The troubled mind: why we make problems for ourselves

Laura Schulz
MIT, ECCL
Cogsci 2018

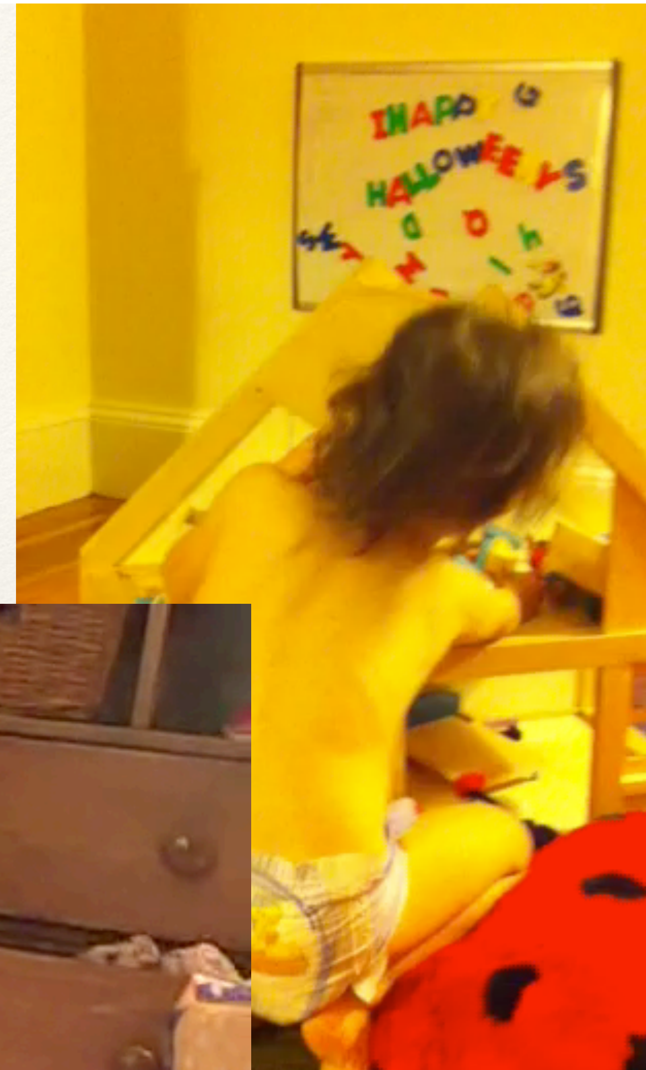

Learning as program induction

- “Coming up with the right hypotheses and theories in the first place is often much harder than ruling among them.”
- How do people, and how can machines, expand their hypothesis spaces to generate wholly new ideas, plans, and solutions?”
- “How do people learn rich representations and action plans (expressable as programs) through observing and interacting with the world?”

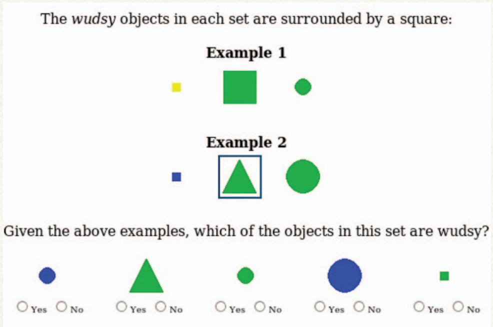

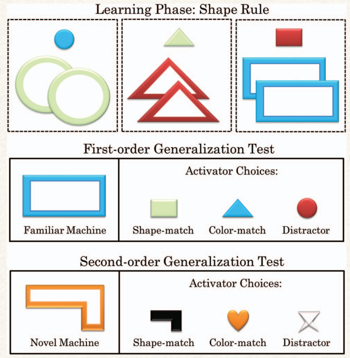

Figure 1. Schematic diagram of materials and procedure for children presented with the machines, which were activated according to a shape-match rule. See the online article for the color version of this figure.

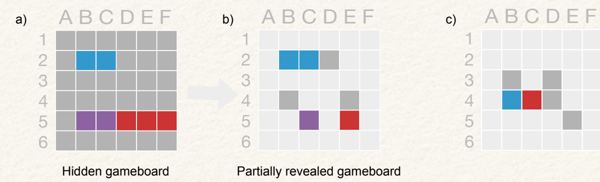

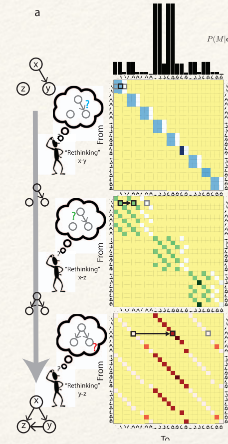

- By "using algorithms that mix stochastic recombination of primitives with memoization and compression to explain data, ask informative questions, and support one- and few-shot-inferences."
- Captures inference across a wide range of conceptual domains (logic, number; magnetism, function words, etc.)
- Not limited to symbolic relations — includes novel grounded simulations, probabilistic inferences, geometric concepts, etc.

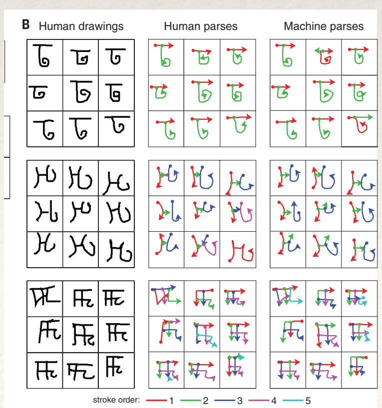

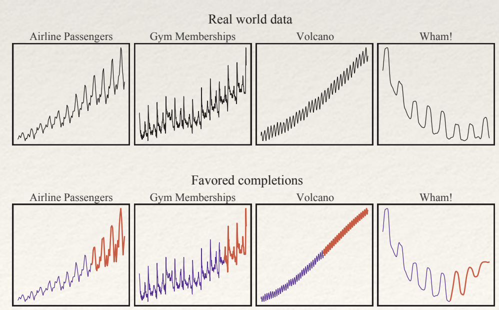

Fig. 6. (Top) Real-world data sets used in Experiment 2b. Descriptions and origin of the data were unknown to participants. (Bottom) Participants were shown the region in blue; most frequently selected completions are shown in red. Note that the periodic composition has been adapted by multiplying it with a radial basis function kernel. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

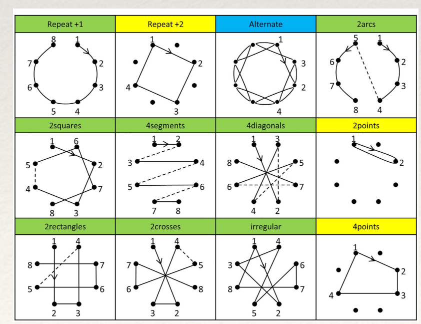

Fig 1. Paradigm. (A) Basic geometrical rules used to create sequences: rotations (+1, +2, -1, -2), axial symmetries (H: horizontal, V: vertical, A,B: oblique) and rotational symmetry (P). From one location of the octagon, each of the 7 others can be reached by the application of one or more primitives. (B) Screen shot from experiment 1. The orange dot appears at successive locations on the octagon, and subjects are asked to predict the next location. (C) Examples of sequences presented to French adults (blue), kids and Munduruku adults (yellow), or both (green).

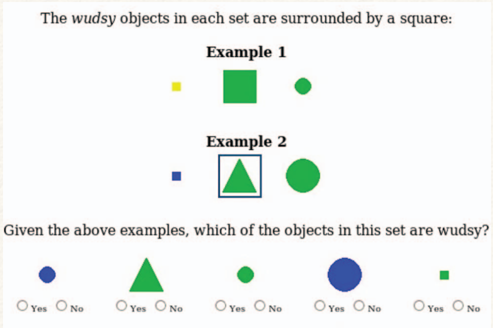

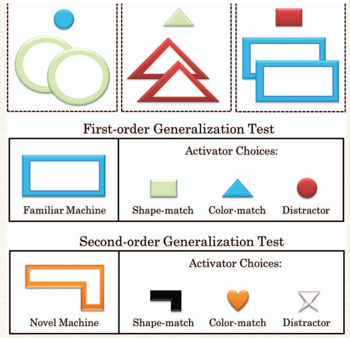

Figure 1. Schematic diagram of materials and procedure for children presented with the machines, which were activated according to a shape-match rule. See the online article for the color version of this figure.

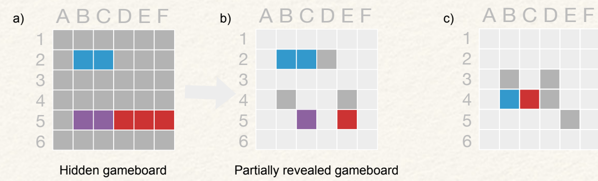

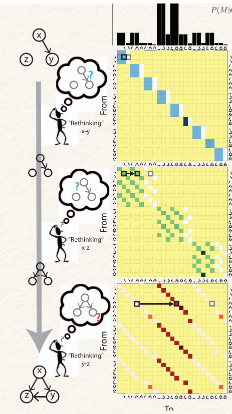

These are clearly in many respects, "wholly new ideas, plans, and solutions".

- Presumably no one has previously written in omniglot, identified the "same shape as exactly one other blue object", asked whether the purple and red ship touch …
- … considered how many times people Googled the band Wham, or predicted the next move in a geometric form across four diagonals before.

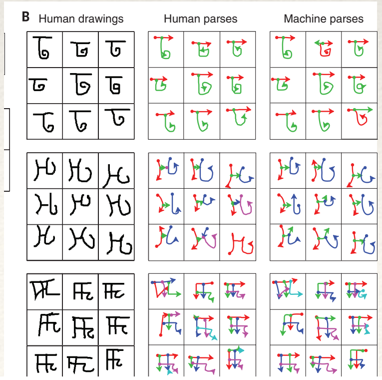

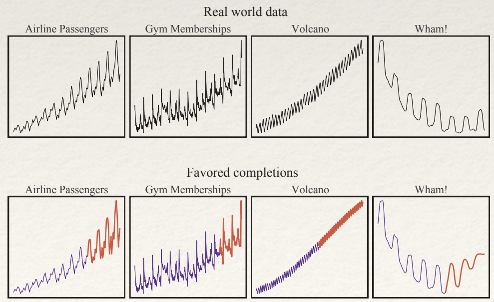

Fig. 6. (Top) Real-world data sets used in Experiment 2b. Descriptions and origin of the data were unknown to participants. (Bottom) Participants were shown the region in blue; most frequently selected completions are shown in red. Note that the periodic composition has been adapted by multiplying it with a radial basis function kernel. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

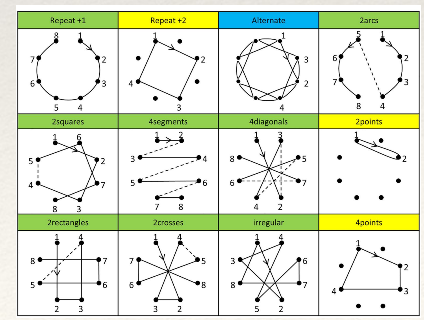

Fig 1. Paradigm. (A) Basic geometrical rules used to create sequences: rotations (+1, +2, -1, -2), axial symmetries (H: horizontal, V: vertical, A,B: oblique) and rotational symmetry (P). From one location of the octagon, each of the 7 others can be reached by the application of one or more primitives. (B) Screen shot from experiment 1. The orange dot appears at successive locations on the octagon, and subjects are asked to predict the next location. (C) Examples of sequences presented to French adults (blue), kids and Munduruku adults (yellow), or both (green).

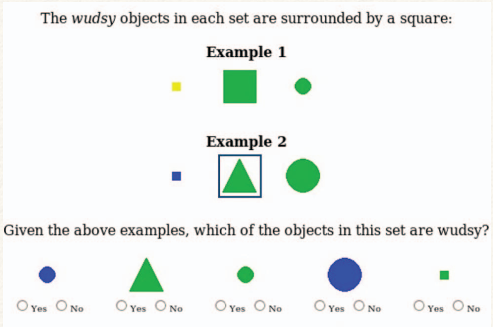

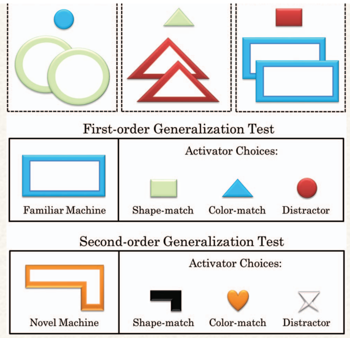

Figure 1. Schematic diagram of materials and procedure for children presented with the machines, which were activated according to a shape-match rule. See the online article for the color version of this figure.

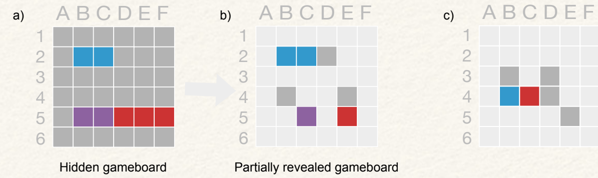

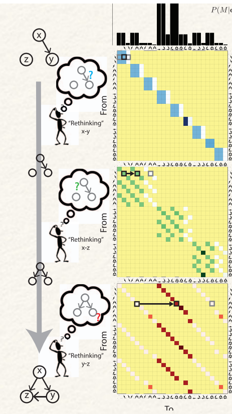

And yet while the ideas, plans and solutions may be wholly new to the learner … they are in some sense, known to the experimenter.

In each case, the training examples (even if only one or a few) are generated from the target hypothesis.

By contrast, in ordinary thought, if we are trying to think of a new idea we, by assumption, do not know the target hypothesis — so we can't rely on examples generated from it.

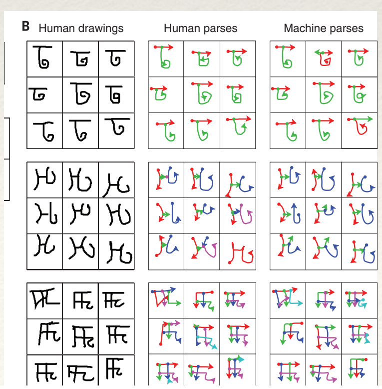

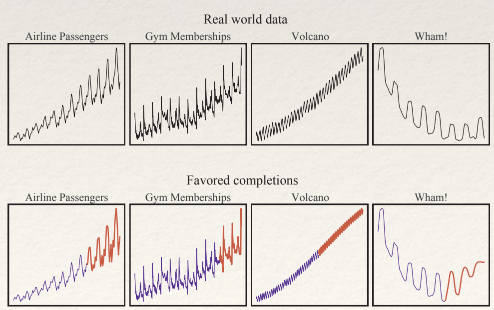

Fig. 6. (Top) Real-world data sets used in Experiment 2b. Descriptions and origin of the data were unknown to participants. (Bottom) Participants were shown the region in blue; most frequently selected completions are shown in red. Note that the periodic composition has been adapted by multiplying it with a radial basis function kernel. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

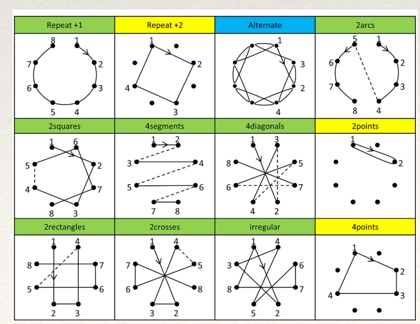

Fig 1. Paradigm. (A) Basic geometrical rules used to create sequences: rotations (+1, +2, -1, -2), axial symmetries (H: horizontal, V: vertical, A,B: oblique) and rotational symmetry (P). From one location of the octagon, each of the 7 others can be reached by the application of one or more primitives. (B) Screen shot from experiment 1. The orange dot appears at successive locations on the octagon, and subjects are asked to predict the next location. (C) Examples of sequences presented to French adults (blue), kids and Munduruku adults (yellow), or both (green).

Thinking new thoughts

- The problem of generating new ideas is not a problem about radical conceptual change or theory change.
- It is a problem of ordinary, everyday thinking: thought is productive.
- We can, quite reliably, make up new – relevant – answers to any *ad hoc* question. These answers may be trivial and they may be false but they are
  - Genuinely new, in that we didn't have them until we thought of them.
  - Genuinely made up, in that we didn't learn them from new evidence or new testimony.
  - Answers to the question. They are not non-sequiters.

Thinking new thoughts

- Why doesn't McDonald's sell hotdogs?
- How would you get chimney swifts out of your chimney?
- What's the origin of the phrase "flotsam and jetsam"?
- Who turned down the 1964 prize for literature?

We are startlingly good at generating possible solutions – to almost any problem

- We quickly converge on ideas, plans and solutions that may not be right but are, at least, wrong (as opposed to redundant, irrelevant, already known, etc.)
- Prior knowledge, a stochastic recombination of primitives, and a bias towards simplicity may *still* not be enough to explain how we come up with wholly new hypotheses and theories on the fly
- And besides, we have access to additional information we could, in principle, use …

We know a lot about our problems …

- Long before we can solve our problems or achieve our goals we may have some sense of …
  - How hard the problem is
  - What might count as an answer or solution
  - What might be desirable in an answer or solution

We know a lot about our problems …

- Long before we can solve our problems or achieve our goals we may have some sense of …
  - **How hard the problem is**
  - What might count as an answer or solution
  - What might be desirable in an answer or solution

Maddie Pelz: Thursday 4:20-:4:40

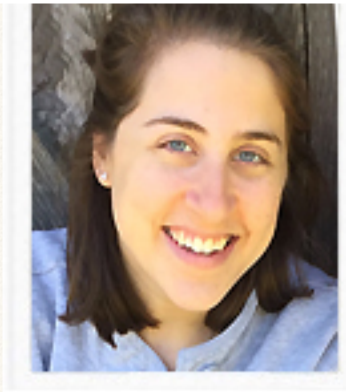

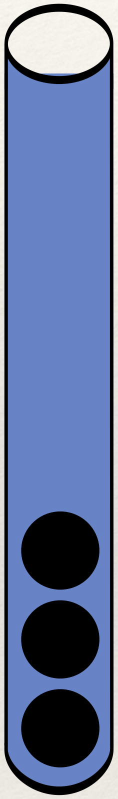

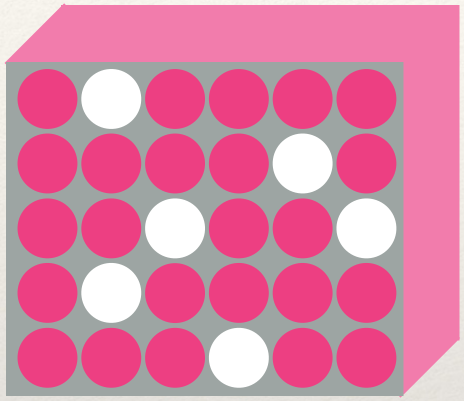

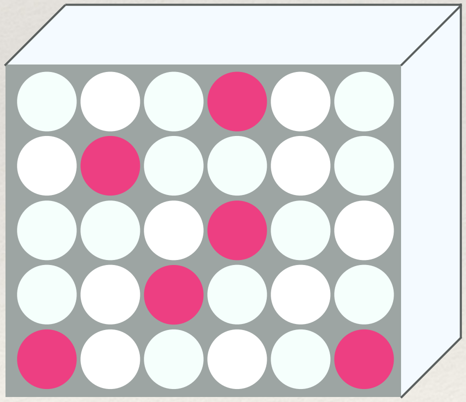

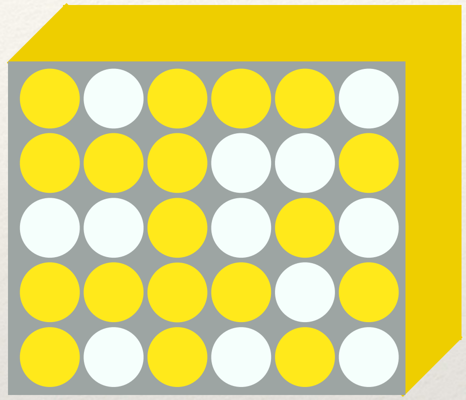

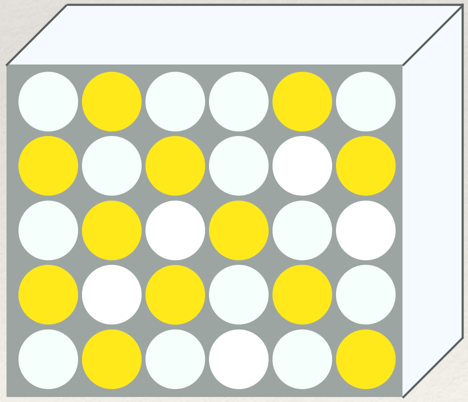

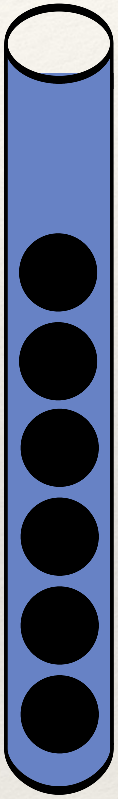

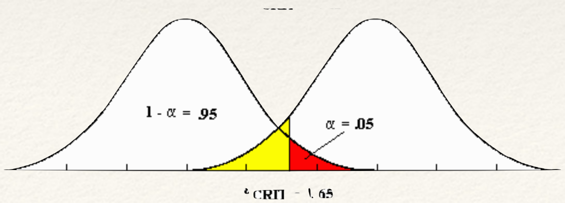

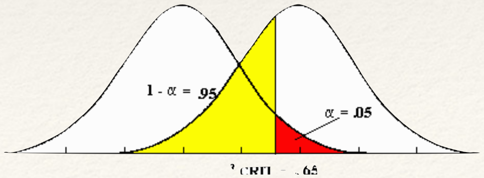

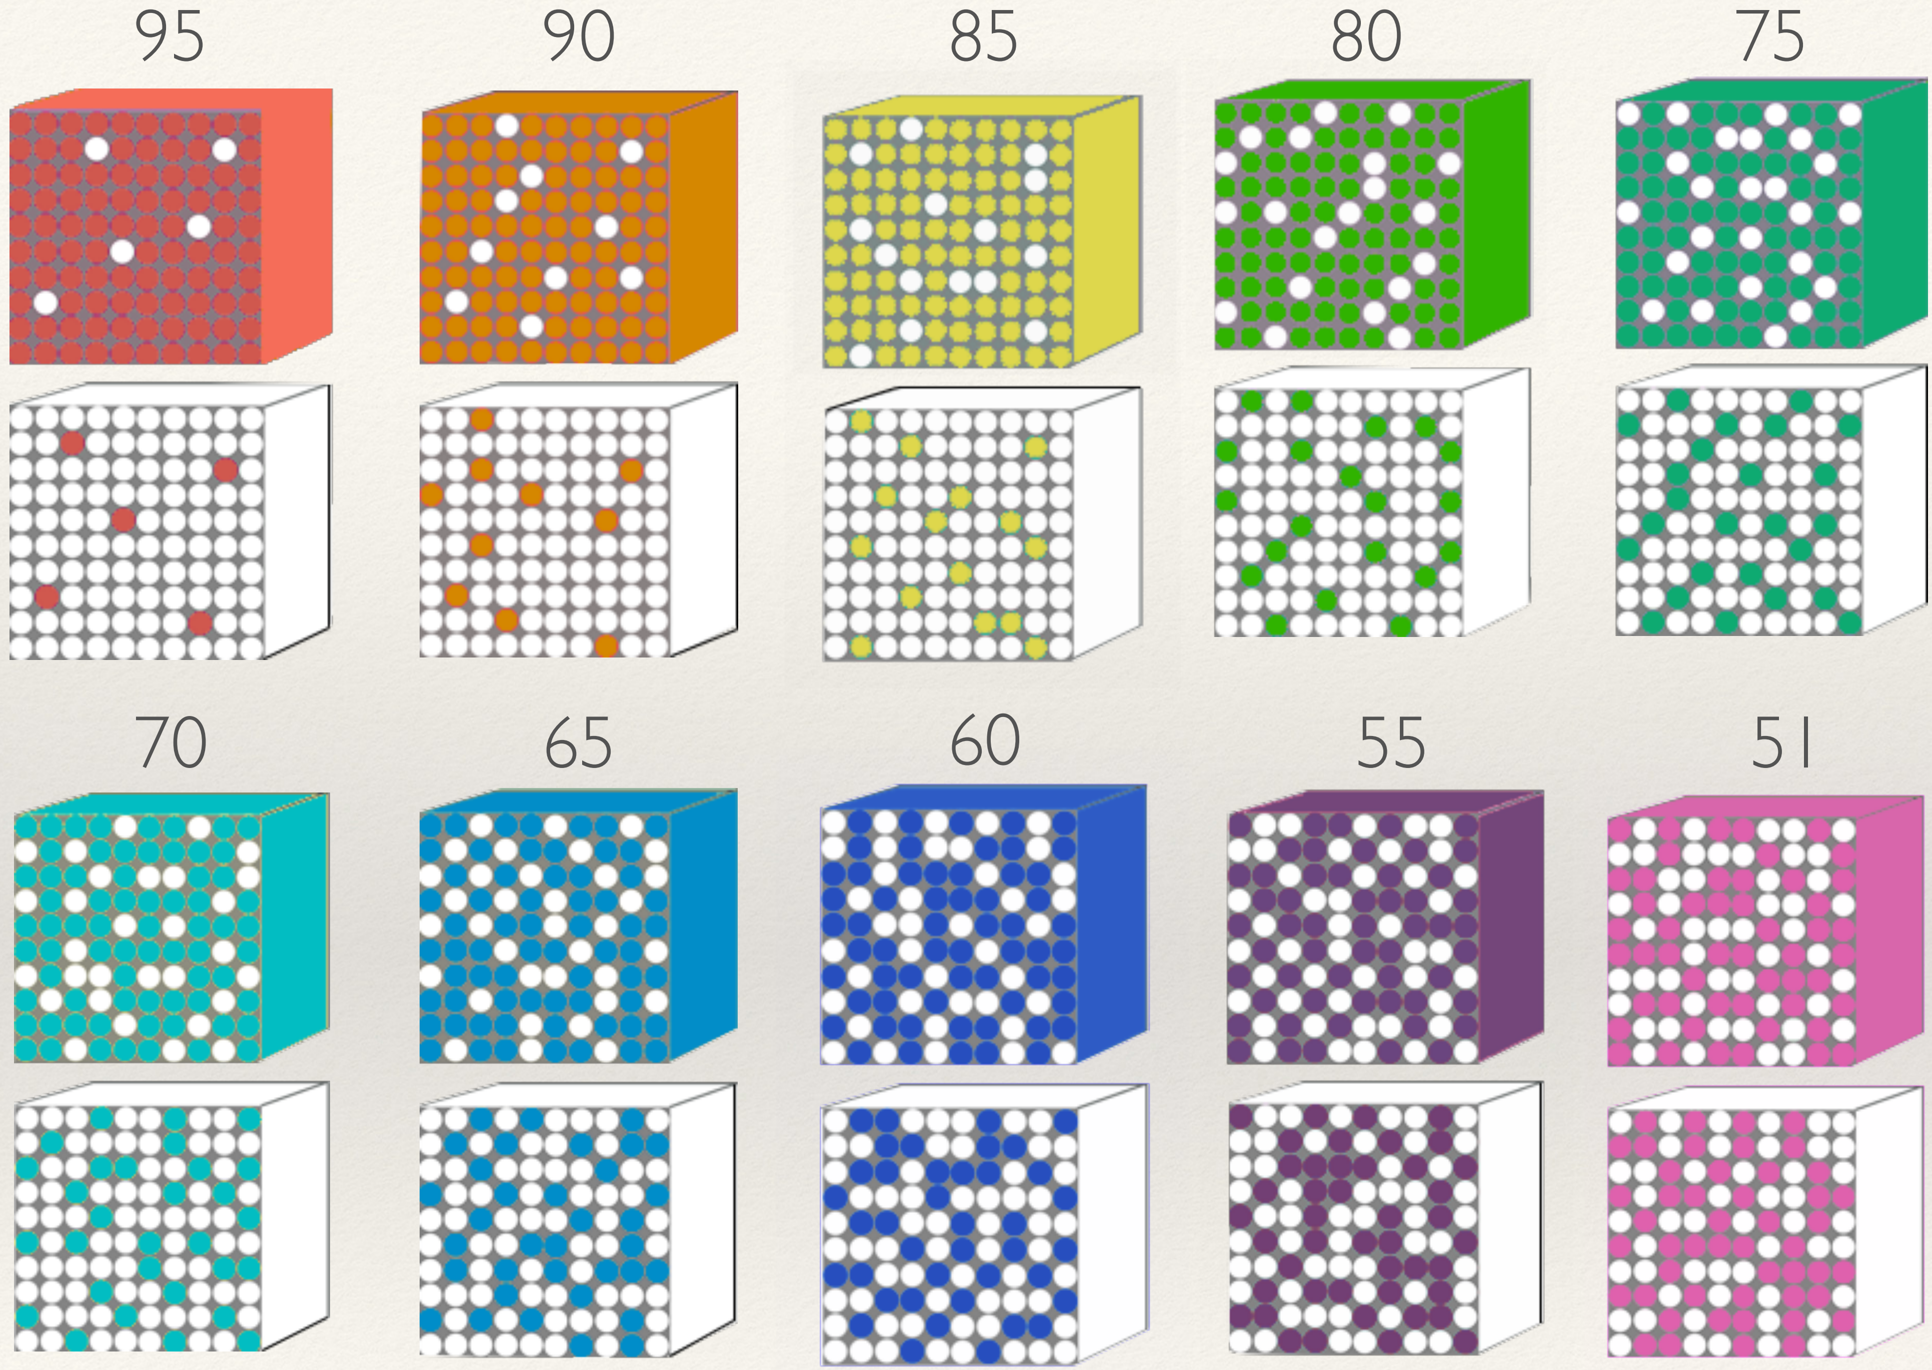
95
90
85
80
75
70
65
60
55
51

Intuitive power analyses

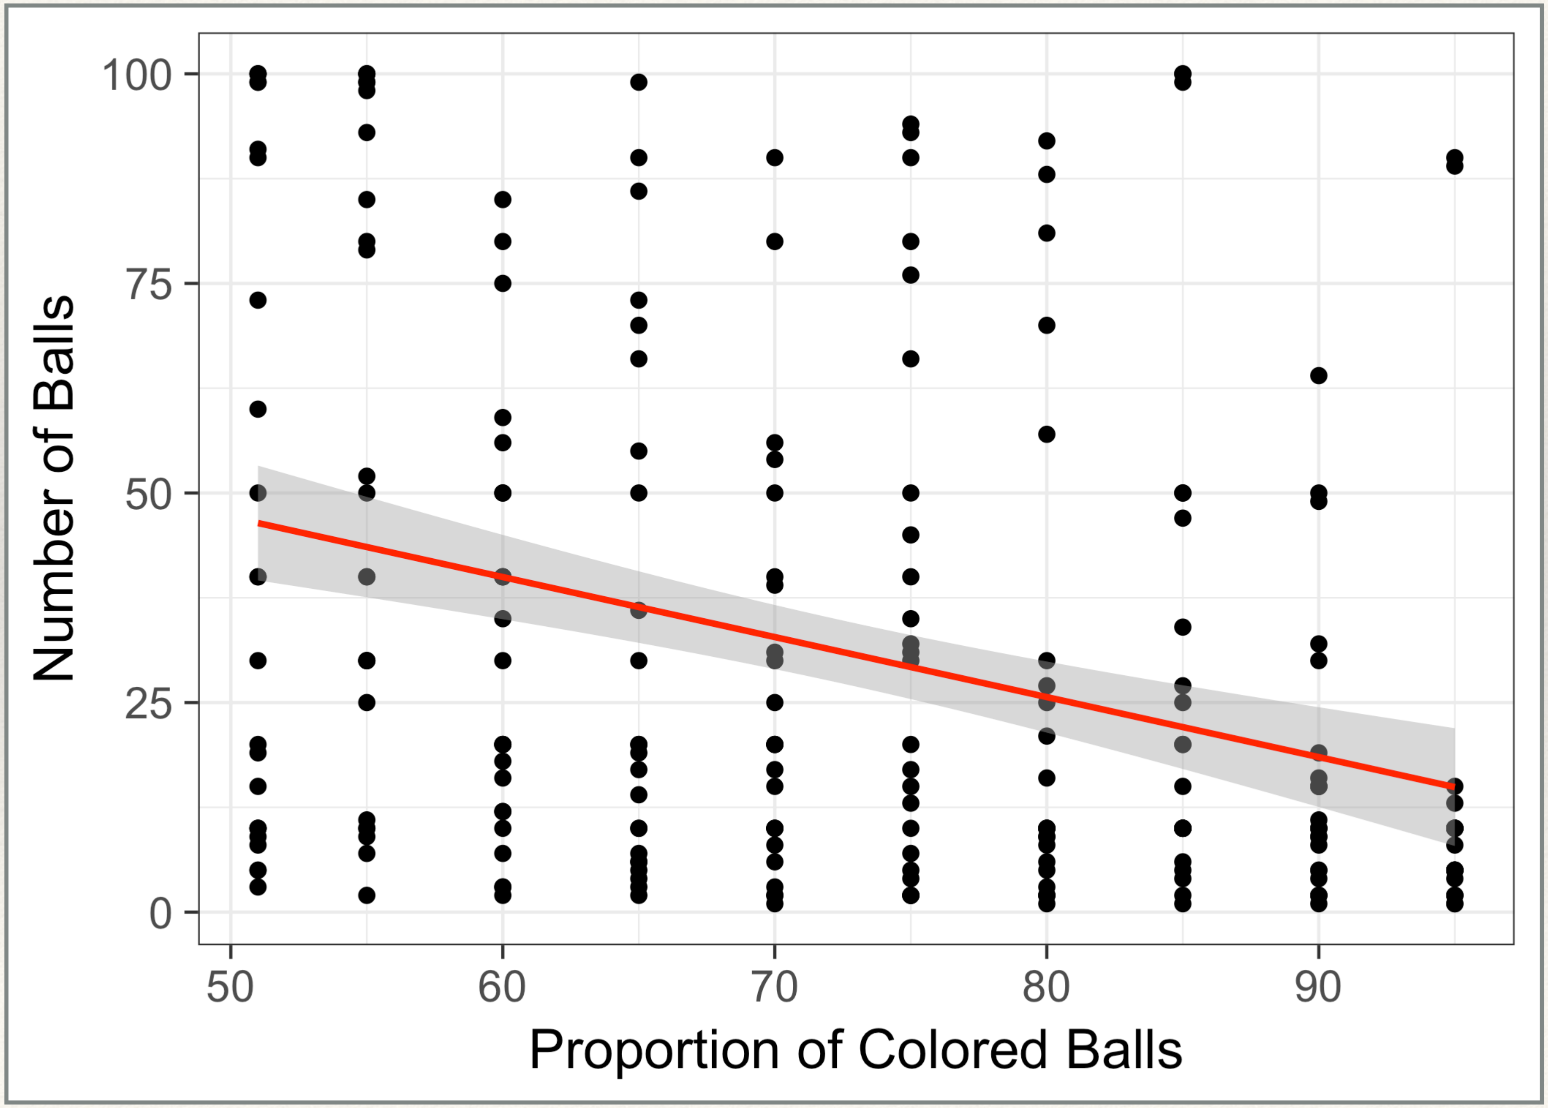

Children ask for more data for harder problems

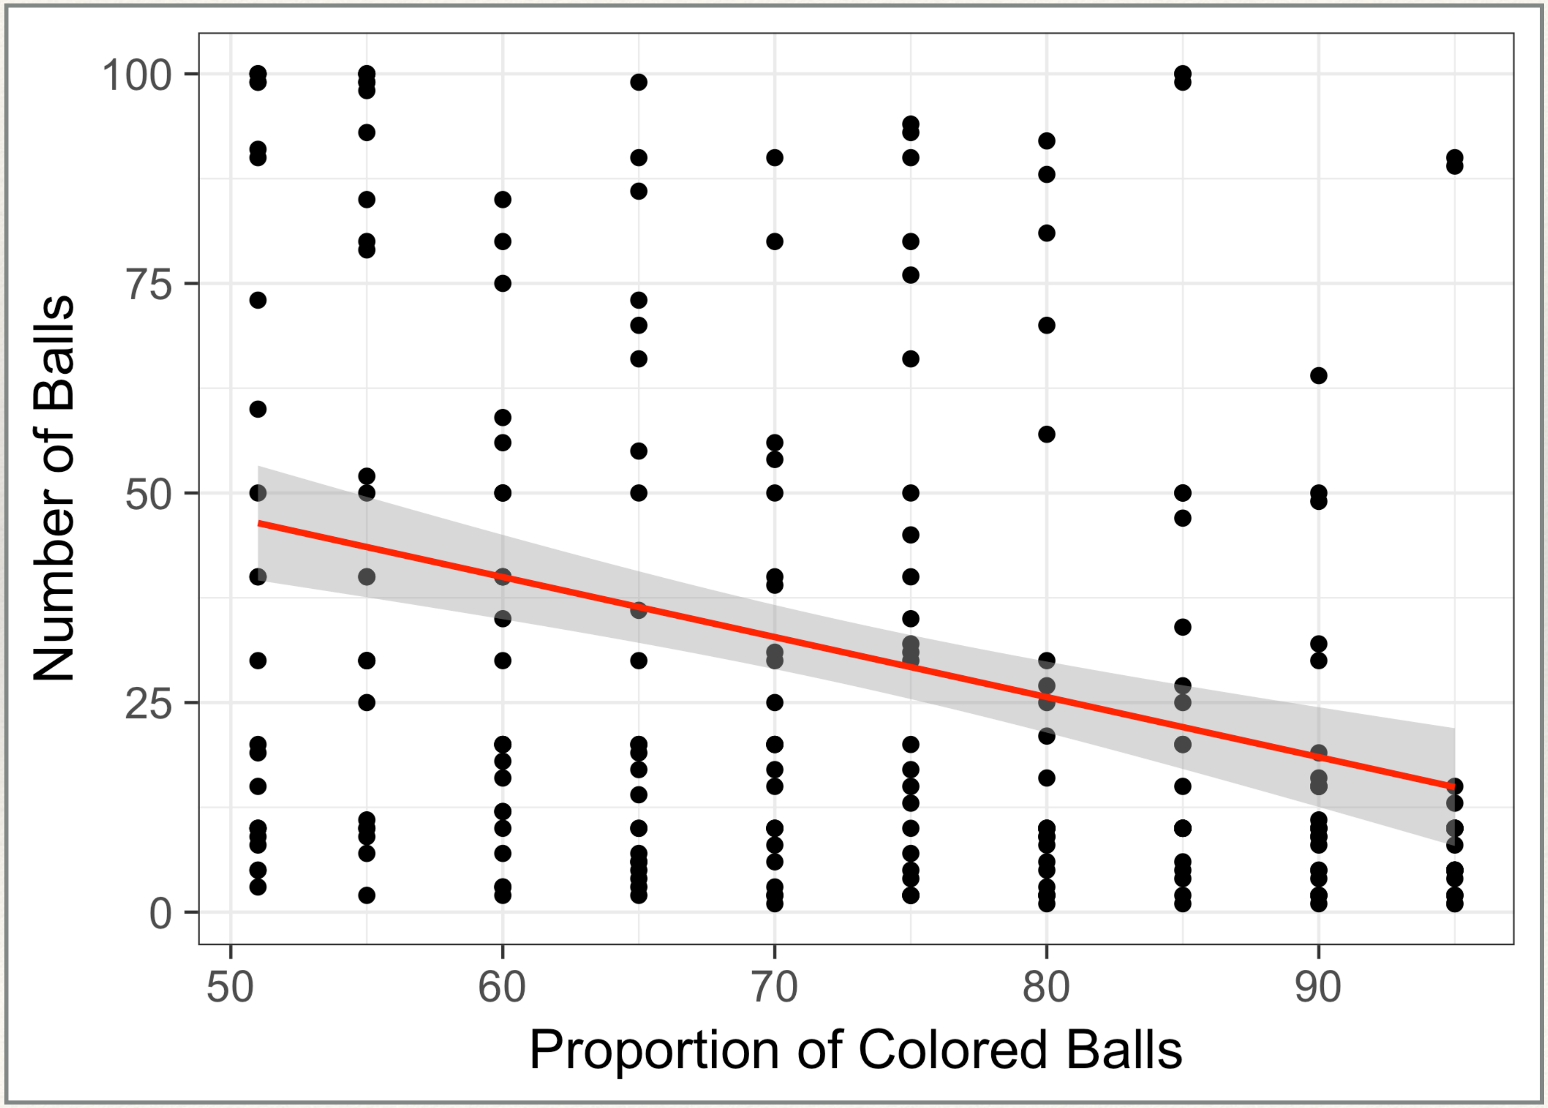

We know a lot about our problems …

- Long before we can solve our problems or achieve our goals we may have some sense of …
  - **How hard the problem is**
  - What might count as an answer or solution
  - What might be desirable in an answer or solution

We know a lot about our problems …

- Long before we can solve our problems or achieve our goals we may have some sense of …
  - **How hard the problem is**
  - **What might count as an answer or solution**
  - What might be desirable in an answer or solution

Abstract form of cause and effect

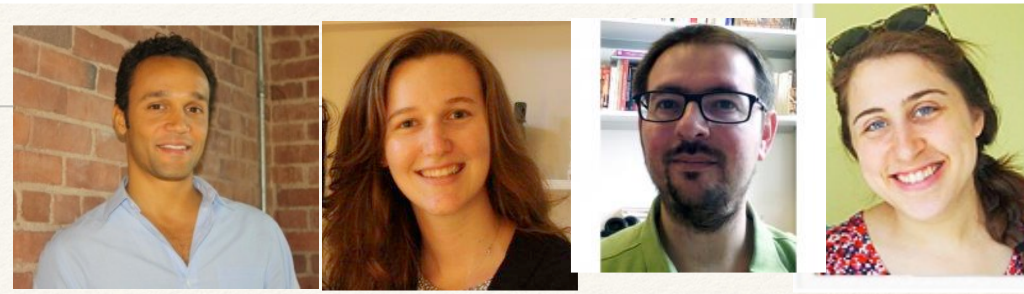

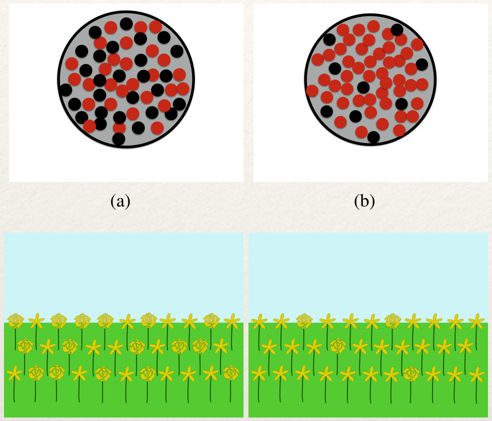

Relative proportion

Discrete vs. continuous

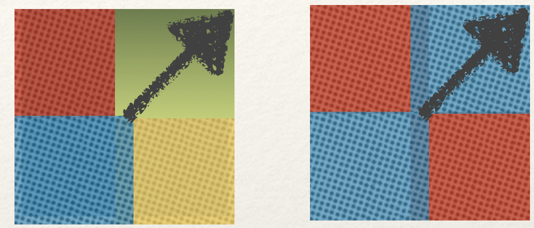

Arity (how many states the cause and effect occupy

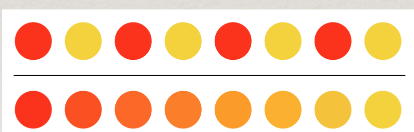

Alternation versus monotonic change

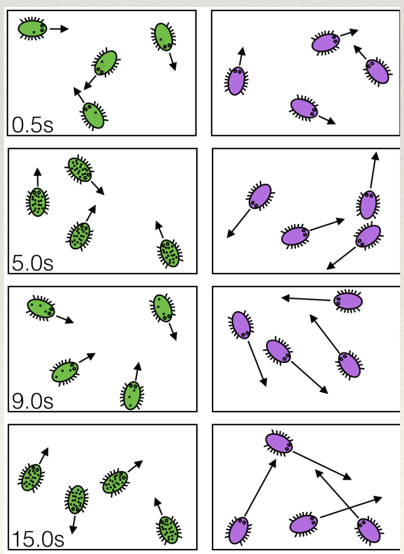

Figure 4: Periodic (green) and monotonic (purple) bugs from stimulus set 2. The green bugs move at constant speed but increase and decrease in number of spots; the purple bugs

Rate of change (fast or slow); cyclic vs. acyclic; exponential vs. linear, etc.

But not just relationship of causes to effects relationships of problems to solutions broadly

Problems are rich in all kinds of information

- Consider the information contained in question words (even before we get to the content of the questions) …

Who?

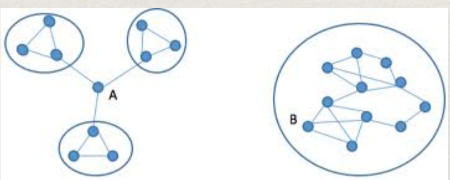

Where?

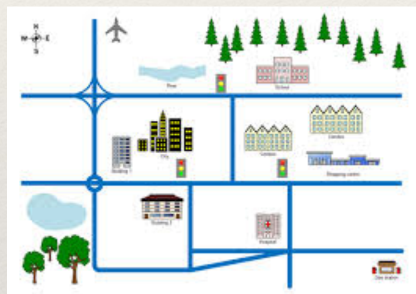

When?

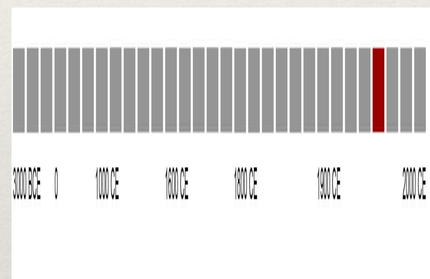

What?

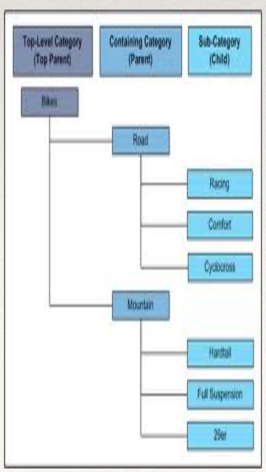

Which?

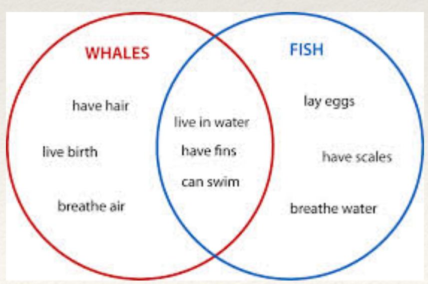

How?

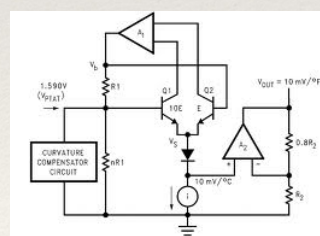

Why?

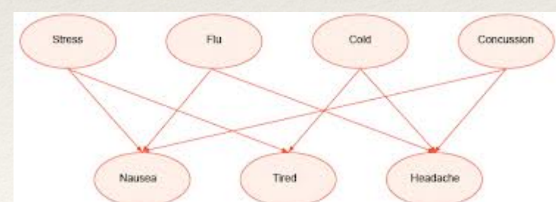

Why does …?

this rule or empirical generalization hold?

…

Why did …?

this unexpected event occur?

…

Why can't …?

some deviation from a rule or generalization occur?

…

Why did she…?

engage in some unexpected thought or intentional action?

…

Why did Trump …?

it's a rant

…

Why did the chicken …?

it's a joke

…

We know a lot about our problems …

- When we do not have an abstract representation of what might count as a solution to a problem we may resort to inefficient and often ineffective searches.
  - Indeed, what it might mean for us to think that a problem is "tractable" or "well-posed" might be to recognize that we don't know the answer the a problem …
  - but the problem does contain enough information to guide the search.

We know a lot about our problems …

- Our ability to represent what "counts" as a solution to a problem before we know what the solution is might explain how:
    - We can have a sense of "being on the right track" well before we can better account for the data.
    - We can think an idea is a great idea – even when we know it is wrong.
- We may be able to constrain our proposals on two separate dimensions:
    - how well they fit the data: "TRUTH"
    - how well they would solve our problems if they were true: "TRUTHINESS"

We know a lot about our problems …

- Long before we can solve our problems or achieve our goals we may have some sense of …
  - **How hard the problem is**
  - **What might count as an answer or solution**
  - **What might be desirable in an answer or solution**

Junyi Chu: Thursday 4-:4:20

Cognitive Pragmatism

- Abundant research suggests children will endorse known, factual, reliable, verified information over uncertain, speculative, unreliable, unverified information.
- But when known, factual, reliable, verified information fails to solve our problems or achieve our goals, we may need to reject it in favor of speculative conjectures —
- that may not have the virtue of being (currently) knowably true, but at least have the virtue of providing answers to our problem.

Cognitive Pragmatism

Here are some small
Daxes and
some big Blickets.
The Big Blickets made hats for
the small Daxes

Question with known answer: Why are the small Daxes wearing hats?

A) Because the Big Blickets made hats for the small Daxes

B) Because the Big Blickets are older than the small Daxes

Cognitive Pragmatism

Here are some small Daxes and some big Blickets.
The Big Blickets made hats for the small Daxes

Question with unknown answer: Why are the Blickets bigger than the Daxes?

A) Because the Big Blickets made hats for the small Daxes

B) Because the Big Blickets are older than the small Daxes

Cognitive Pragmatism

We may even accept conjectures that **contradict** known "facts" if the conjecture provides a possible solution to our problem

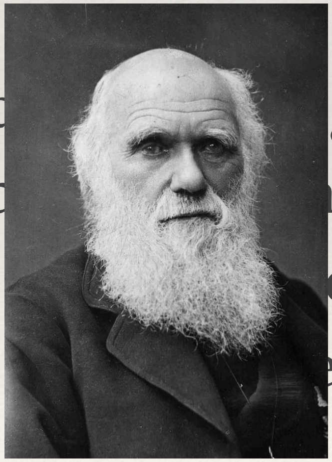

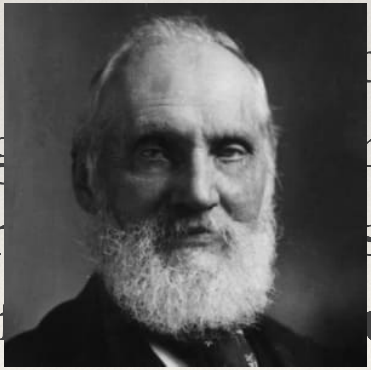

Inc gree olution
sup ble g ncrease
od th se it
ent o ue.

probability — or utility?

- Sally is a counselor at a children's summer theater camp. She has to shout a lot to be heard over the kids. She has had a sore throat all week.
- She turns on the news and hears about a new virus — V1-09. Fifteen people have been hospitalized with it so far. A sore throat is one of its symptoms.
- Diagnosis: There's a blood test available that can diagnose the presence of V1-09 with 98% accuracy.
- Intervention: There's a new medication available that's now being sold at drug stores nationwide.

Sore throat from yelling all day

New V1-09 Virus

Many factors affect the utility of a proposal — and these could be used to guide the construction of new programs, not just their evaluation

Josh Rule

Accuracy is so important that solutions with low accuracy hardly count as solutions at all.

Concision reduces the chance of errors and the cost to discover and store a solution.

Efficiency respects limits in time and computational power that slow users from solving their many problems.

Generality lets a few solutions apply to many problems, reducing the costs of storing many distinct solutions.

Modularity breaks a system at its semantic joints into composable parts that can be optimized and reused independently.

Reusability reduces the total solution complexity with partial solutions that can be reused to solve many problems.

Elegance by way of symmetry and minimalism is common among mature solutions and signals that each component plays a non-trivial role in the solution.

Clarity makes a program easier to learn and explain while also revealing the essential structure of the problem, which may lead to further improvements.

Robustness allows solutions to degrade gracefully, recover from errors, and accept many input formats, increasing the user's ability to focus on other problems.

Cleverness allows a problem solver to discover solutions to otherwise unsolvable problems.

Figure 2: A list of traits common to good programs.

We know a lot about our problems …

- Long before we can solve our problems or achieve our goals we may have some sense of …
  - **How hard the problem is**
  - **What might count as an answer or solution**
  - **What might be desirable in an answer or solution**

Why do we have so many problems?

We populate the world with problems of our own making— we want to end poverty, cure cancer, write the Great American novel, achieve enlightenment, eat more hot dogs than anyone else …

Why do we have so many problems?

- Maybe it's not that we're smart enough to generate new problems and goals …
- Maybe it's that having problems and goals is what allows us to be smart …
- They constrain the search space
- And the value of the solutions we generate may far exceed the generalizability or merits of any given problem or goal.

Learning as program induction

- How do people, and how can machines, expand their hypothesis spaces to generate wholly new ideas, plans, and solutions?”
- “How do people learn rich representations and action plans (expressable as programs) through observing and interacting with the world?
- Not only by “using algorithms that mix stochastic recombination of primitives with memoization and compression …”
- But also by using the information in our problems to bootstrap our ways towards solutions.

Massachusetts Institute of Technology

brain+cognitive sciences

Thanks!

BOSTON CHILDREN'S MUSEUM

MIT EARLY CHILDHOOD COGNITION LAB
PlayLab
AT BOSTON CHILDREN'S MUSEUM

NSF

The Center for Brains, Minds & Machines