CUMULATIVE LEARNING WITH CAUSAL-RELATIONAL MODELS

Kristinn R. Thórisson[1,2] & Arthur Talbot[1,3]

1 *Center for Analysis and Design of Intelligent Agents, Reykjavik U.*
2 *Icelandic Institute for Intelligent Machines, Reykjavik, Iceland*
3 *Ecole normale supérieure Paris-Saclay, France*

cadia.ru.is

JOINT MULTI-CONFERENCE ON HUMAN-LEVEL ARTIFICIAL INTELLIGENCE
ARTIFICIAL GENERAL INTELLIGENCE | PRAGUE | 24 August 2018

ICELANDIC INSTITUTE FOR INTELLIGENT MACHINES
VITVÉLASTOFNUN ÍSLANDS
www.iiim.is

Key Topics

- Models & Modeling
- Environment, variables, relations
- **Cumulative Modeling**
- **Cumulative Learning**

Long-Term Focus

- How to create a machine with general intelligence

Long-Term Focus

- How to create a machine with general intelligence
- Controller + body that
  - achieves complex goals
  - in novel task-environments

Complex Task-Environment

- Large number of variables, relations, and transformation functions
- Novel states are common

Complex Task-Environment

- Large number of variables, relations, and transformation functions
- Novel states are common

Current Research Question

- What kind of information structures are needed for a controller to deal with novel real-world environments?

Task-Environment

- $E = \{V, F, R, S_0\}$
  - $V = \{v_1, v_2, \ldots v_{||V||}\}$
  - $F = \{f_1, f_2, \ldots f_n\}$
- Regularity, dynamics

Task-Environment

- $E = \{V,F,R,S_0\}$
  - $V = \{v_1, v_2, \ldots v_{||V||}\}$
  - $F = \{f_1, f_2, \ldots f_n\}$
- Regularity, dynamics
  - Variables
    - observable
    - manipulatable

Agent Controller

- E = {V,F,R,S_0}
  - V = {v_1, v_2, … $v_{||V||}$}
  - F = {f_1, f_2, … f_n}

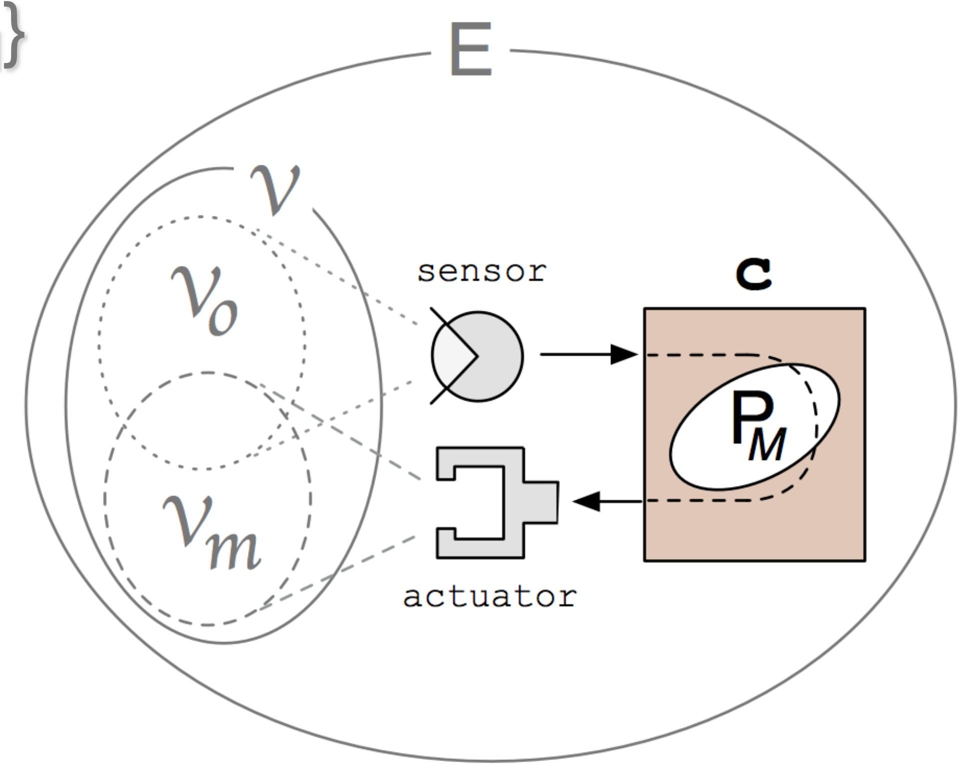

- Controller implements a *Process* that models relations between manipulated and observed variables

Requirements on Controller

- Plans are necessary for achieving goals
- Controller must <u>learn</u> to identify the variables relevant for the job, *on the job*
- Reasoning – abduction, deduction, generalization – are necessary tools due to <u>limited resources</u>
- **Learning:** *Adaptation through experience in service of goal-seeking*
  - …in a way that meets the above requirements

Requirements on Controller

- Automatic generation and revision of knowledge,
- supporting on-job learning and reasoning.
- What kinds of knowledge structures?

Requirements on Controller

- Automatic generation and revision of knowledge,
- supporting on-job learning and reasoning.
- <u>What kinds of knowledge structures?</u>

Fundamentals

- **Good Regulator Theorem**
  - every good controller of a system *must* capture a model of that system

    R. C. Conant and W. R. Ashby (1970). Every good regulator of a system must be a model of that system. International Journal of Systems Science 1(2), 89–97.

- “good”: the controller achieves its goals, as measured in light of its specification

On-Job Modeling: CRMs

- Explicit causal-relational models

Causal-Relational Models

- CRMs are
  - created based on experience
  - Capture relations between observed variables
  - Model relations between variables and elements in *E*

$$\alpha_{t_1} \text{—} \ \beta_{t_{1+d}}$$

HUMANOBS project AERA system

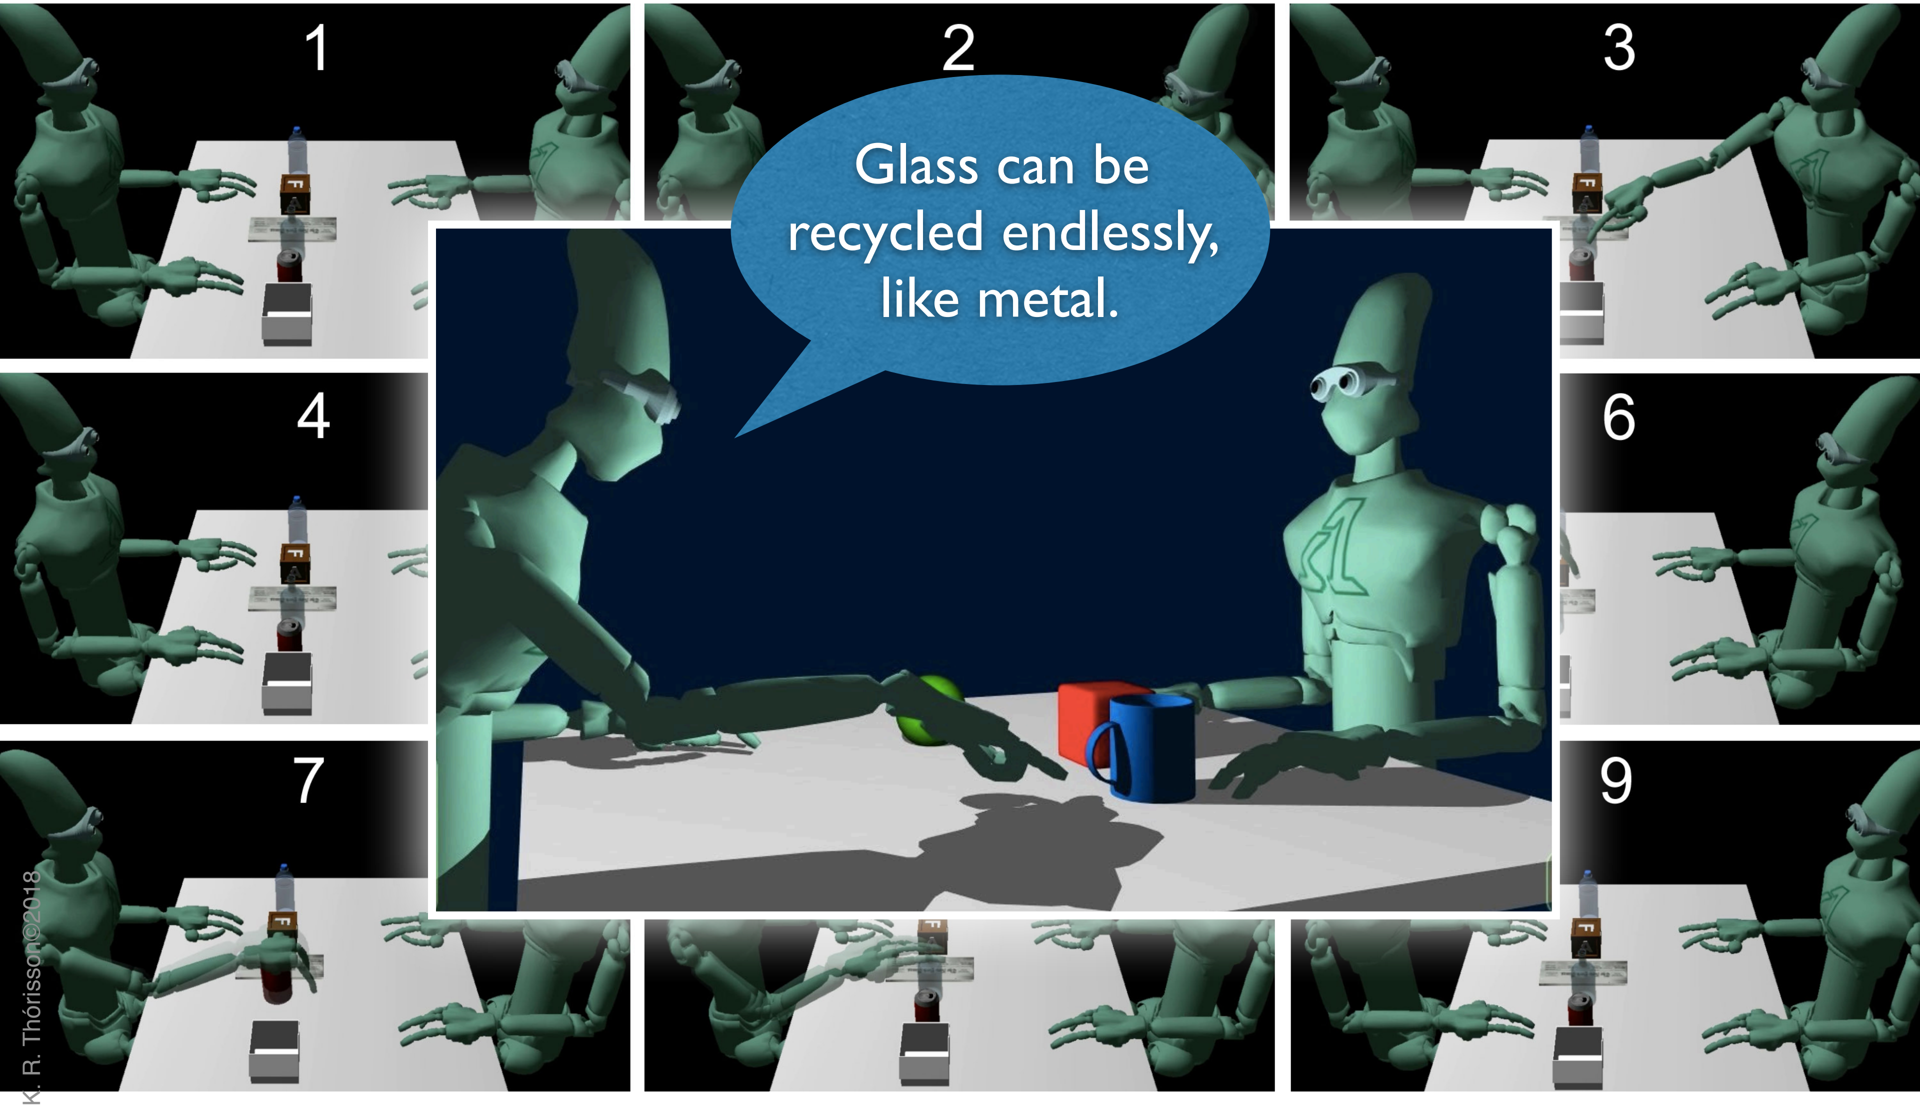

Causal-Relational Models

- The process of creating and using models involves testing them to evaluate their *usefulness:*
- the more accurately they help the controller achieve goals the more useful they are
- This is done continuously

Generating CRMs

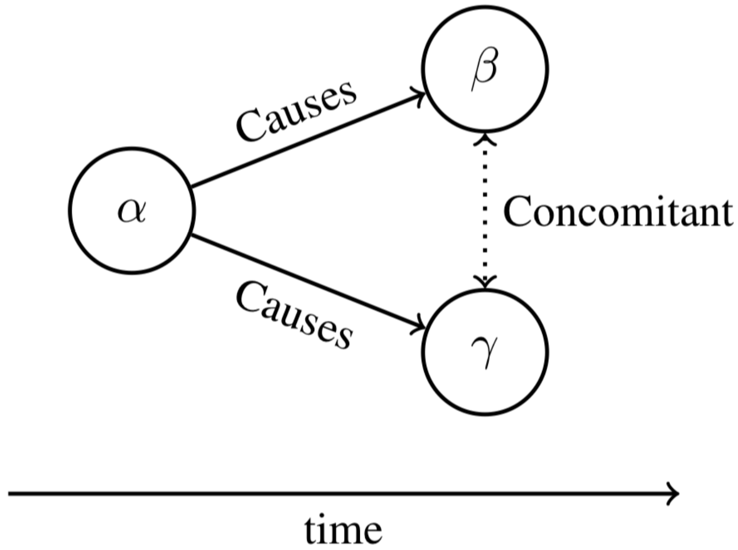

α, β, γ : observable variables

$$\alpha \Rightarrow \beta$$
$$\alpha \Rightarrow \gamma$$

Actual phenomenon in the physical world

Generating CRMs

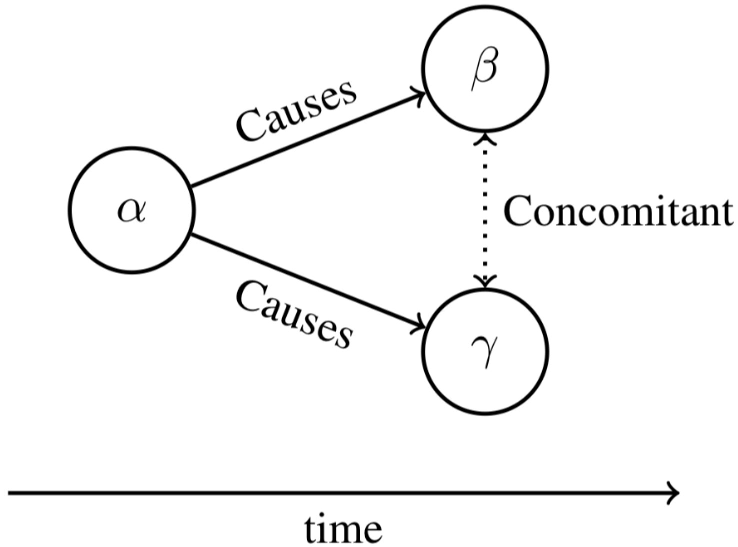

MODELS

M_1: $\alpha \Rightarrow \beta$

M_2: $\alpha \Rightarrow \gamma$

M_3: $\gamma \Rightarrow \beta$

M_4: $\beta \Rightarrow \gamma$

α, β, γ : observable variables

$$\alpha \Rightarrow \beta$$
$$\alpha \Rightarrow \gamma$$

Actual phenomenon in the physical world

Generating CRMs

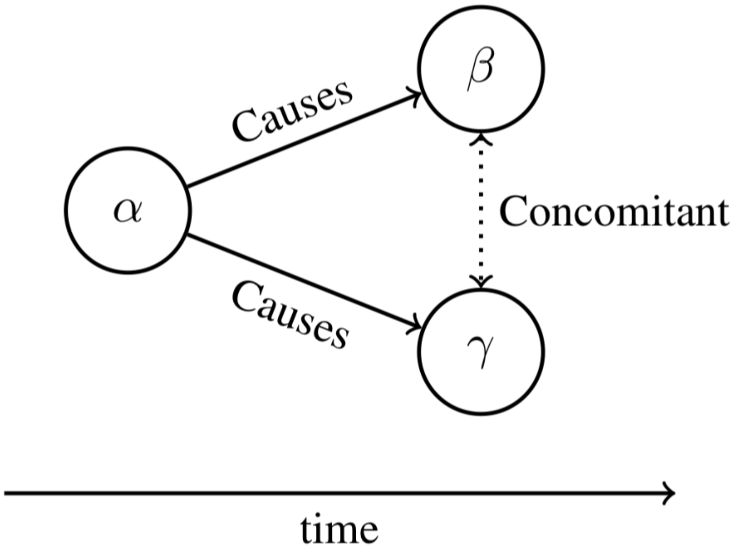

MODELS

M_1: $\alpha \Rightarrow \beta$

M_2: $\alpha \Rightarrow \gamma$

M_3: $\gamma \Rightarrow \beta$

M_4: $\beta \Rightarrow \gamma$

- Any of these will <u>predict</u> observed events correctly:
  - If you see β you will see γ, and vice versa
  - If you see α you will see β and γ

Generating CRMs

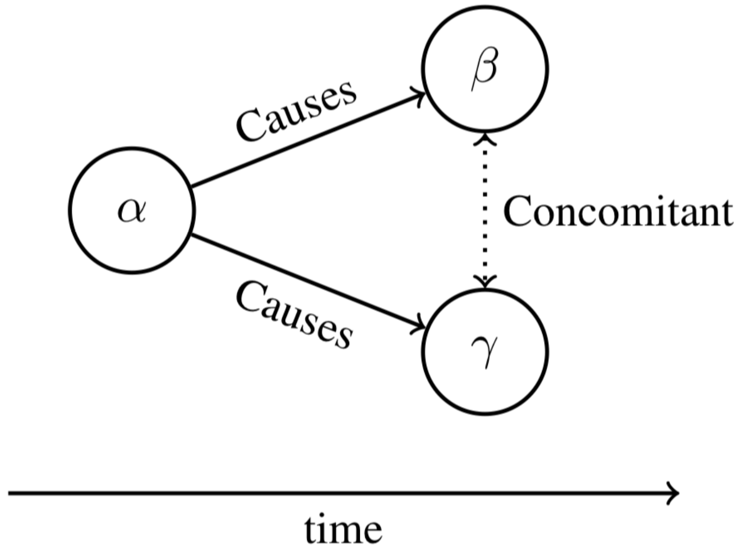

MODELS

M_1: $\alpha \Rightarrow \beta$

M_2: $\alpha \Rightarrow \gamma$

M_3: $\gamma \Rightarrow \beta$

M_4: $\beta \Rightarrow \gamma$

- However, if you want to stop seeing γ it does not help to make β go away, or vice versa
- To achieve a sub-goal of removing β (or γ) the only variable that can achieve that is α, captured only by M_1 (or M_2)

Probability: Doing ≠ Seeing

"... causality deals with how probability functions change in response to influences (e.g., new conditions or interventions) that originate from outside the probability space,
while
probability theory … cannot tell us how that function would change under such external influences.

Thus, 'doing' is not reducible to 'seeing', and there is no point trying to fuse the two together."

J. Pearl (2001).
Bayesianism and causality, or, why I am only a half-Bayesian.
In D. Corfield & J. Williamson (Eds.), Foundations of Bayesianism, vol. 12, 19–36.
p. 36

Causality ≠ Probability

"... causality deals with how probability functions change in response to influences (e.g., new conditions or interventions) that originate from outside the probability space,
while
probability theory … cannot tell us how that function would change under such external influences.

Thus, 'doing' is not reducible to 'seeing', and there is no point trying to fuse the two together."

J. Pearl (2001).
Bayesianism and causality, or, why I am only a half-Bayesian.
In D. Corfield & J. Williamson (Eds.), Foundations of Bayesianism, vol. 12, 19–36.
p. 36

CRMs are Bi-Directional

- CRMs are bi-directional in that
  - read “forward” (left to right) they mean that A *may cause* B
  - read “backward” they mean that B *may have been caused by* A

Bi-Directional CRMs

- CRMs are bi-directional in that
  - read “forward” (left to right) they mean that A *may cause* B
  - read “backward” they mean that B *may have been caused by* A

Bi-Directional CRMs

- CRMs are bi-directional in that
  - read “forward” (left to right) they mean that A *may cause* B
  - read “backward” they mean that B *may have been caused by* A

DEDUCTION

Bi-Directional CRMs

- CRMs are bi-directional in that
  - read "forward" (left to right) they mean that A *may cause* B
  - read "backward" they mean that B *may have been caused by* A

DEDUCTION

ABDUCTION

Bi-Directional CRMs

- CRMs are bi-directional in that
  - read "forward" (left to right) they mean that A *may cause* B
  - read "backward" they mean that B *may have been caused by* A

DEDUCTION

PREDICTION

ABDUCTION

Bi-Directional CRMs

- CRMs are <u>bi-directional</u> in that
  - read “forward” (left to right) they mean that A *may cause* B
  - read “backward” they mean that B *may have been caused by* A

DEDUCTION
PREDICTION
ABDUCTION
PLANNING

Model Generation

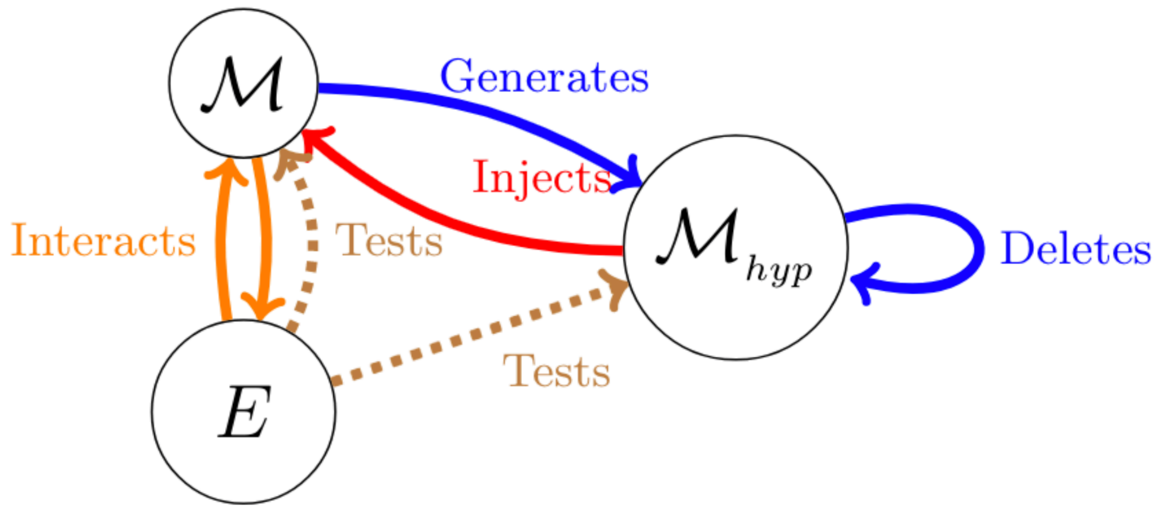

Continuous:

- model creation
- model testing
- comparison of new models' to old models' performance
- replacement of worse models with better models

Model Generation

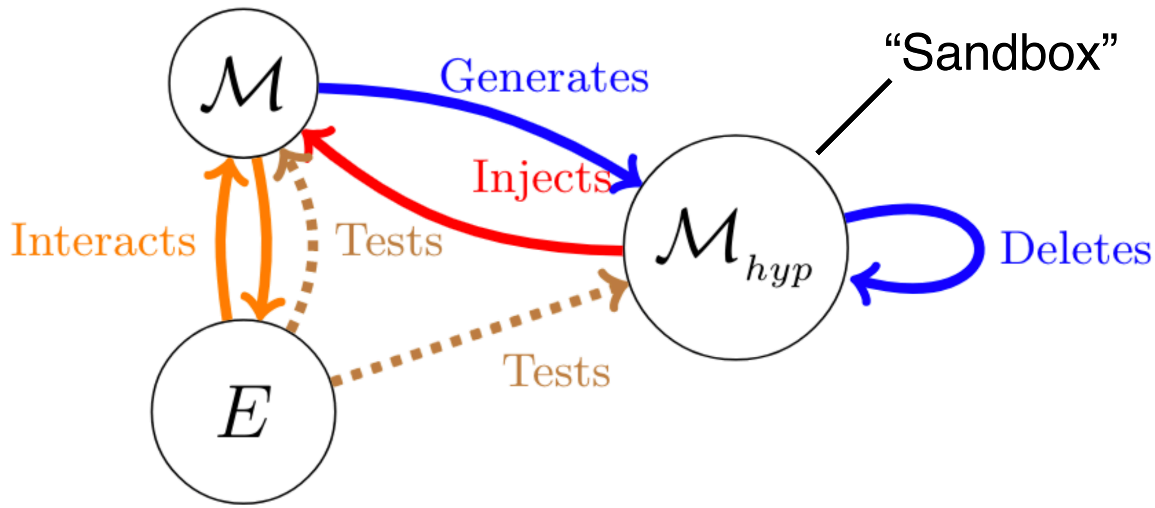

Continuous:

- model creation
- model testing
- comparison of new models' to old models' performance
- replacement of worse models with better models

Model Generation

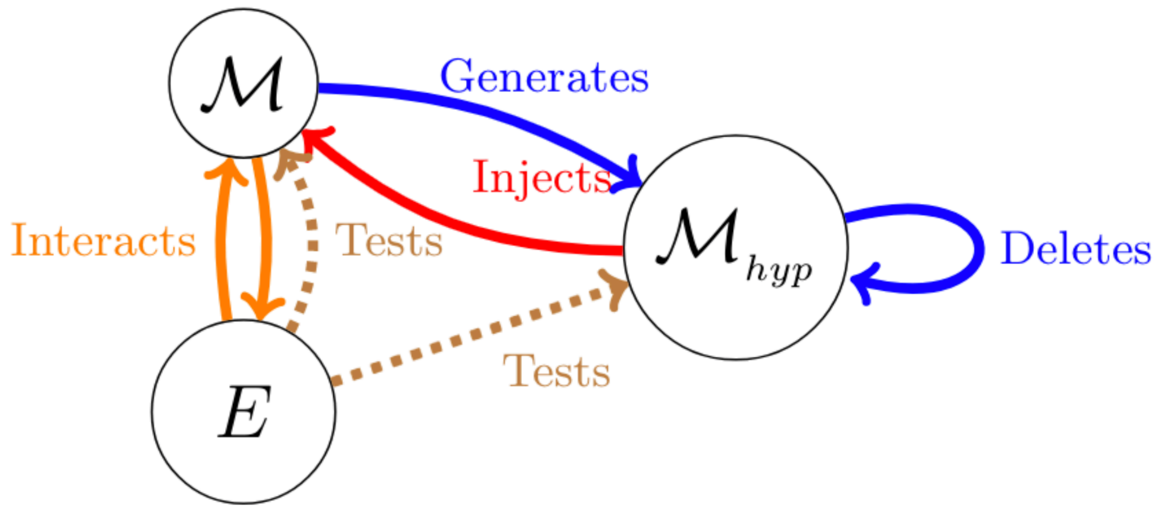

Continuous:

- model creation
- model testing
- comparison of new models' to old models' performance
- replacement of worse models with better models
- Each improved model introduces a small but noticeable improvement on overall system performance

A CUMULATIVE MODELER:

- continuously creates, evaluates, and uses models,
- racking up a model set, *M*, that
- becomes increasingly accurate over time, and thus increasingly useful.

A CUMULATIVE MODELER:

- continuously creates, evaluates, and uses models,
- racking up a model set, *M*, that
- becomes increasingly accurate over time, and thus increasingly useful.
- Upon repeated usage, every modification of the model set *M*, using the methods above, makes it more reliable for predicting and achieving goals within *E*.

A CUMULATIVE MODELER:

- continuously creates, evaluates, and uses models,
- racking up a model set, *M*, that
- becomes increasingly accurate over time, and thus increasingly useful.
- Upon repeated usage, every modification of the model set *M*, using the methods above, makes it more reliable for predicting **and** achieving goals within *E*.

A CUMULATIVE MODELER:

- continuously creates, evaluates, and uses models,
- racking up a model set, *M*, that
- becomes increasingly accurate over time, and thus increasingly useful.
- Upon repeated usage, every modification of the model set *M*, using the methods above, makes it more reliable for <u>predicting</u> **and** <u>achieving</u> goals within *E*.
- Repeated usage and testing of the models increases the overall reliability of the set as a whole, in small steps, cumulatively, as they capture the target phenomena in increasing detail and accuracy.

A CUMULATIVE MODELER:

- continuously creates, evaluates, and uses models
- racking up a model set, *M*, that
- becomes increasingly accurate over time, and thus increasingly useful.
- Upon repeated usage, every modification of the model set *M*, using the methods above, makes it more reliable for predicting **and** achieving goals within *E*.
- Repeated usage and testing of the models increases the overall reliability of the set as a whole, in small steps, cumulatively, as they capture the target phenomena in increasing detail and accuracy.

Conclusions

- Unified abduction and deduction in causal-relational models:
  - enables incremental improvement of the model set,
  - to increasingly contain models that approach actual causal relations between variables,
  - incrementally improving the performance of the overall system.

Conclusions

- Unified abduction and deduction in causal-relational models:
  - enables incremental improvement of the model set,
  - to increasingly contain models that approach actual causal relations between variables,
  - incrementally improving the performance of the overall system.
- Even in partially non-deterministic worlds, having models that approximate causal relations between variables is better than having only statistical information because it explicitly identifies the relevant variables affected by any action (to the extent possible), and thus provides direct support for goal-achievement.

Conclusions

- The result is **domain-independent cumulative learning**, where
  - knowledge increasingly represents causal and other important relations between observable variables in an environment
  - new knowledge is integrated with old knowledge in a logical process based on deduction and abduction - good models support both correctly,
    - serving foresight (for producing expectations) and
    - causal modeling (for producing plans and sub-goals) in the service of goal achievement

HUMANOBS project   AERA system

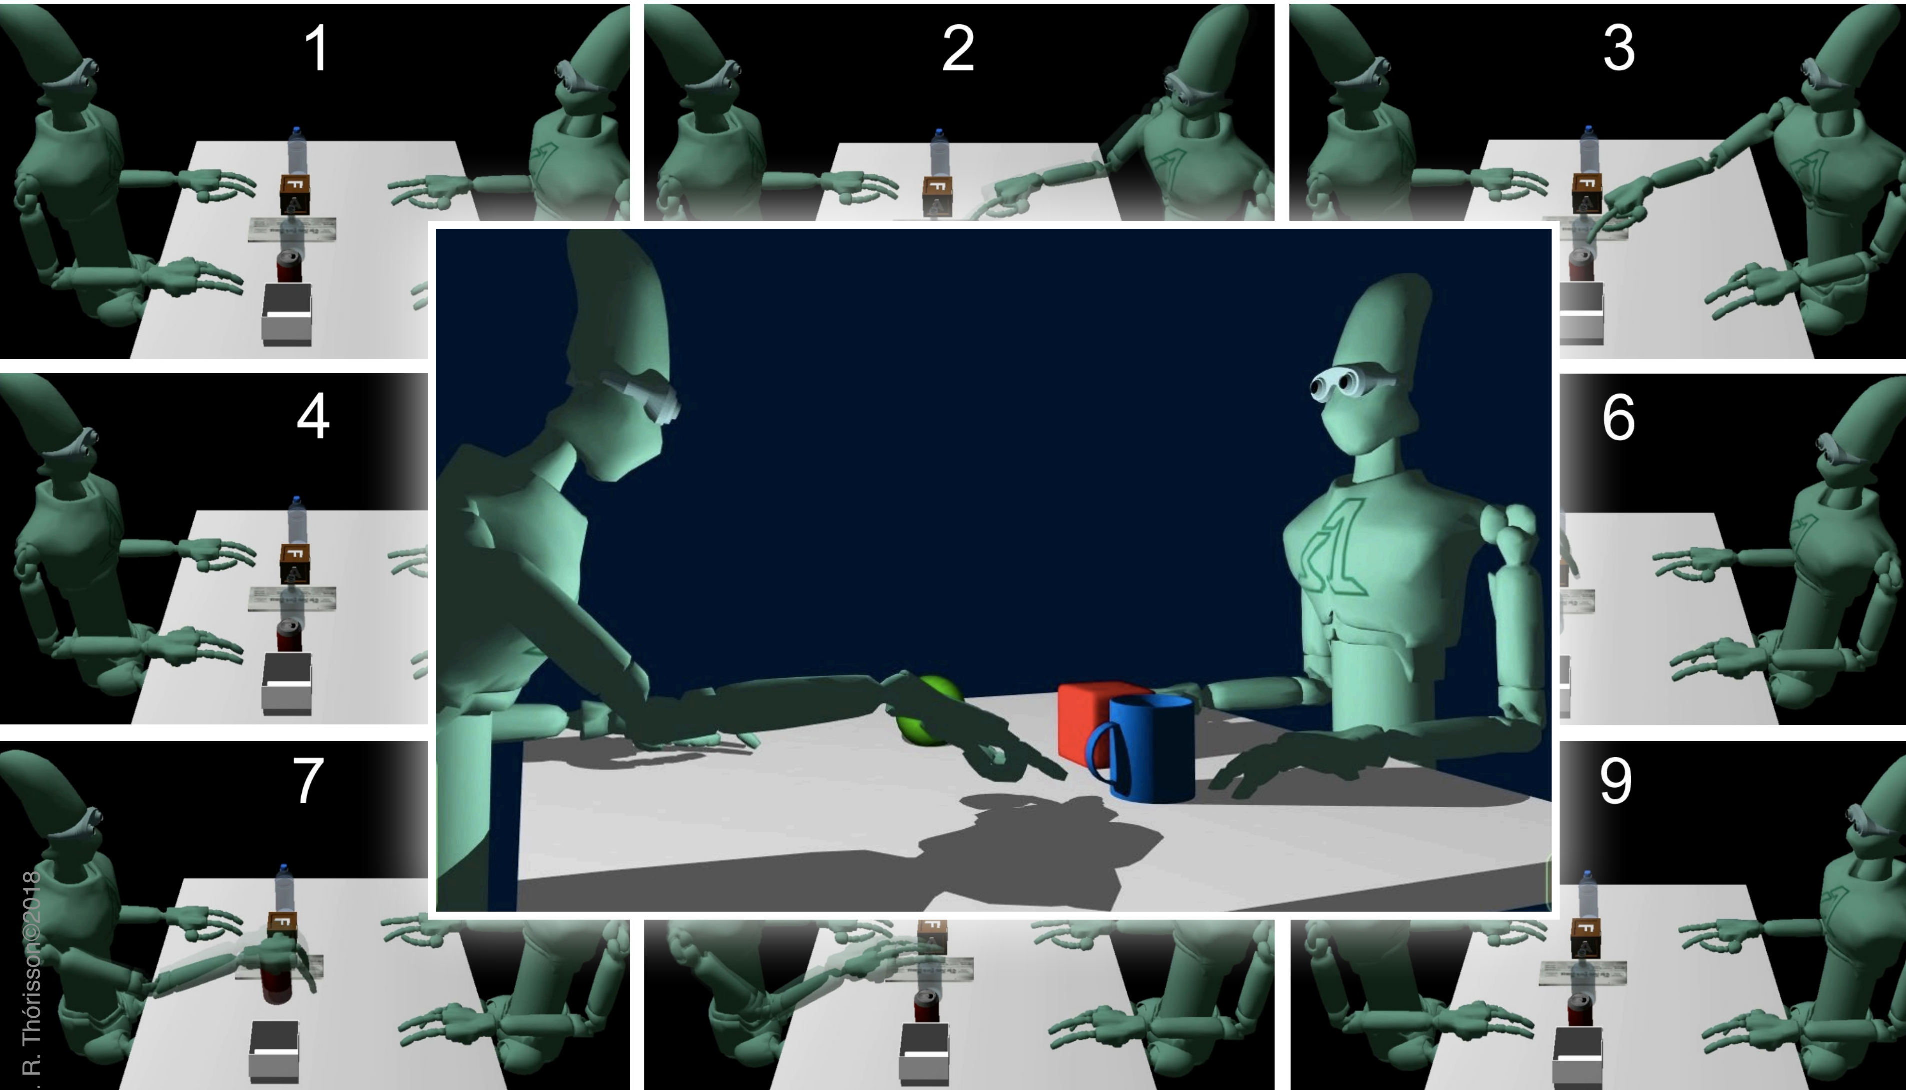

THANKS