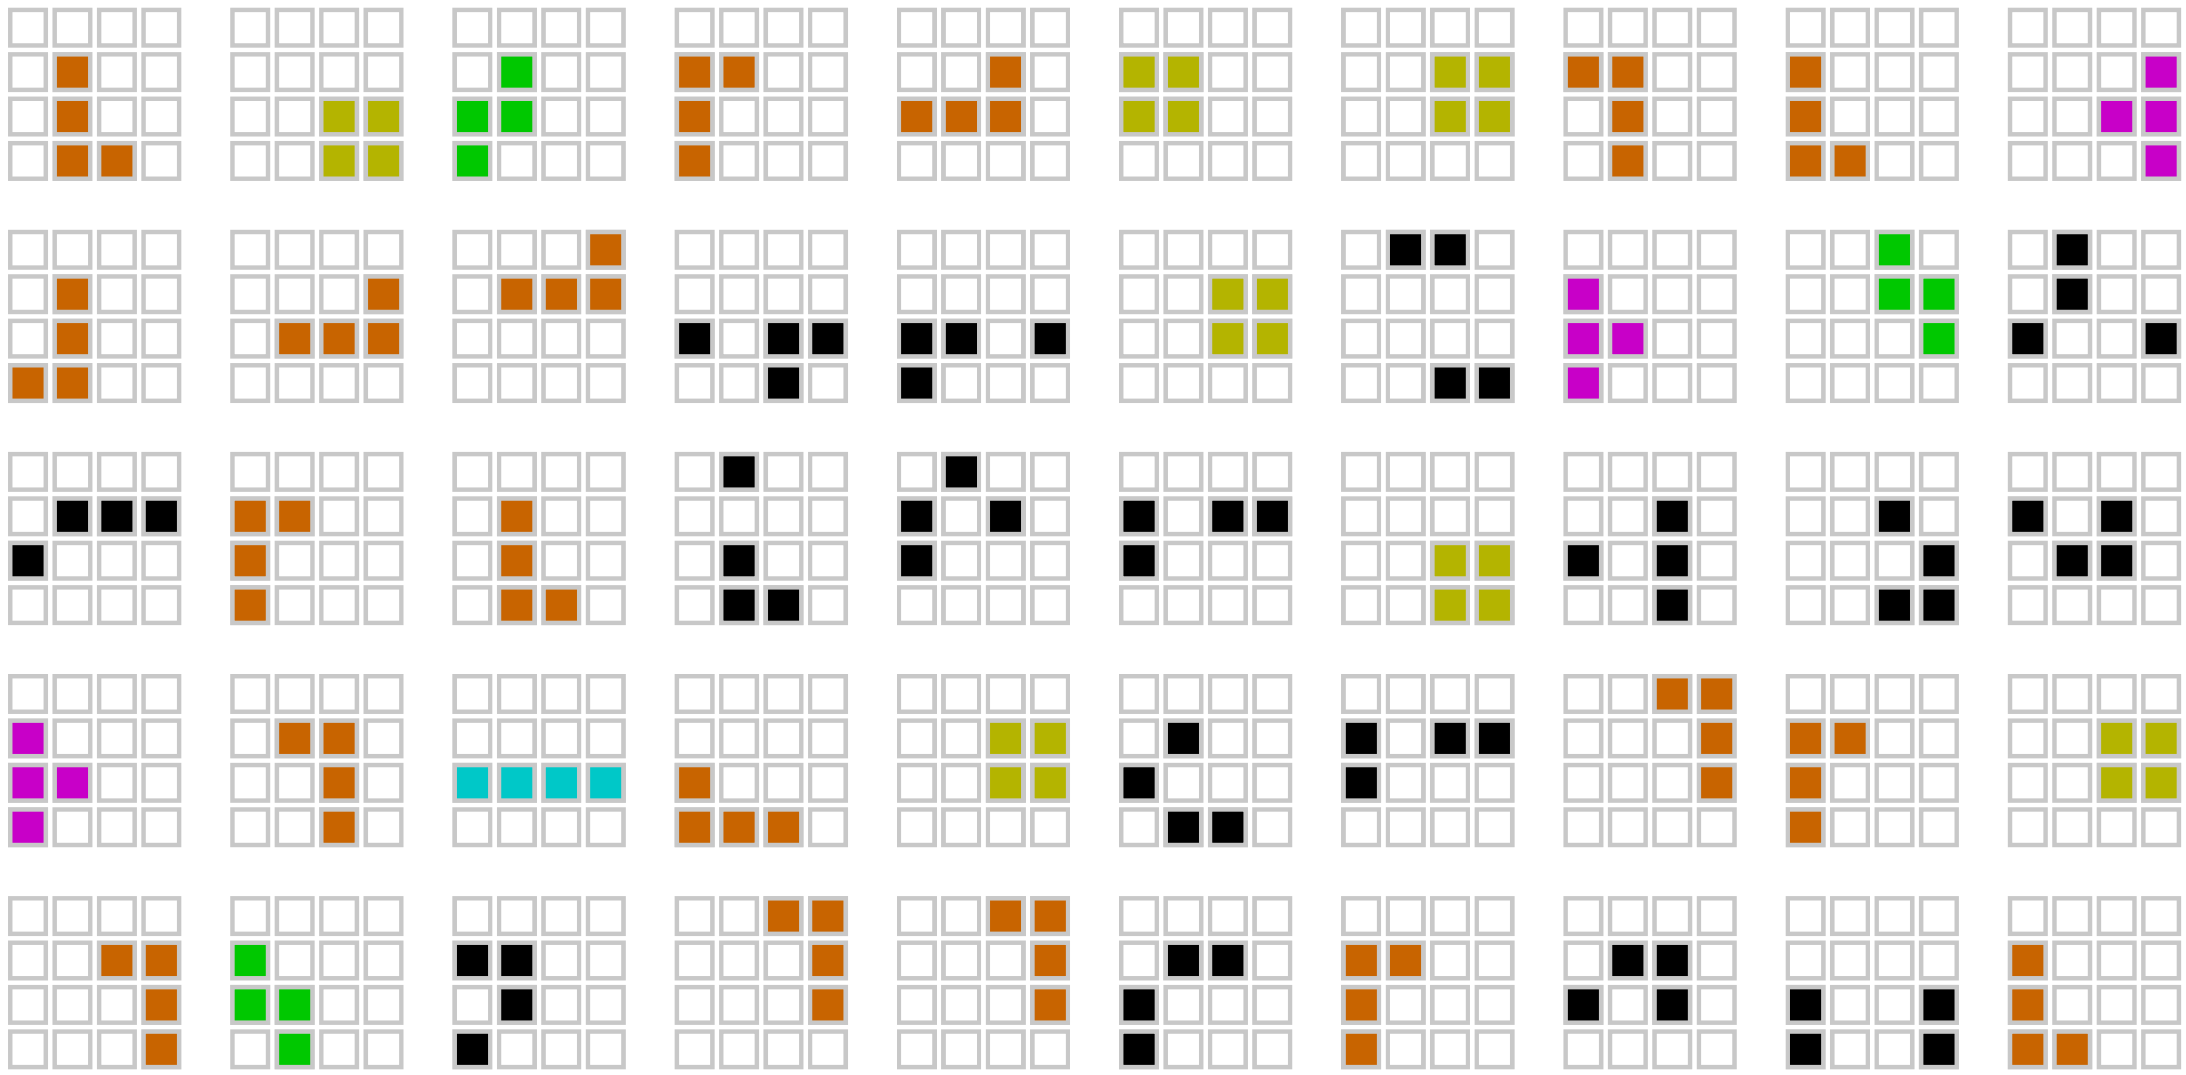

4:4

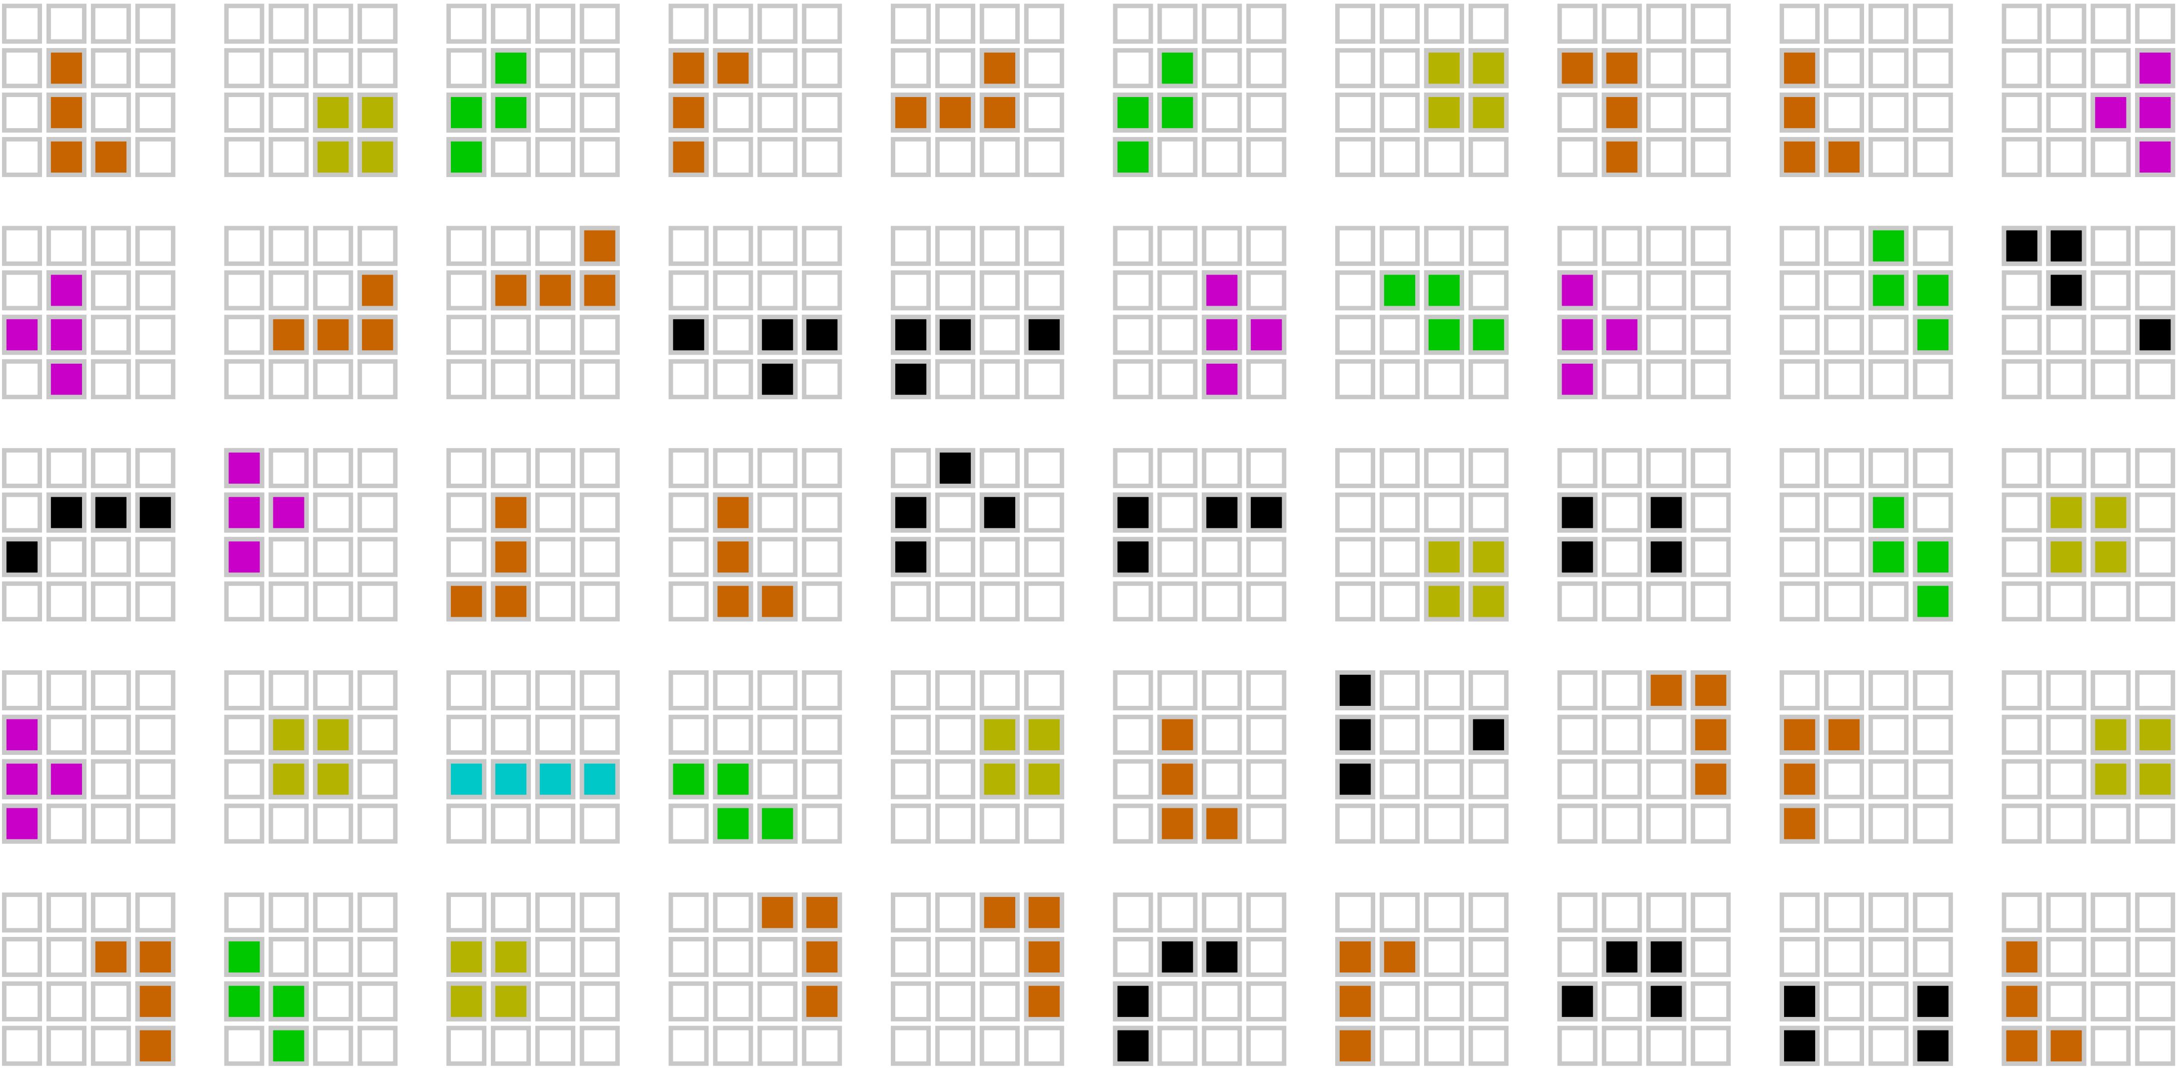

4:5

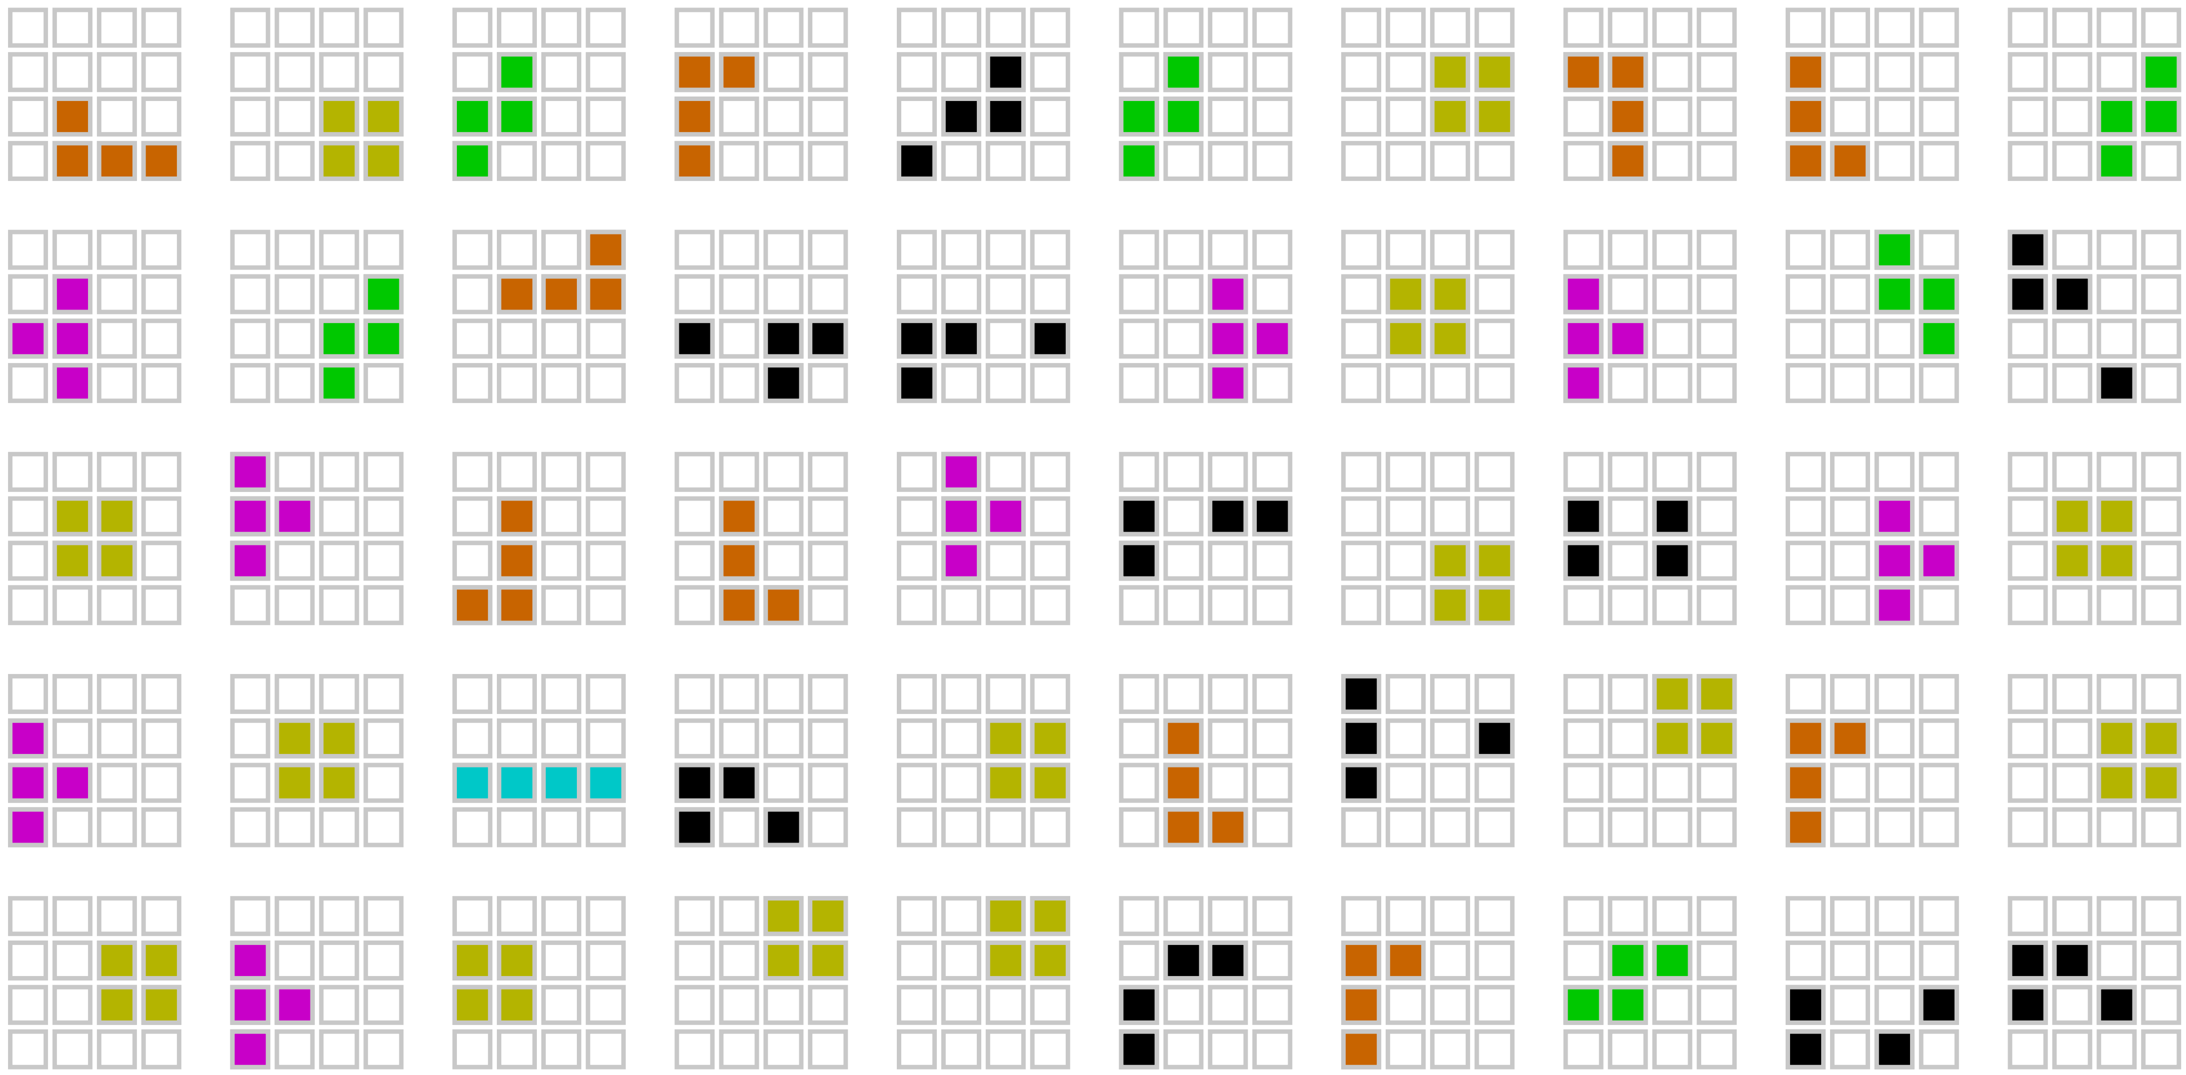

4:6

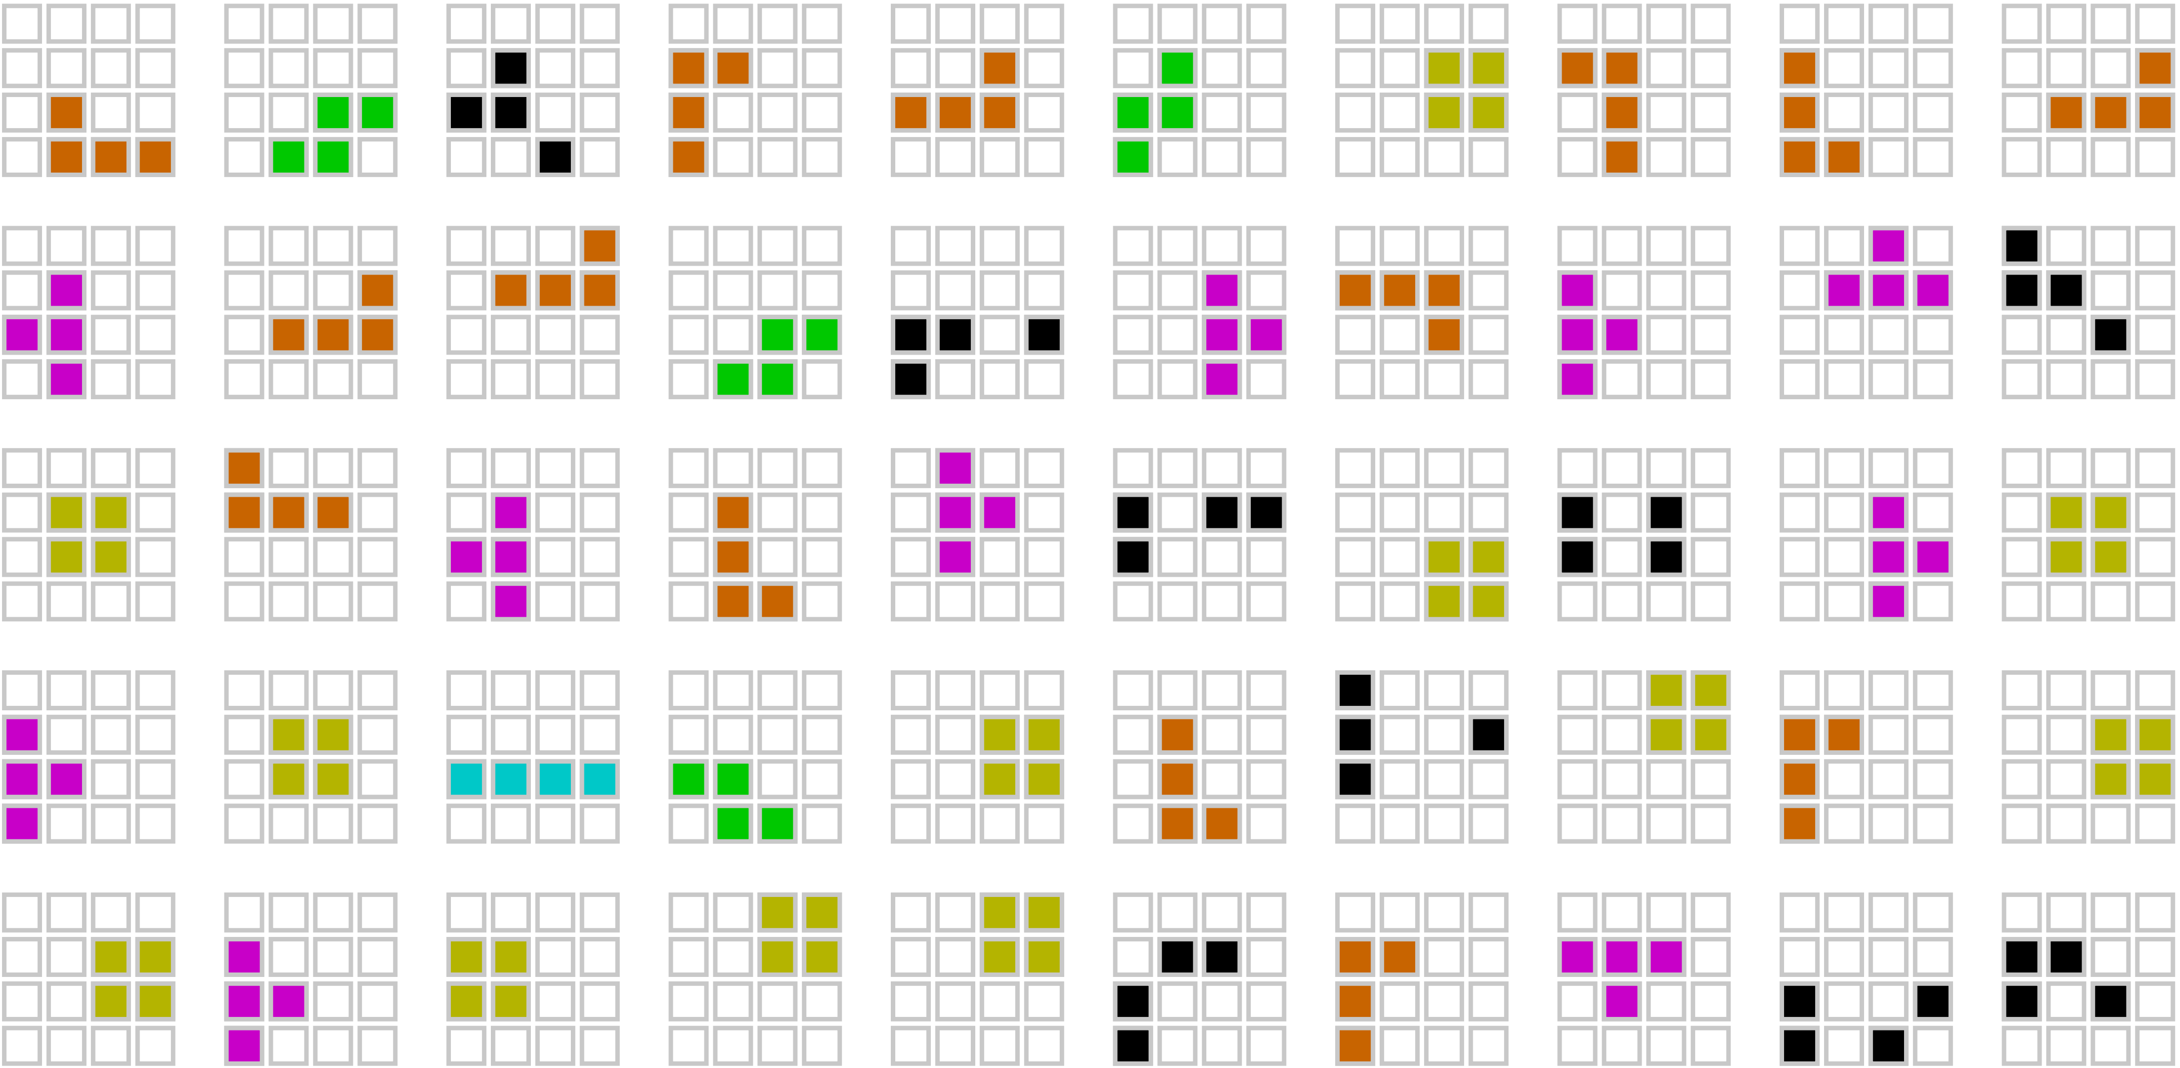

4:7

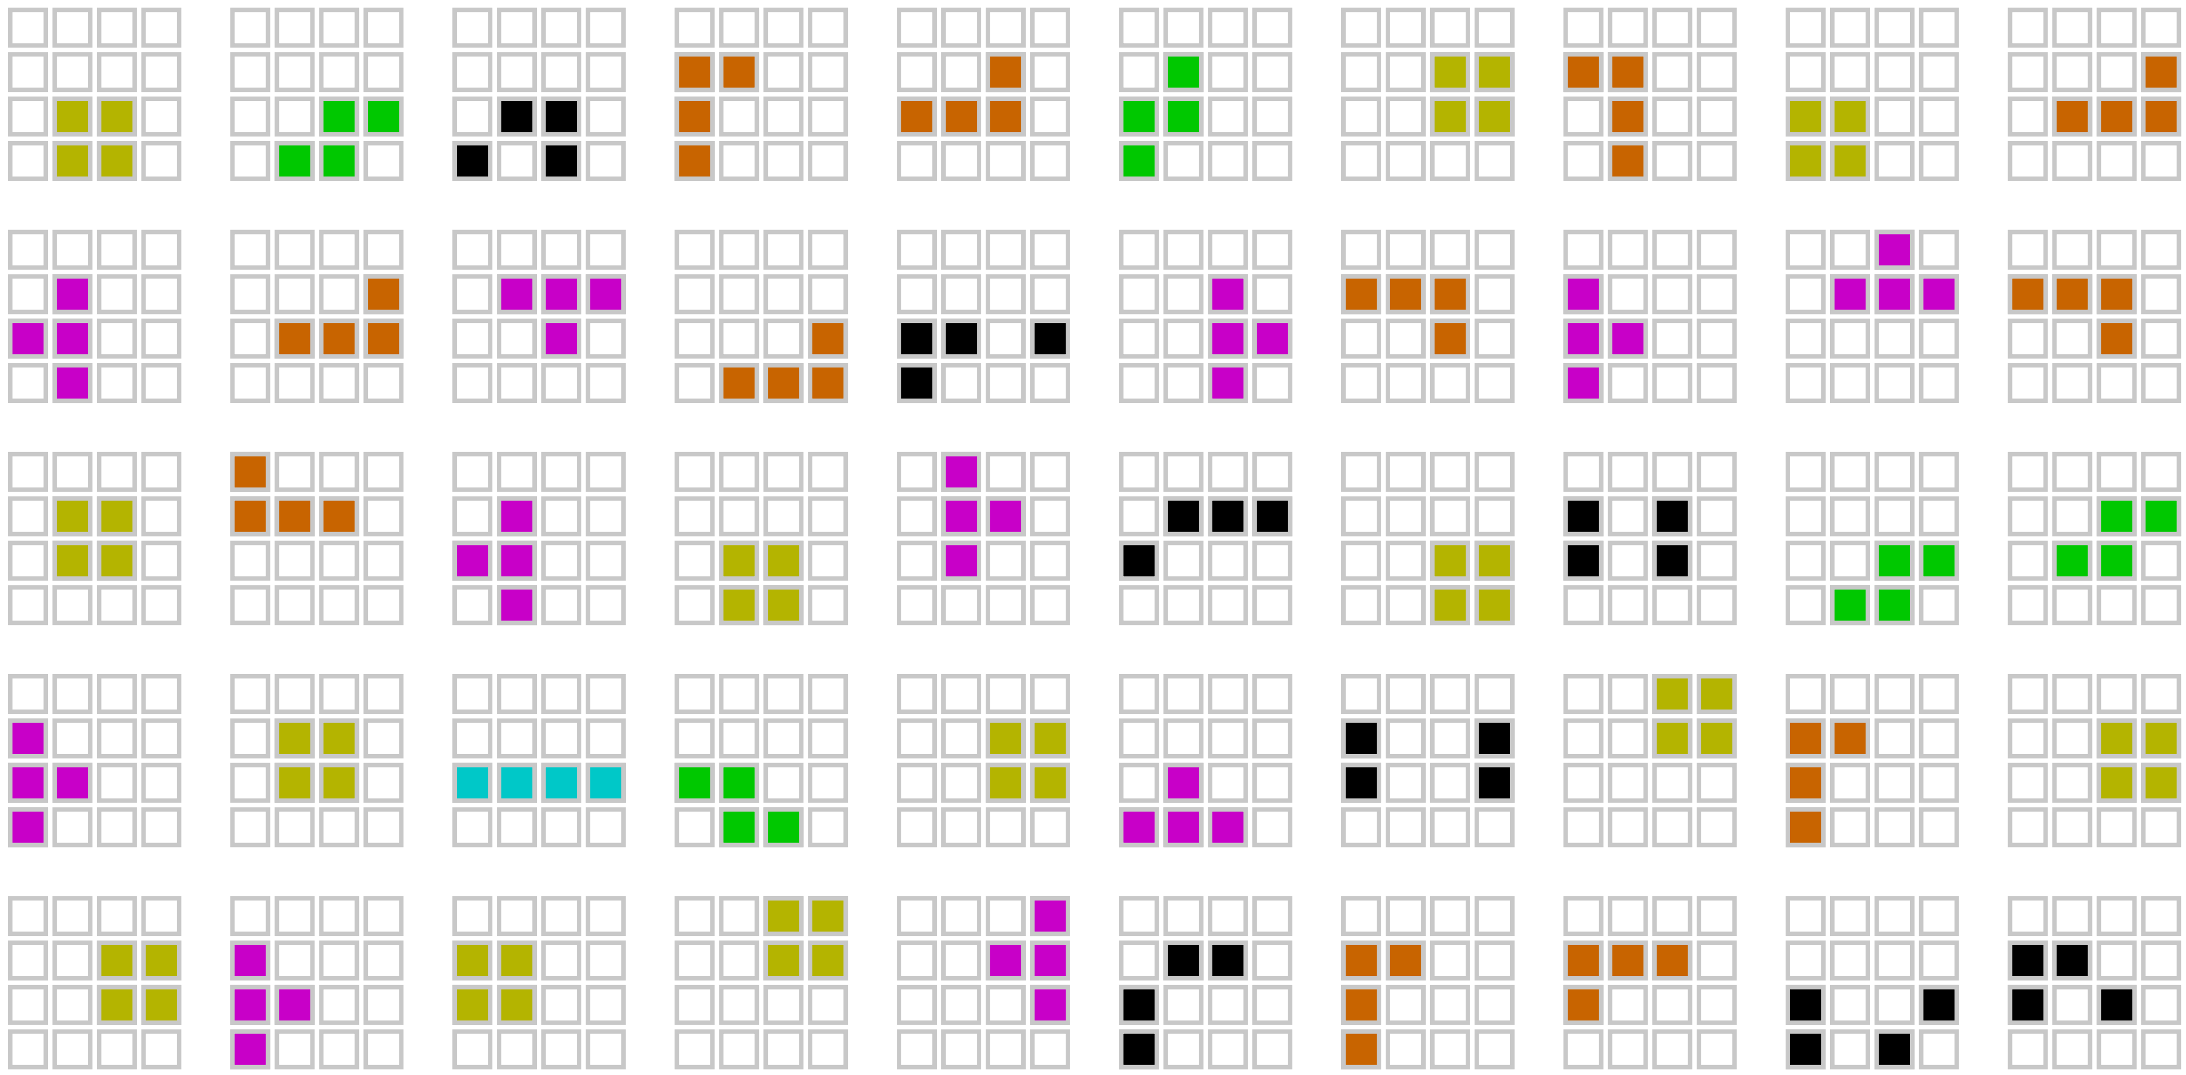

4:8

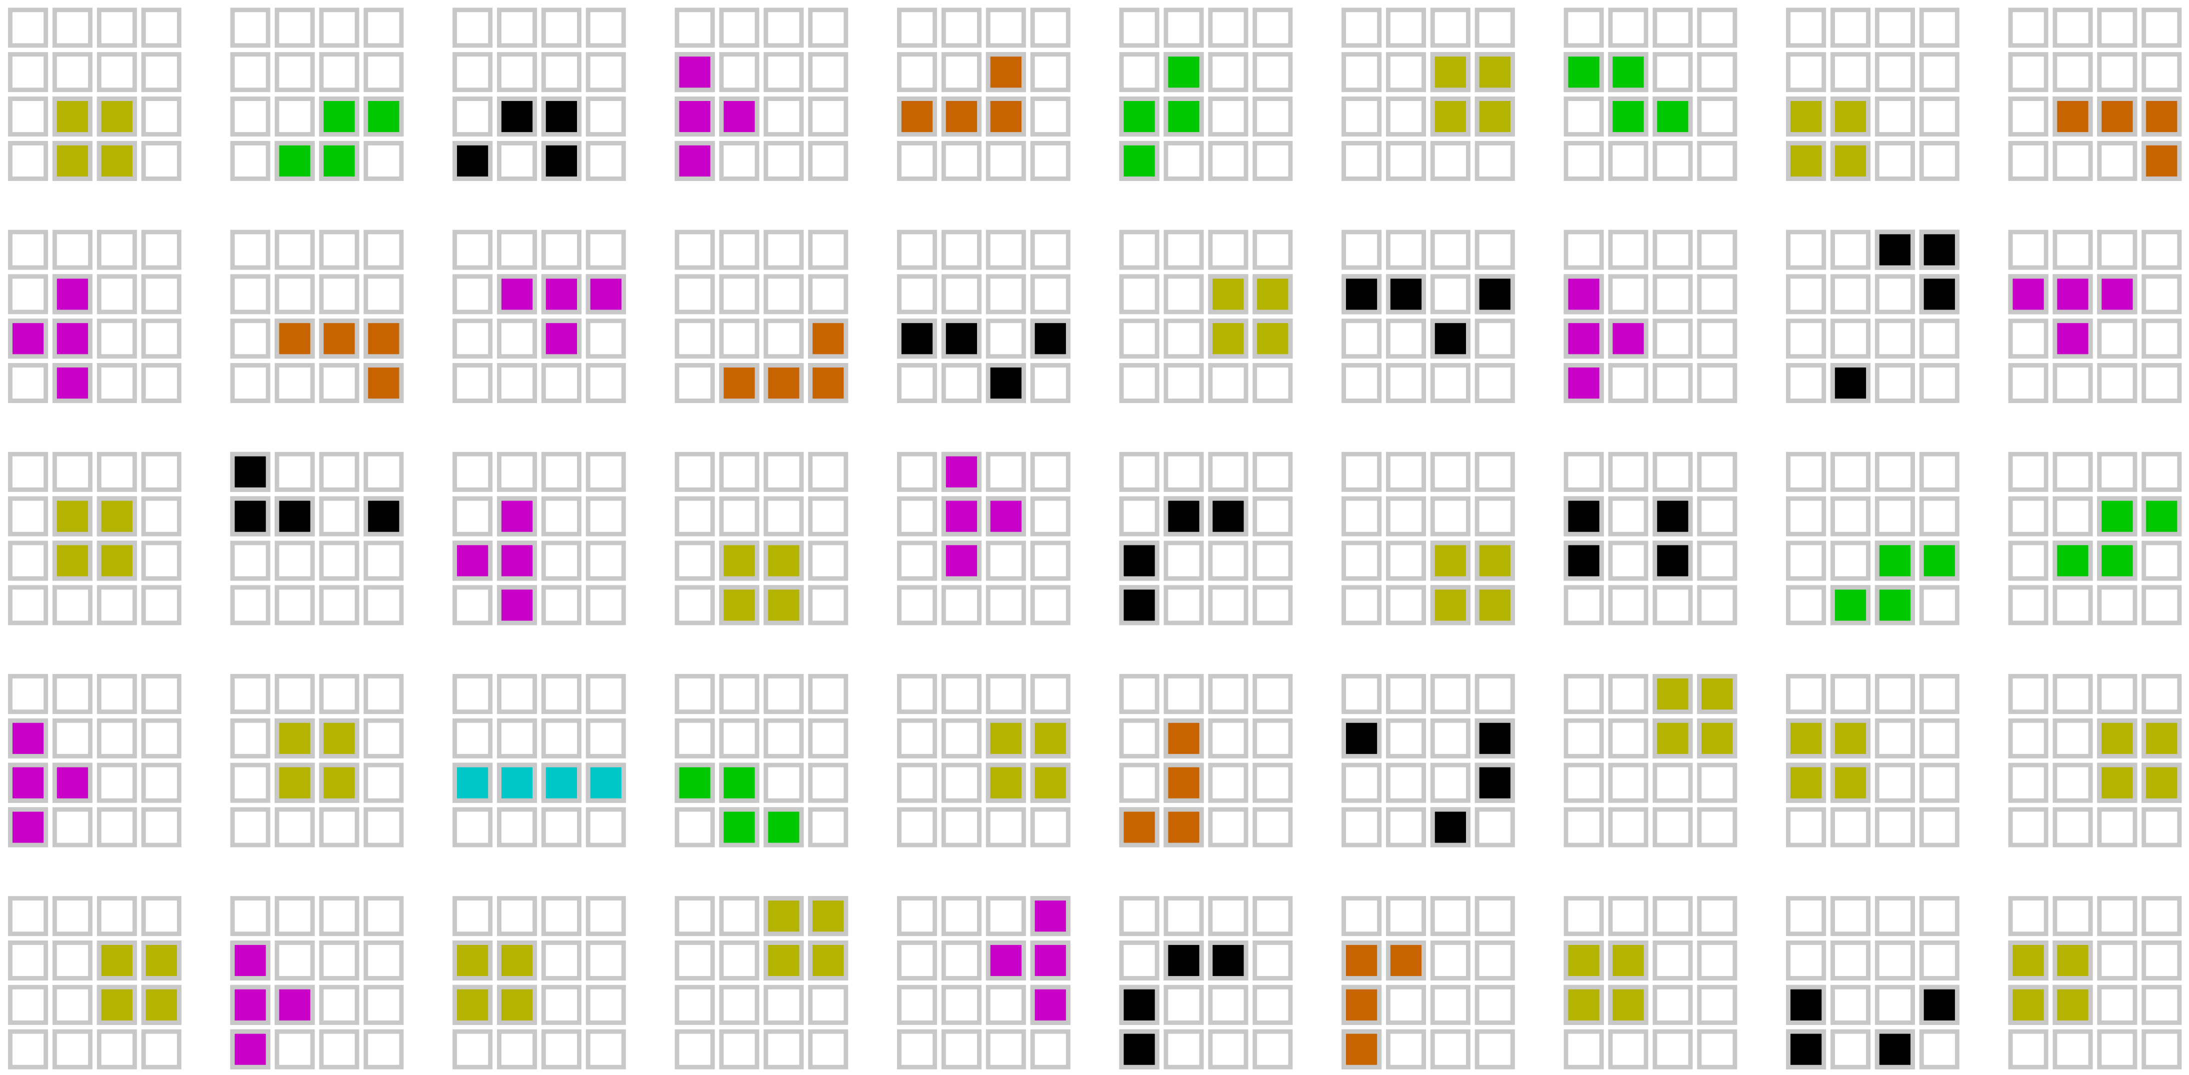

4:9

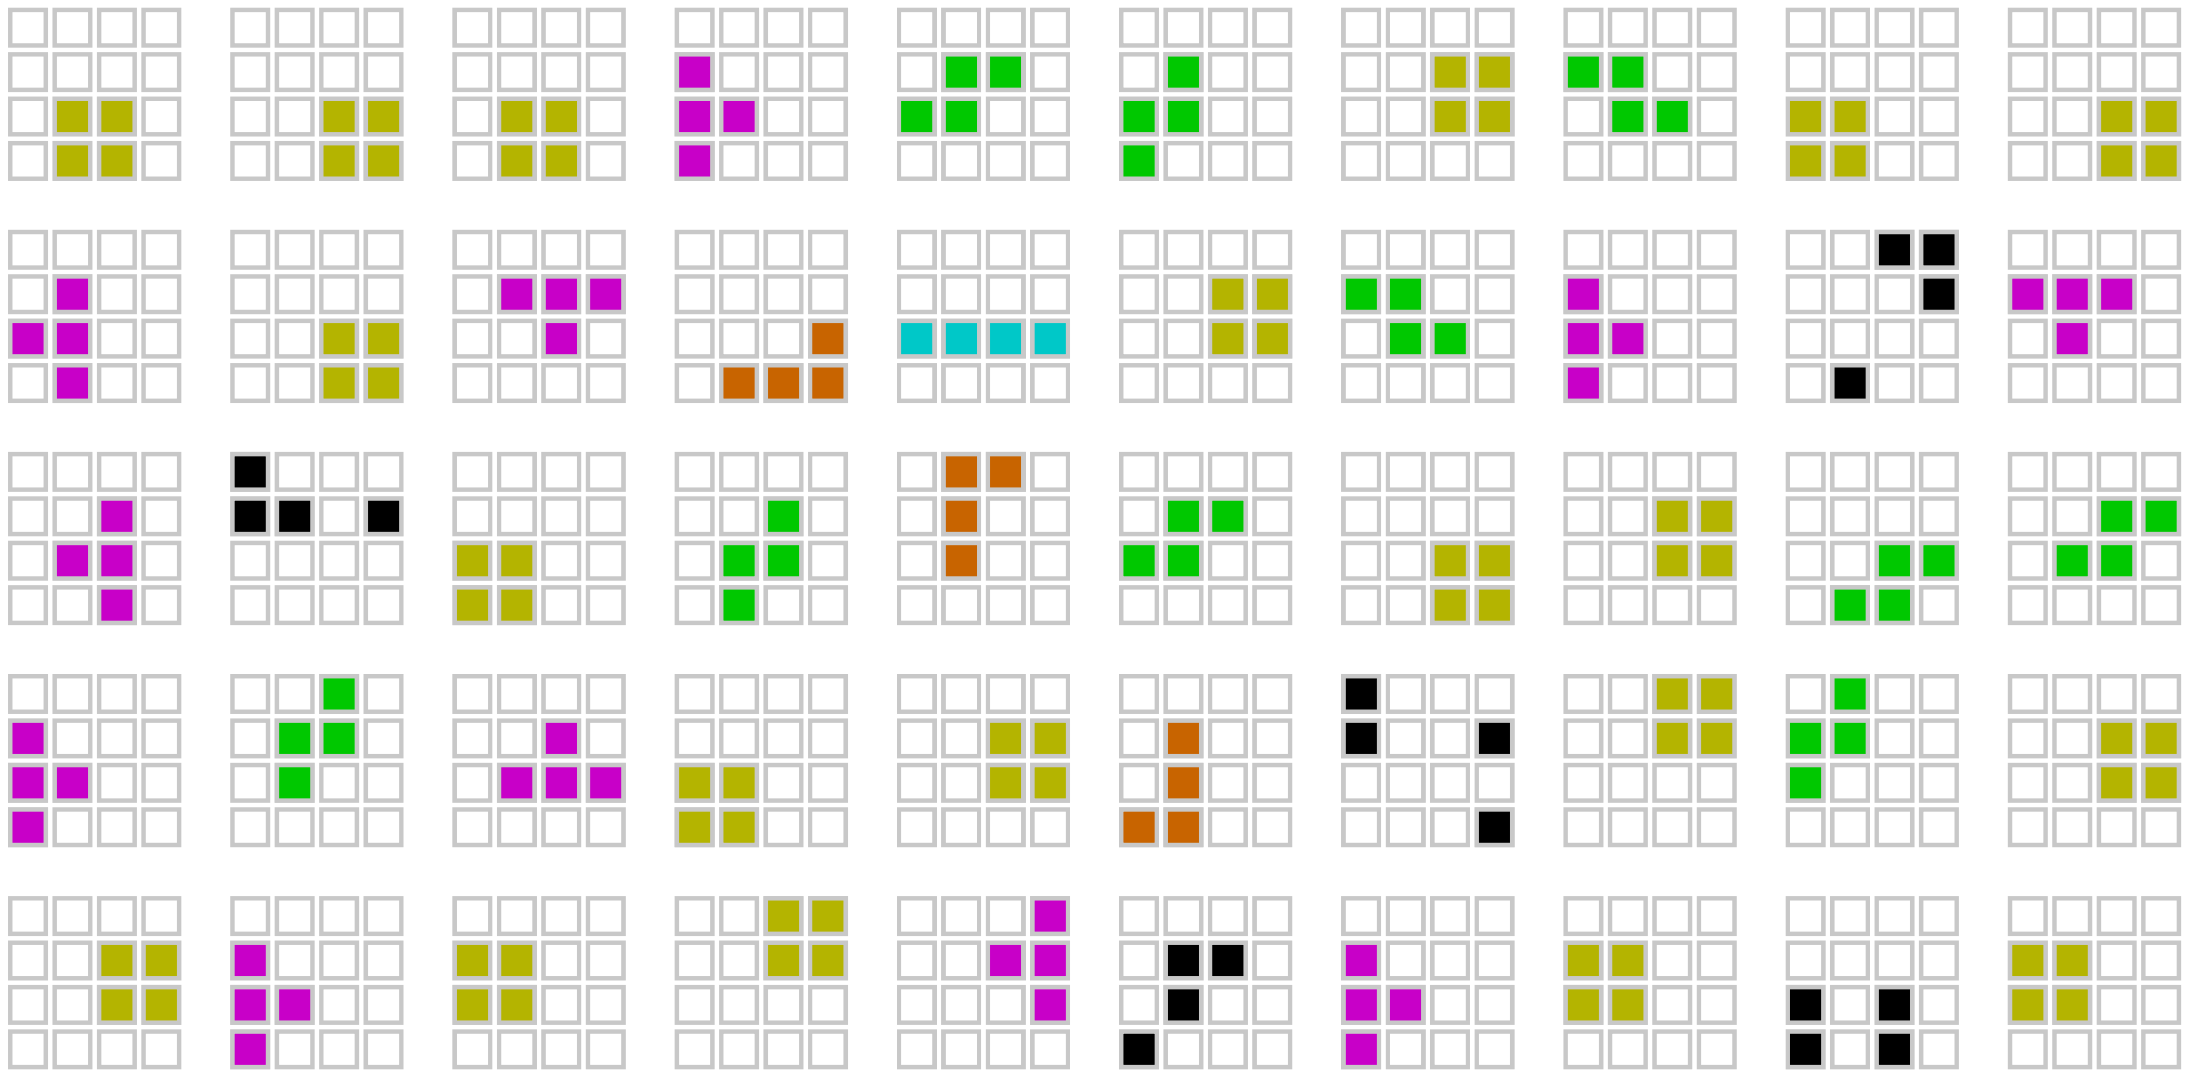

4:10

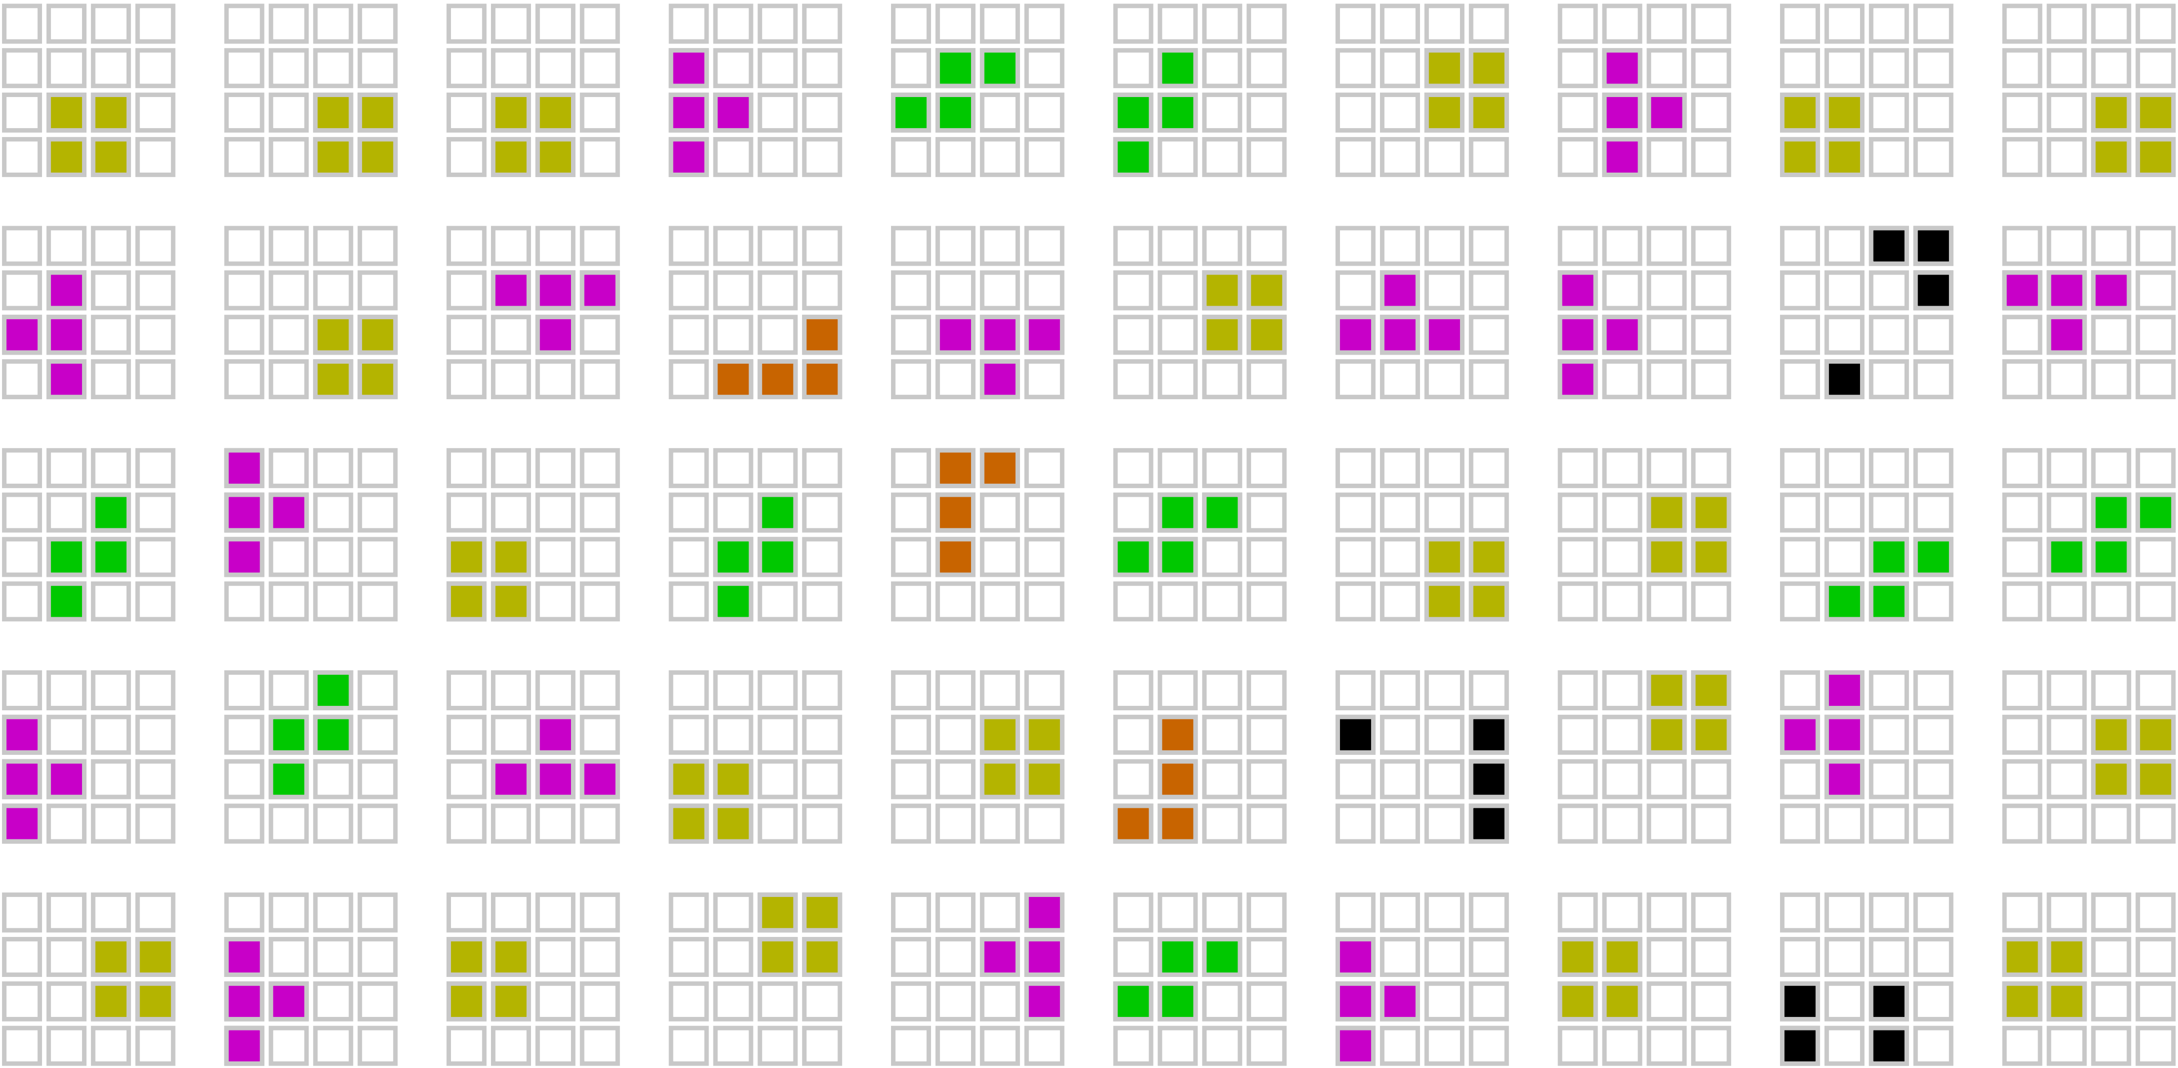

4:11

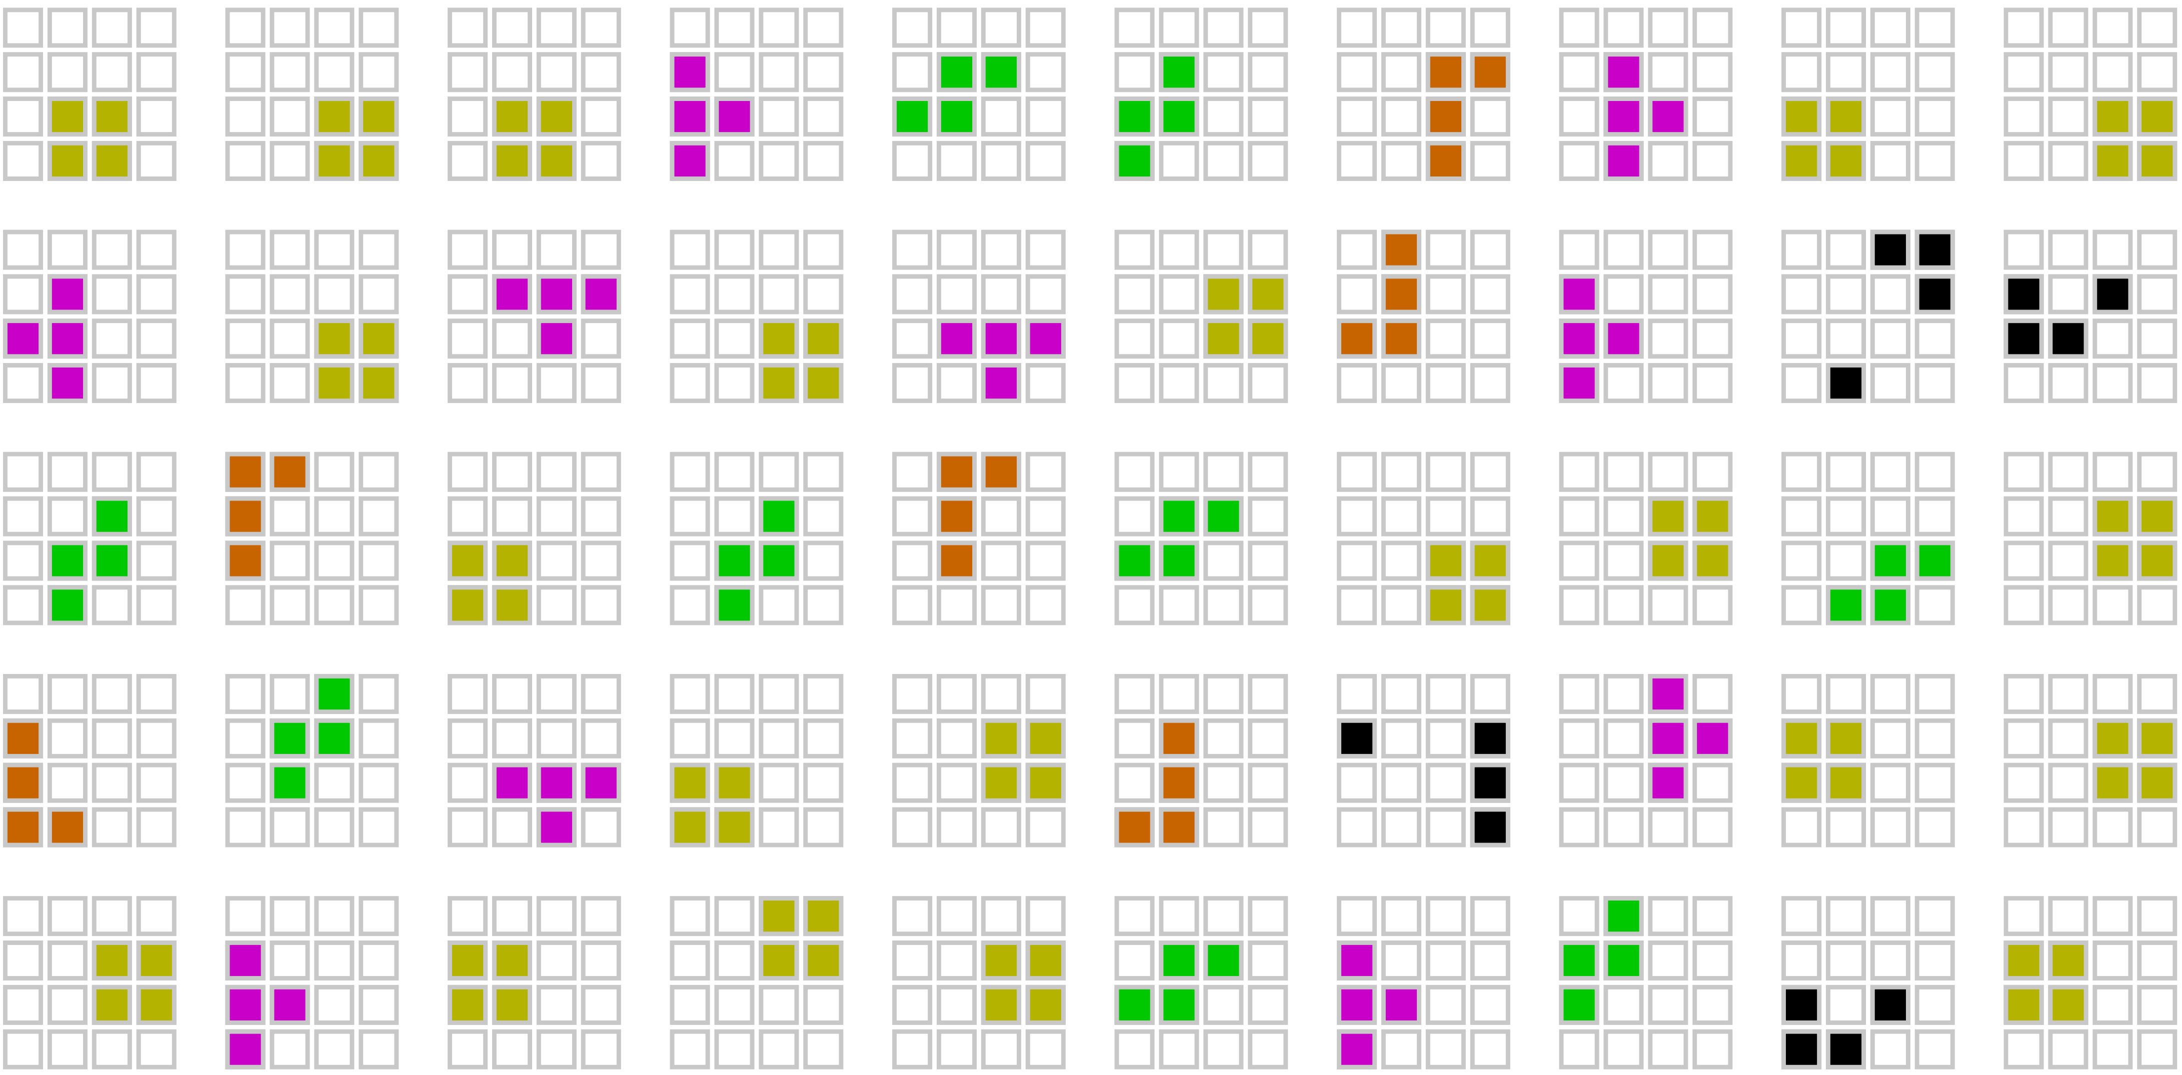

4:12

STUDY 5: TETROMINOS

- Tetrominos: finite subset of the regular square tiling with a connected interior using 4 squares

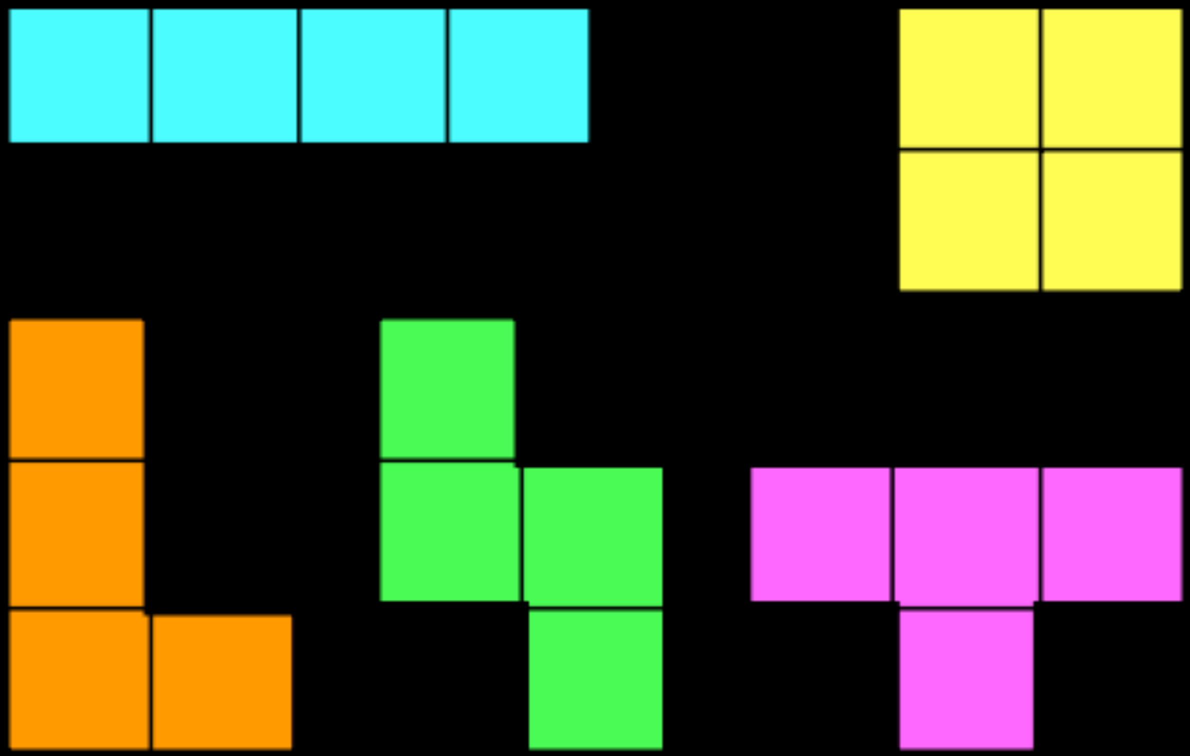

STUDY 5: TETROMINOS

- Tetrominos: finite subset of the regular square tiling with a connected interior using 4 squares

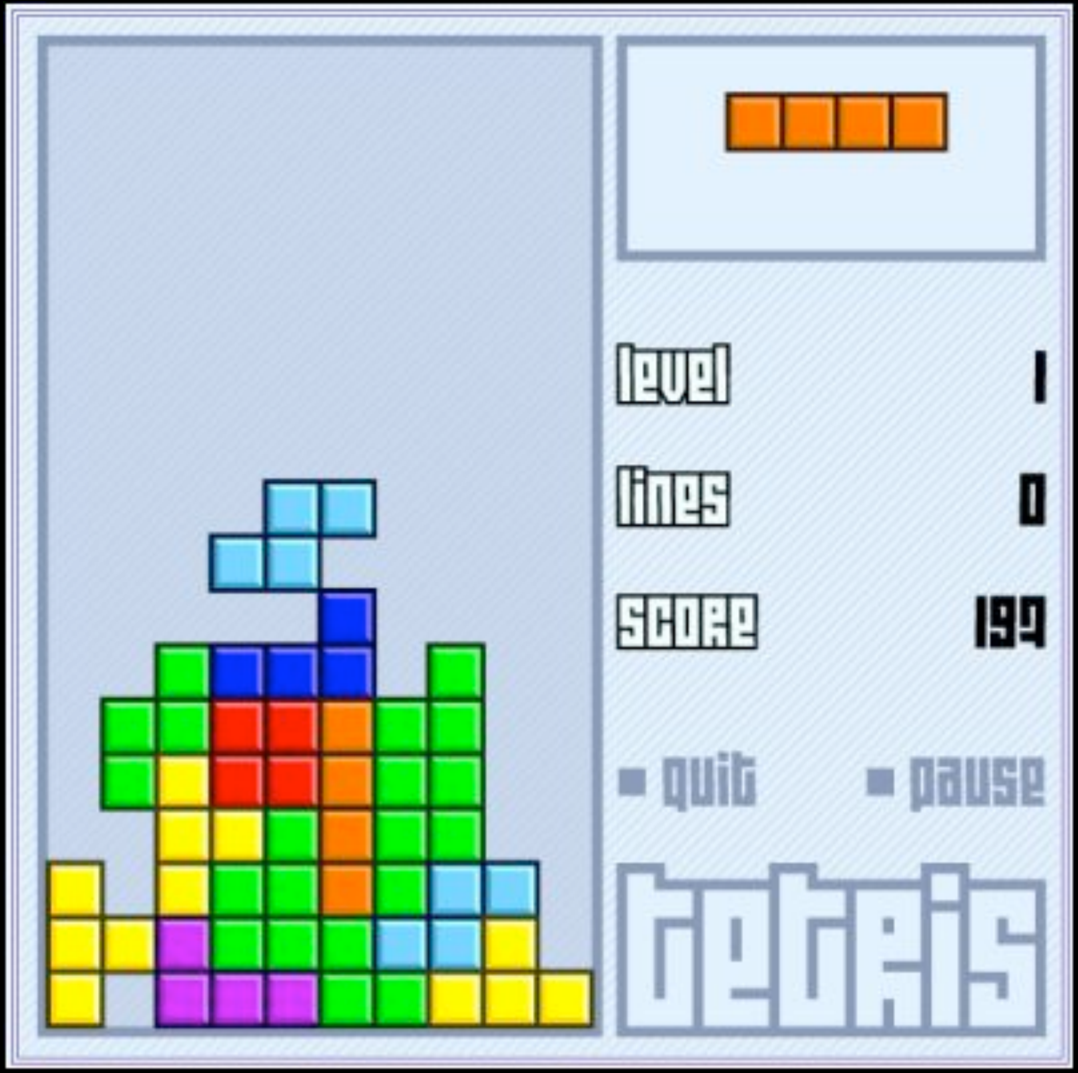

STUDY 5: EXPECTED TETROMINOS

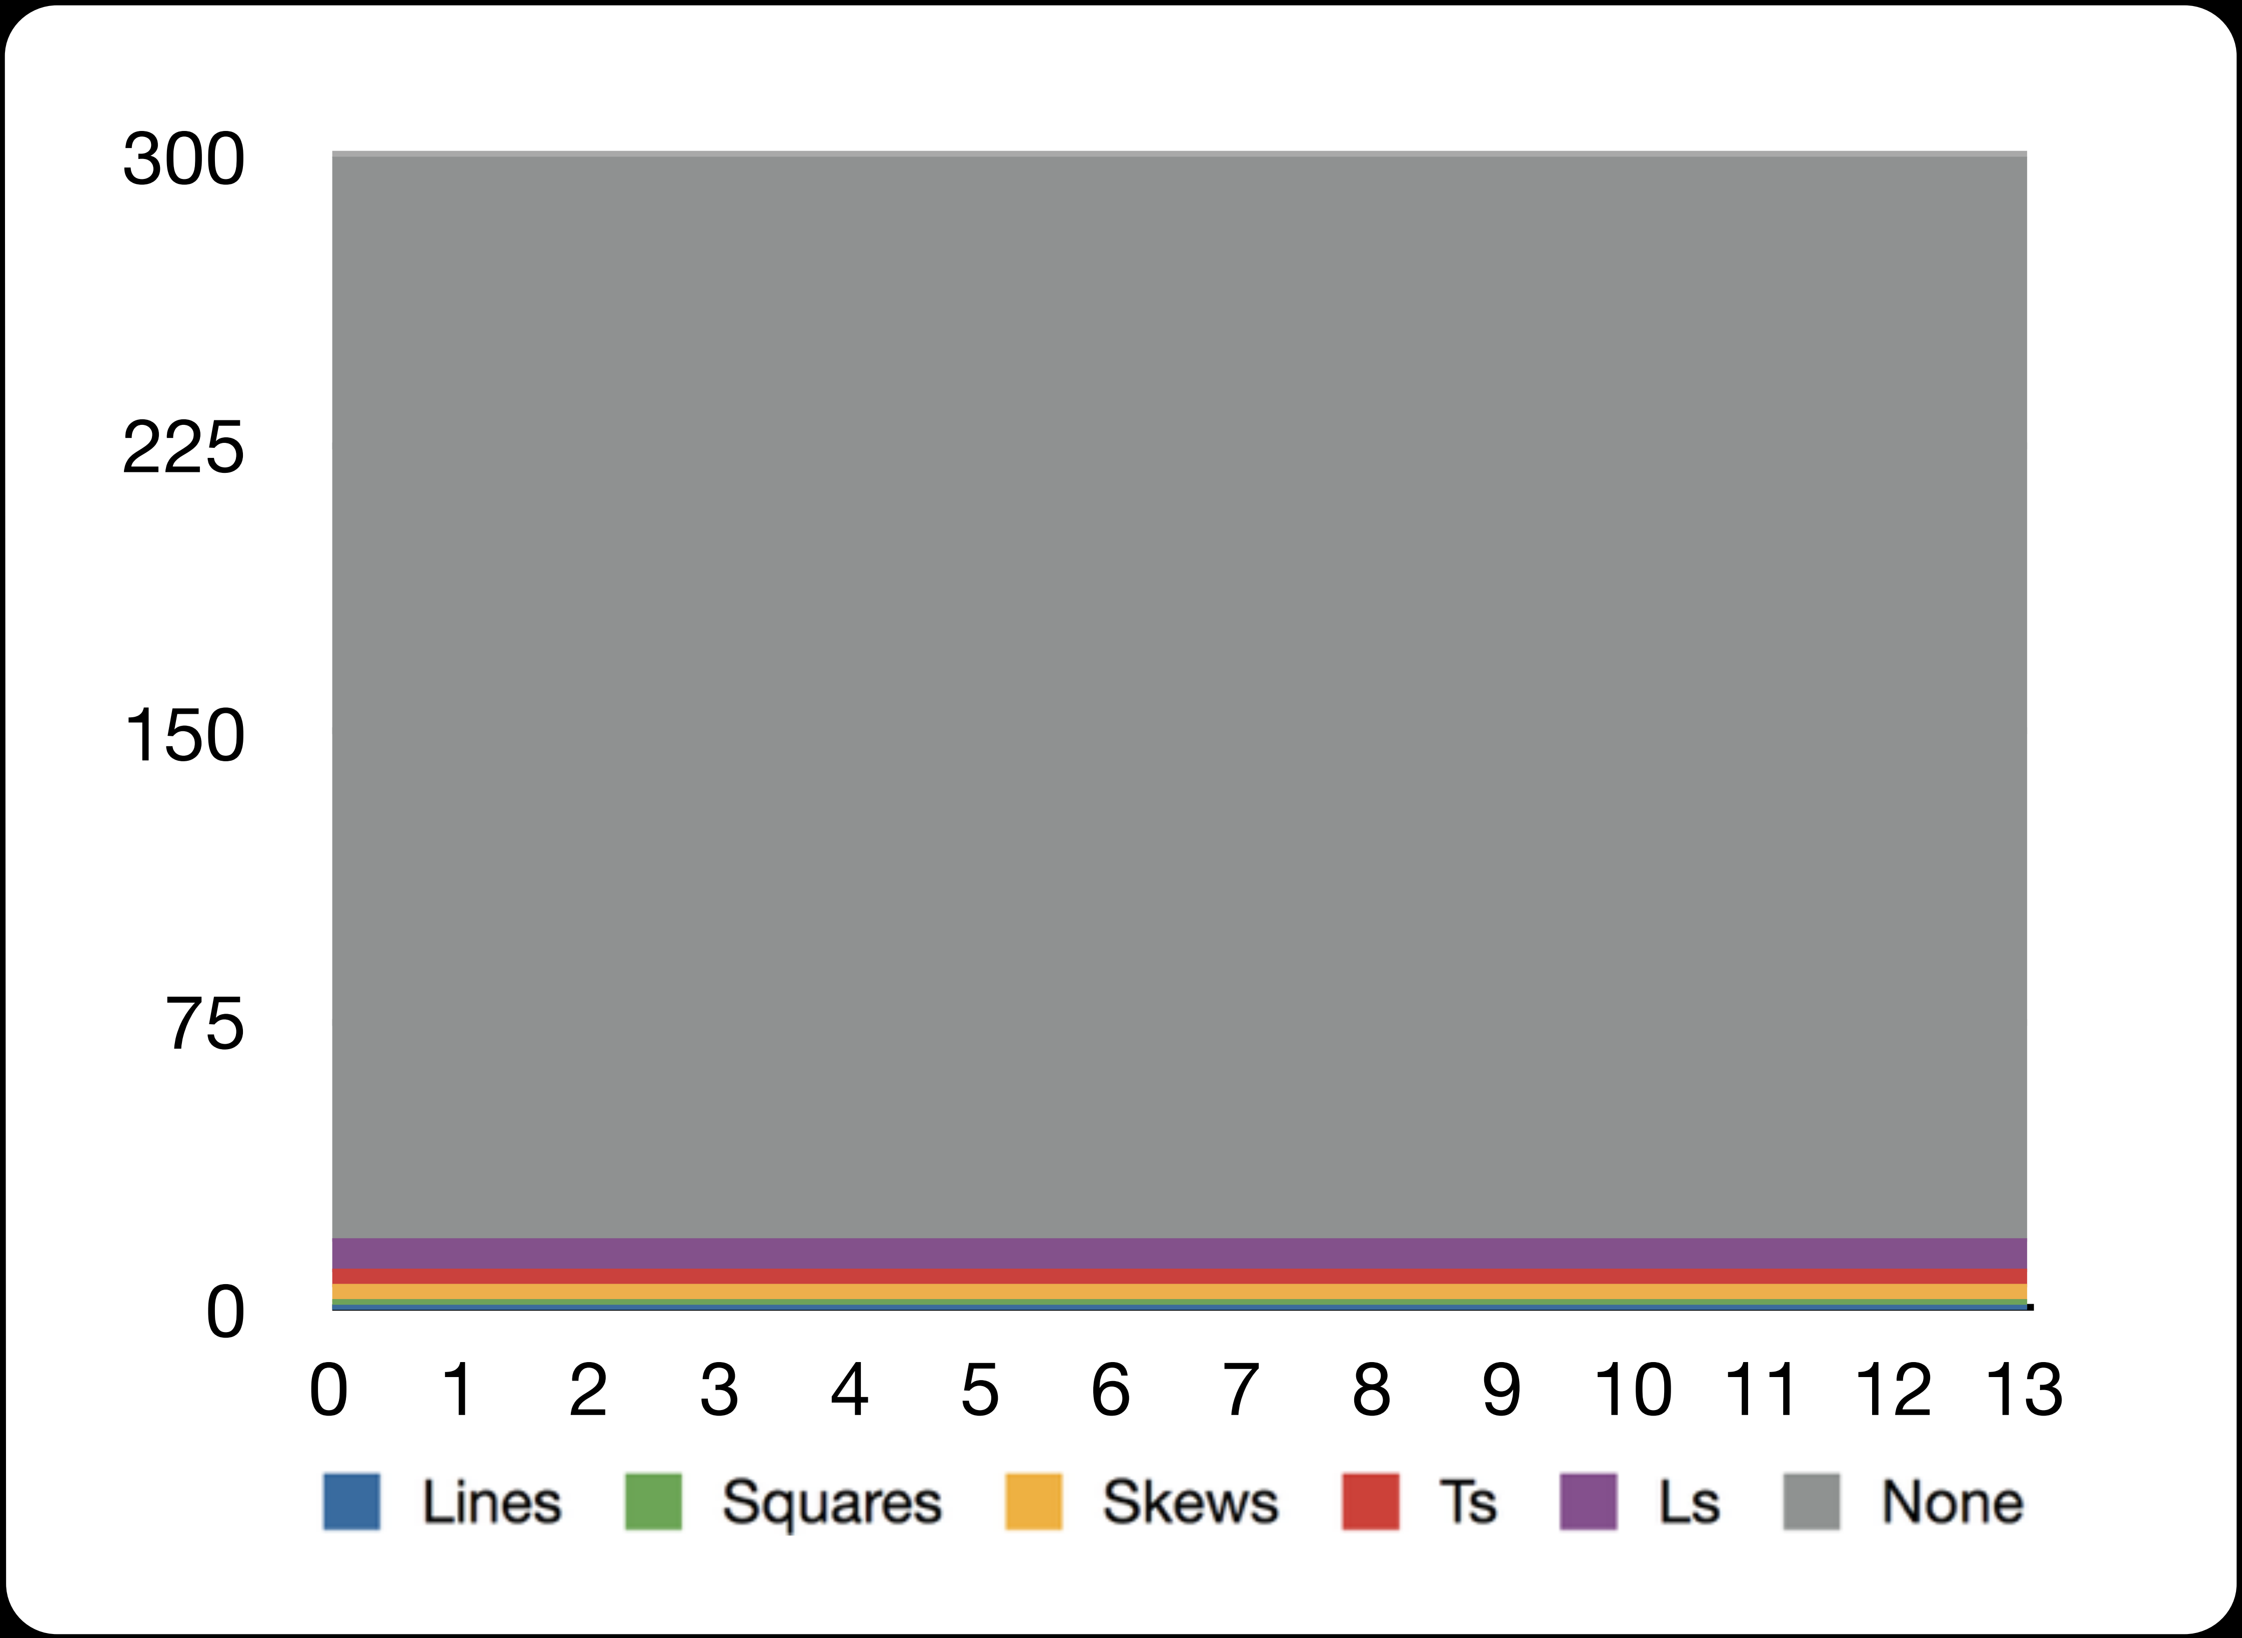

STUDY 5: ACTUAL TETROMINOS

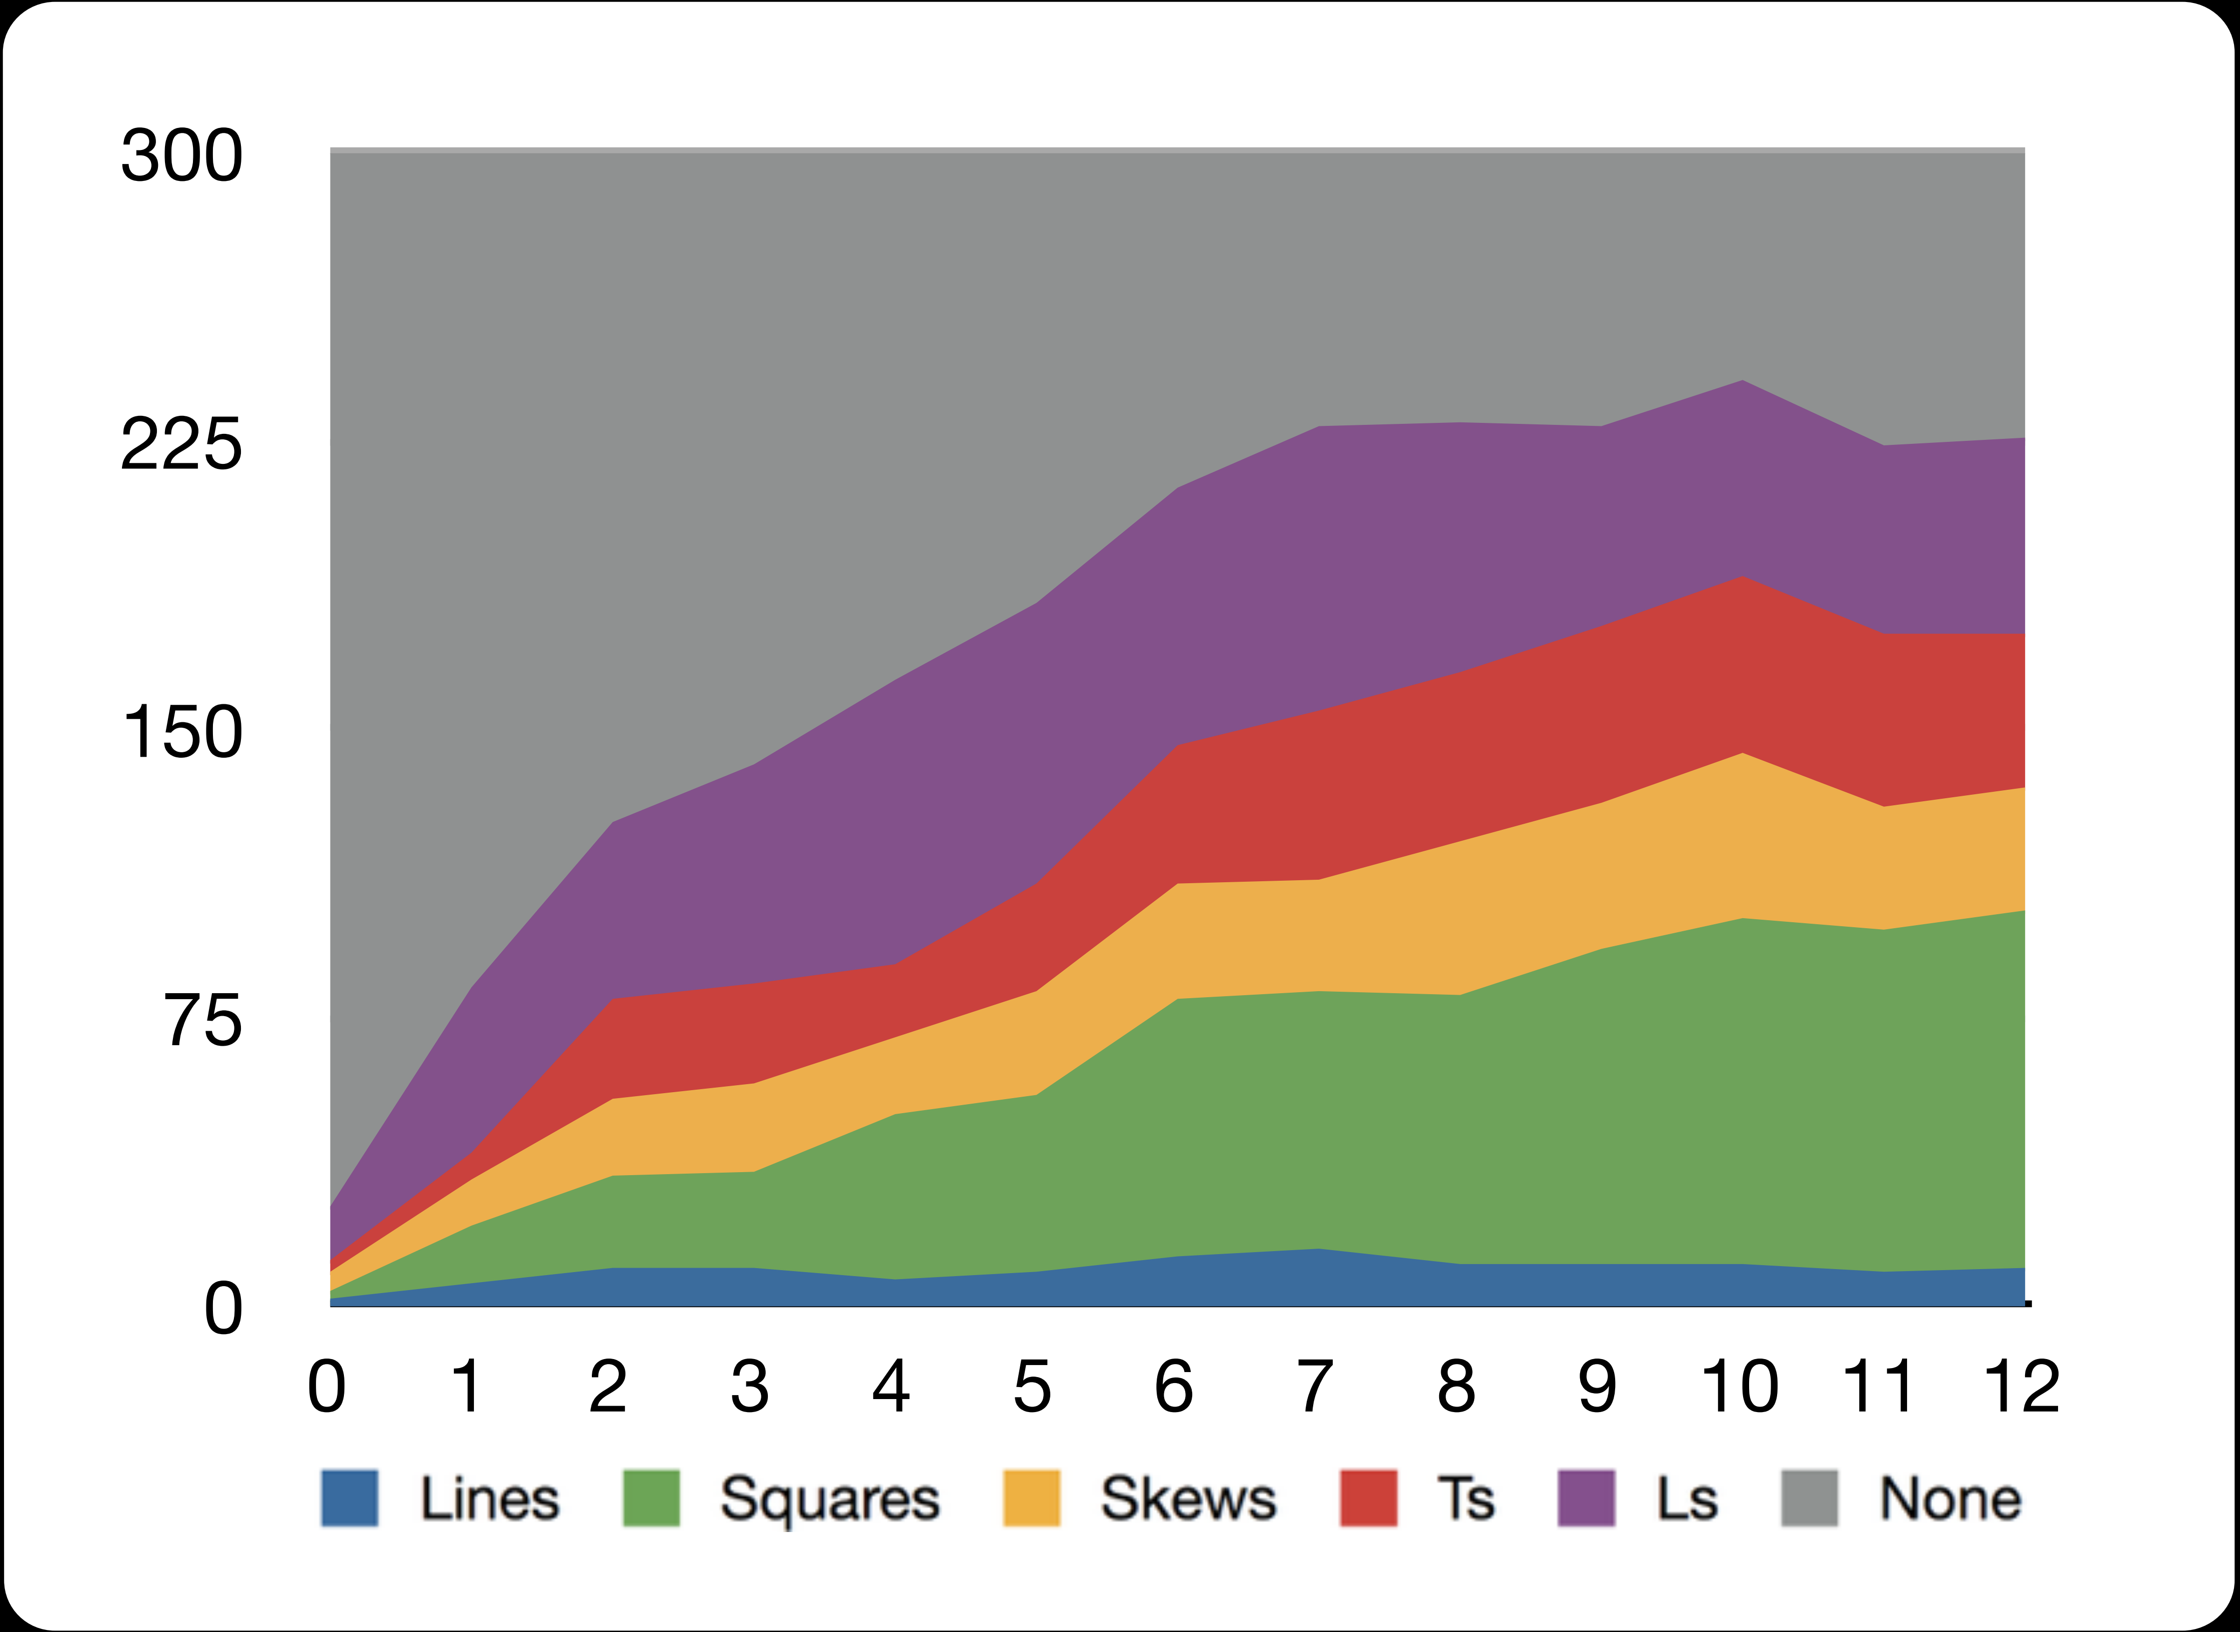

ARE TETROMINOS JUST EASIER? NO!

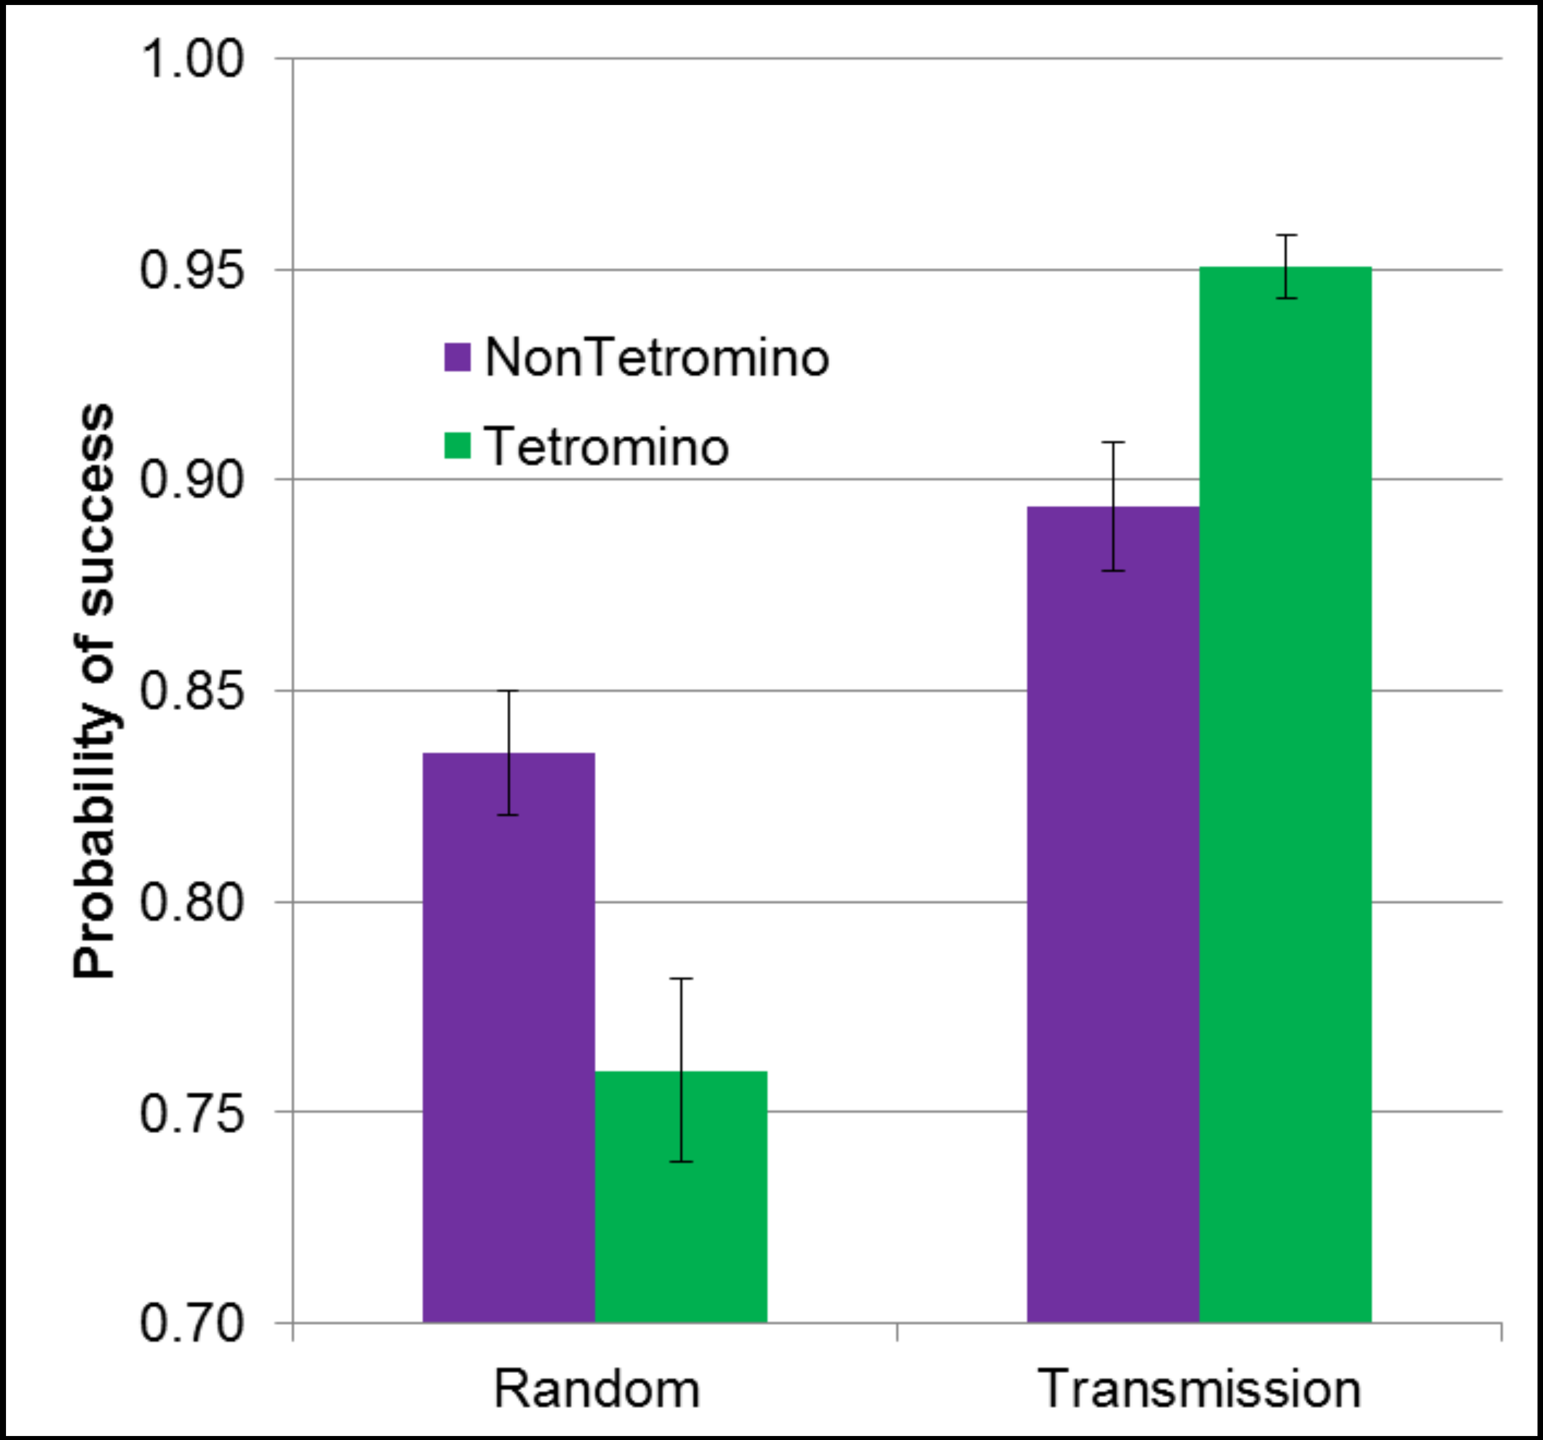

SUMMARY OF EXPERIMENTS

- **Miniature language learning**
  Emergence of compositionality with pressure from learning and expressivity. Matches simulation results well.

- **Embodied communication game**
  Exploitation of embodied behaviour to signal meaning. Would be nice to have a model of this.

- **Artificial sign language learning**
  Conventionalised systematicity. Relates well to models, but we need better understanding of iconicity.

- **Simon game**
  Emergence of a system. Need better tools for analysing structure.

- **Baboon iterated learning**
  Systematicity from transmission not species-unique. Need model to understand why tetrominos!

OVERALL SUMMARY

A RICH BEHAVIOURAL REPERTOIRE

+

TRANSMITTED THROUGH ITERATED LEARNING

→

SYSTEMATIC STRUCTURE

WHY DOES THIS WORK?

- Iterated learning of sets of behaviours involves repeated transmission through an informational bottleneck
- Behaviours adapt in order to better pass through this bottleneck
- Actual structure depends on a number of factors (including what the behaviour is used for)
- Overarching universal: **compressible** behaviours pass through the bottleneck more easily

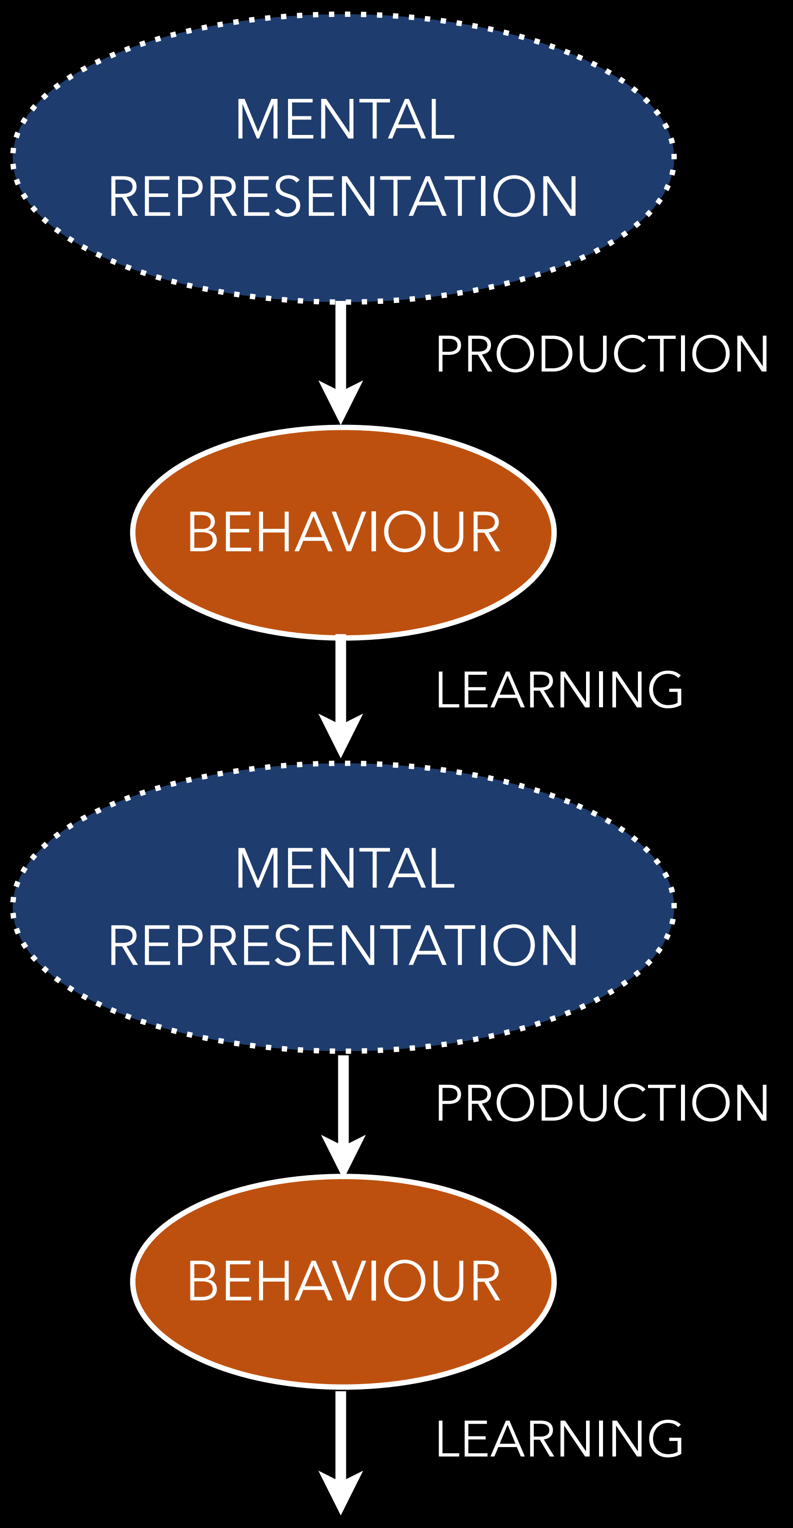

CONCLUSIONS

- Cultural transmission of sets of behaviours leads to systematic structure
- Wherever there is iterated learning of behavioural repertoires, we should find compressible, systematic structure emerging

- Ongoing work:
  extend experiments to music (with Andrea Ravignani & Tania Delgado)
  relate experimental results in humans and Zebra Finches (with Olga Feher)

Thanks!