

Machine-Learning Machines Don't Learn; AI Needs Real Learning, eg Learning *Ex Nihilo*

Selmer Bringsjord

Rensselaer AI & Reasoning (RAIR) Lab
Department of Cognitive Science
Department of Computer Science
Lally School of Management & Technology
Rensselaer Polytechnic Institute (RPI)
Troy, New York 12180 USA

IFLAI2
Oct 7 2021



Machine-Learning Machines Don't Learn; AI Needs Real Learning, eg Learning *Ex Nihilo*

Selmer Bringsjord

Rensselaer AI & Reasoning (RAIR) Lab
Department of Cognitive Science
Department of Computer Science
Lally School of Management & Technology
Rensselaer Polytechnic Institute (RPI)
Troy, New York 12180 USA

IFLAI2
Oct 7 2021



Machine-Learning Machines Don't Learn; AI Needs Real Learning, eg Learning *Ex Nihilo*


Selmer Bringsjord

Rensselaer AI & Reasoning (RAIR) Lab
Department of Cognitive Science
Department of Computer Science
Lally School of Management & Technology
Rensselaer Polytechnic Institute (RPI)
Troy, New York 12180 USA


IFLA12
Oct 7 2021



Advanced Logician (Real) Machine Learning



Advanced Logician (Real) Machine Learning



Given This, Do Machine-Learning Machines Learn? No.

Given This, Do Machine-Learning Machines Learn? No.

Do Machine-Learning Machines Learn?

Selmer Bringsjord and Naveen Sundar Govindarajulu and Shreya Banerjee and John Hummel

Abstract We answer the present paper's title in the negative. We begin by introducing and characterizing “real learning” ($\mathcal{RL}$) in the formal sciences, a phenomenon that has been firmly in place in homes and schools since at least Euclid. The defense of our negative answer pivots on an integration of *reductio* and proof by cases, and constitutes a general method for showing that any contemporary form of machine learning (ML) isn't real learning. Along the way, we canvass the many different conceptions of “learning” in not only AI, but psychology and its allied disciplines; none of these conceptions (with one exception arising from the view of cognitive development espoused by Piaget), aligns with real learning. We explain in this context by four steps how to broadly characterize and arrive at a focus on $\mathcal{RL}$.

Selmer Bringsjord
Rensselaer Polytechnic Institute, 110 8th Street Troy, NY USA 12180, e-mail: selmerbringsjord@gmail.com

Naveen Sundar Govindarajulu
Rensselaer Polytechnic Institute, 110 8th Street Troy, NY USA 12180, e-mail: Naveen.Sundar.G@gmail.com

Shreya Banerjee
Rensselaer Polytechnic Institute, 110 8th Street Troy, NY USA 12180, e-mail: shreyabbanerjee@gmail.com

John Hummel
901 West Illinois Street, Urbana, IL 61801, e-mail: jehummel@illinois.edu

Do Machine-Learning Machines Learn?

17

8 Appendix: The Formal Method

The following deduction uses fonts in an obvious and standard way to sort between functions (f), agents (a), and computing machines (m) in the Arithmetical Hierarchy. Ordinary italicized Roman is used for particulars under these sorts (e.g. f is a particular function). In addition, ‘ C ’ denotes any collection of conditions constituting jointly necessary-and-sufficient conditions for a form of current ML, which can come from relevant textbooks (e.g. Luger, 2008; Russell and Norvig, 2009) or papers; we leave this quite up to the reader, as no effect upon the validity of the deductive inference chain will be produced by the preferred instantiation of ‘ C .’ It will perhaps be helpful to the reader to point out that the deduction eventuates in the proposition that no machine in the ML fold that in this style learns a relevant function f thereby also real-learns f . We encode this target as follows:

$$(\star) \neg \exists m \exists f [\phi := MLearns(m, f) \wedge \psi := RLearns(m, f) \wedge C_\phi(m, f) \vdash^* (ci') \neg (cii)_{\psi}(m, f)]$$

Note that $(\star)$ employs meta-logical machinery to refer to particular instantiations of C for a particular, arbitrary case of ML (ϕ is the atomic sub-formula that can be instantiated to make the particular case), and particular instantiations of the triad $(ci') \neg (cii)$ for a particular, arbitrary case of $\mathcal{RL}$ (ψ is the atomic sub-formula that can be instantiated to make the particular case). Meta-logical machinery also allows us to use a provability predicate to formalize the notion that real learning is produced by the relevant instance of ML. If we “pop” ϕ/ψ to yield ϕ'/ψ' we are dealing with the particular instantiation of the atomic sub-formula.

The deduction, as noted in earlier when the informal argument was given, is indirect proof by cases; accordingly, we first assume $\neg(\star)$, and then proceed as follows under this supposition.

(1)	$\forall f, a [f : \mathbb{N} \mapsto \mathbb{N} \rightarrow (RLearns(a, f) \rightarrow (i) \neg (iii))]$	Def of Real Learning
(2)	$MLearns(m, f) \wedge RLearns(m, f) \wedge f : \mathbb{N} \mapsto \mathbb{N}$	supp (for $\exists$ elim on $(\star)$)
(3)	$\forall m, f [f : \mathbb{N} \mapsto \mathbb{N} \rightarrow (MLearns(m, f) \leftrightarrow C(m, f))]$	Def of ML
(4)	$\forall f [f : \mathbb{N} \mapsto \mathbb{N} \rightarrow (TurComp(f) \vee TurUncomp(f))]$	theorem
(5)	$TurUncomp(f)$	supp; Case 1
(6)	$\neg \exists m \exists f [(f : \mathbb{N} \mapsto \mathbb{N} \wedge TurUncomp(f) \wedge C(m, f))]$	theorem
$\therefore$ (7)	$\neg \exists m MLearns(m, f)$	(6), (3)
$\therefore$ (8)	$\perp$	(7), (2)
(9)	$TurComp(f)$	supp; Case 2
$\therefore$ (10)	$C_{\phi'}(m, f)$	(2), (3)
$\therefore$ (11)	$(ci') \neg (cii)_{\psi'}(m, f)$	from supp for $\exists$ elim on $(\star)$ and provability
$\therefore$ (12)	$\neg (ci') \neg (cii)_{\psi'}(m, f)$	inspection: proofs wholly absent from C
$\therefore$ (13)	$\perp$	(11), (12)
$\therefore$ (14)	$\perp$	<i>reductio</i> ; proof by cases

The Four-Step Road to Real Learning

The Four-Step Road to Real Learning

- Step 1: Observe the acute discontinuity of human vs. nonhuman cognition. (Only humans understand and employ e.g. abstract reasoning schemas unaffected by the physical; layered quantification; recursion; and infinite structures/infinitary reasoning.)

The Four-Step Road to Real Learning

- Step 1: Observe the acute discontinuity of human vs. nonhuman cognition. (Only humans understand and employ e.g. abstract reasoning schemas unaffected by the physical; layered quantification; recursion; and infinite structures/infinity reasoning.)
- Step 2: Exclude forms of “learning” made possible via exclusive use of reasoning and communication capacities in nonhuman animals (i.e. exclude forms of “learning” that don’t eventuate in bona fide jtb *knowledge*).

The Four-Step Road to Real Learning

- Step 1: Observe the acute discontinuity of human vs. nonhuman cognition. (Only humans understand and employ e.g. abstract reasoning schemas unaffected by the physical; layered quantification; recursion; and infinite structures/infinitary reasoning.)
- Step 2: Exclude forms of “learning” made possible via exclusive use of reasoning and communication capacities in nonhuman animals (i.e. exclude forms of “learning” that don’t eventuate in bona fide jtb *knowledge*).
- Step 3: Within the focus arising from Step 2, further narrow the focus to $HL^{\geq}$ reasoning and communication sufficiently powerful to perceive, and be successfully applied to, both (i) cohesive bodies of declarative content, and (ii) sophisticated natural-language content. Dub this **RC**.

The Four-Step Road to Real Learning

- Step 1: Observe the acute discontinuity of human vs. nonhuman cognition. (Only humans understand and employ e.g. abstract reasoning schemas unaffected by the physical; layered quantification; recursion; and infinite structures/infinity reasoning.)
- Step 2: Exclude forms of “learning” made possible via exclusive use of reasoning and communication capacities in nonhuman animals (i.e. exclude forms of “learning” that don’t eventuate in bona fide jtb *knowledge*).
- Step 3: Within the focus arising from Step 2, further narrow the focus to $HL^{\geq}$ reasoning and communication sufficiently powerful to perceive, and be successfully applied to, both (i) cohesive bodies of declarative content, and (ii) sophisticated natural-language content. Dub this **RC**.
- Step 4: Real Learning (RL) is the acquisition of genuine knowledge via **RC**.

But how is this mechanizable?

Well, how about a new form of machine learning?
(by reasoning)

Novel Form of Machine Learning:

Learning *Ex Nihilo*

(or Learning *Ex Minima*)

Novel Form of Machine Learning:

Learning *Ex Nihilo*

(or Learning *Ex Minima*)



EPiC Series in Computing

Volume 72, 2020, Pages 1–27

GCAI 2020. 6th Global Conference
on Artificial Intelligence (GCAI 2020)



Learning *Ex Nihilo*

Selmer Bringsjord¹, Naveen Sundar Govindarajulu¹, John Licato², and Michael Giancola¹

¹ Rensselaer AI & Reasoning (RAIR) Lab,
Rensselaer Polytechnic Institute, Troy NY 12180 USA

² Advancing Machine and Human Reasoning (AMHR) Lab,
University of South Florida, Tampa FL 33620 USA

{ selmer.bringsjord, naveen.sundar.g, mike.j.giancola, john.licato } @gmail.com

Abstract

This paper introduces, philosophically and to a degree formally, the novel concept of learning *ex nihilo*, intended (obviously) to be analogous to the concept of creation *ex nihilo*. Learning *ex nihilo* is an agent's learning “from nothing”, by the suitable employment of inference schemata for deductive and inductive reasoning. This reasoning must be in machine-verifiable accord with a formal proof/argument theory in a *cognitive calculus* (i.e., here, roughly, an intensional higher-order multi-operator quantified logic), and this reasoning is applied to percepts received by the agent, in the context of both some prior knowledge, and some prior and current interests. Learning *ex nihilo* is a challenge to contemporary forms of ML, indeed a severe one, but the challenge is here offered in the spirit of seeking to stimulate attempts, on the part of non-logicist ML researchers and engineers, to collaborate with those in possession of learning-*ex nihilo* frameworks, and eventually attempts to integrate directly with such frameworks at the implementation level. Such integration will require, among other things, the symbiotic interoperation of state-of-the-art automated reasoners and high-expressivity planners, with statistical/connectionist ML technology.

1 Introduction

This paper introduces, philosophically and to a degree logico-mathematically, the novel concept of learning *ex nihilo*, intended (obviously) to be analogous to the concept of creation *ex nihilo*.¹ Learning *ex nihilo* is an agent's learning “from nothing,” by the suitable employment of inference schemata for deductive and inductive² (e.g., analogical, enumerative-inductive, abductive, etc.) reasoning. This reasoning must be in machine-verifiable accord with a formal

¹No such assumption as that creation *ex nihilo* is real or even formally respectable is made or needed in the present paper. The concept of creation *ex nihilo* is simply for us an intellectual inspiration — but as a matter of fact, the literature on it in analytic philosophy does provide some surprisingly rigorous accounts. In the present draft of the present paper, we don't seek to mine these accounts.

²Not to be confused with inductive logic programming (about which more will be said later), or inductive deductive techniques and schemas (e.g. mathematical induction, the induction schema in Peano Arithmetic, etc.). As we explain later, learning *ex nihilo* is in part powered by non-deductive inference schemata seen in inductive logic. An introductory overview of inductive logic is provided in [39].


Bringsjord, S., Govindarajulu, N.S., Licato, J. & Giancola, M. (2020)
“Learning *Ex Nihilo*” *Proceedings of the 6th Global Conference on Artificial Intelligence (GCAI 2020)*, within *International Conferences on Logic and Artificial Intelligence at Zhejiang University (ZJULogAI)*, in Danoy, G., Pang, J. & Sutcliffe, G., eds., *EPiC Series in Computing* **72**: 1–27 (Manchester, UK: EasyChair Ltd), ISSN: 2398-7340.
<https://easychair.org/publications/paper/NzWVG>

Bringsjord, S., Govindarajulu, N.S., Licato, J. & Giancola, M. (2020)
“Learning *Ex Nihilo*” *Proceedings of the 6th Global Conference on Artificial Intelligence (GCAI 2020)*, within *International Conferences on Logic and Artificial Intelligence at Zhejiang University (ZJULogAI)*, in Danoy, G., Pang, J. & Sutcliffe, G., eds., *EPiC Series in Computing* **72**: 1–27 (Manchester, UK: EasyChair Ltd), ISSN: 2398-7340.

<https://easychair.org/publications/paper/NzWVG>

Bringsjord, S., Govindarajulu, N.S., Licato, J. & Giancola, M. (2020)
“Learning *Ex Nihilo*” *Proceedings of the 6th Global Conference on Artificial Intelligence (GCAI 2020)*, within *International Conferences on Logic and Artificial Intelligence at Zhejiang University (ZJULogAI)*, in Danoy, G., Pang, J. & Sutcliffe, G., eds., *EPiC Series in Computing* **72**: 1–27 (Manchester, UK: EasyChair Ltd), ISSN: 2398-7340.

<https://easychair.org/publications/paper/NzWVG>



I 3-Strength-Value Continuum

Certain

Evident

Overwhelmingly Likely

Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not

Unlikely

Overwhelmingly Unlikely

Beyond Reasonable Belief

Evidently False

Certainly False

I 3-Strength-Value Continuum

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

I 3-Strength-Value Continuum

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

..... Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

I 3-Strength-Value Continuum

Epistemically Positive

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

I 3-Strength-Value Continuum

Epistemically Positive

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

Epistemically Negative

I 3-Strength-Value Continuum

Epistemically Positive

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

Epistemically Negative

I 3-Strength-Value Continuum

Epistemically Positive

- (6) Certain
- (5) Evident
- (4) Overwhelmingly Likely
- (3) Beyond Reasonable Doubt
- (2) Likely
- (1) More Likely Than Not
- (0) Counterbalanced
- (-1) More Unlikely Than Not
- (-2) Unlikely
- (-3) Overwhelmingly Unlikely
- (-4) Beyond Reasonable Belief
- (-5) Evidently False
- (-6) Certainly False

Epistemically Negative

Or ... 11 Strength Factors

Acceptable An agent a at time t finds ϕ acceptable *iff* withholding ϕ is not more reasonable than believing in ϕ .

$$\mathbf{B}^1(a, t, \phi) \Leftrightarrow \begin{cases} \mathbf{W}(a, t, \phi) \not\succ_t^a \mathbf{B}(a, t, \phi); \text{ or} \\ \left(\neg \mathbf{B}(a, t, \phi) \wedge \neg \mathbf{B}(a, t, \neg \phi) \right) \not\succ_t^a \mathbf{B}(a, t, \phi) \end{cases}$$

Some Presumption in Favor An agent a at time t has some presumption in favor of ϕ *iff* believing ϕ at t is more reasonable than believing $\neg \phi$ at time t :

$$\mathbf{B}^2(a, t, \phi) \Leftrightarrow \left(\mathbf{B}(a, t, \phi) \succ_a^t \mathbf{B}(a, t, \neg \phi) \right)$$

Beyond Reasonable Doubt An agent a at time t has beyond reasonable doubt in ϕ *iff* believing ϕ at t is more reasonable than withholding ϕ at time t :

$$\mathbf{B}^3(a, t, \phi) \Leftrightarrow \begin{cases} \mathbf{B}(a, t, \phi) \succ_a^t \mathbf{W}(a, t, \phi); \text{ or} \\ \left(\mathbf{B}(a, t, \phi) \succ_t^a \left(\neg \mathbf{B}(a, t, \phi) \wedge \neg \mathbf{B}(a, t, \neg \phi) \right) \right) \end{cases}$$


Evident A formula ϕ is evident to an agent a at time t *iff* ϕ is beyond reasonable doubt and if there is a ψ such that believing ψ is more reasonable for a at time t than believing ϕ , then a is certain about ψ at time t .

$$\mathbf{B}^4(a, t, \phi) \Leftrightarrow \begin{cases} \mathbf{B}^3(a, t, \phi) \wedge \\ \exists \psi : \left[\begin{array}{l} \mathbf{B}(a, t, \psi) \succ_t^a \mathbf{B}(a, t, \phi) \\ \Rightarrow \mathbf{B}^5(a, t, \psi) \end{array} \right] \end{cases}$$


Certain An agent a at time t is certain about ϕ *iff* ϕ is beyond reasonable doubt and there is no ψ such that believing ψ is more reasonable for a at time t than believing ϕ .

$$\mathbf{B}^5(a, t, \phi) \Leftrightarrow \begin{cases} \mathbf{B}^3(a, t, \phi) \wedge \\ \neg \exists \psi : \mathbf{B}(a, t, \psi) \succ_t^a \mathbf{B}(a, t, \phi) \end{cases}$$

Demo: “Sully-esque AI Performs Miracle on the Hudson”



Demo: “Sully-esque AI Performs Miracle on the Hudson”






Michael Giancola
Graduate Research Assistant (PhD)

TEB






LGA




TEB

LGA







"Birds"





TEB






ATC





LGA





$S(atc, capt, t_1, \text{Land}(capt, t_1, lga_{13}))$
 $\therefore B(capt, t_1, B(atc, t_1, \text{Land}(capt, t_1, lga_{13}))) \quad [I_{12}] \quad \checkmark$
 $\therefore B^1(capt, t_1, \text{Land}(capt, t_1, lga_{13})) \quad [B^1\text{-def}] \quad \checkmark$

















$\therefore \text{Land}(\text{capt}, t_2, \text{teb}) \succ_{t_2}^{\text{capt}} \text{Land}(\text{capt}, t_2, \text{lga}_{13}) \text{ [} \succ_t^a \text{ -def] } \checkmark$
 $\therefore \mathbf{B}^2(\text{capt}, t_2, \text{Land}(\text{capt}, t_2, \text{teb})) \text{ [} \mathbf{B}^2\text{-def] } \checkmark$









$$S(atc, capt, t_3, \text{Land}(capt, t_3, teb_1))$$

$$\therefore B(capt, t_3, B(atc, t_3, \text{Land}(capt, t_3, teb_1))) \quad [I_{12}] \quad \checkmark$$

$$\therefore B^1(capt, t_3, \text{Land}(capt, t_3, teb_1)) \quad [B^1\text{-def}] \quad \checkmark$$

$$\therefore \text{Land}(capt, t_3, hud) \succ_{t_3}^{capt} \text{Land}(capt, t_3, teb_1) \quad [\succ_t^a\text{-def}] \quad \checkmark$$


$$\therefore B^2(capt, t_3, \text{Land}(capt, t_3, hud)) \quad [B^2\text{-def}] \quad \checkmark$$





“Okay, TEB”

ATC



$$\mathbf{S}(atc, capt, t_3, \text{Land}(capt, t_3, teb_1))$$
$$\therefore \mathbf{B}(capt, t_3, \mathbf{B}(atc, t_3, \text{Land}(capt, t_3, teb_1))) \quad [I_{12}] \quad \checkmark$$
$$\therefore \mathbf{B}^1(capt, t_3, \text{Land}(capt, t_3, teb_1)) \quad [\mathbf{B}^1\text{-def}] \quad \checkmark$$


$$\therefore \text{Land}(capt, t_3, hud) \succ_{t_3}^{capt} \text{Land}(capt, t_3, teb_1) \quad [\succ_t^a\text{-def}] \quad \checkmark$$
$$\therefore \mathbf{B}^2(capt, t_3, \text{Land}(capt, t_3, hud)) \quad [\mathbf{B}^2\text{-def}] \quad \checkmark$$






“Okay, TEB”

ATC




ShadowAdjudicator




ShadowAdjudicator

- A nascent automated reasoner for generating and adjudicating arguments




ShadowAdjudicator

- A nascent automated reasoner for generating and adjudicating arguments
- Builds upon ShadowProver




ShadowAdjudicator

- A nascent automated reasoner for generating and adjudicating arguments
- Builds upon ShadowProver
 - Uses ShadowProver for sub-proofs of modal/FOL/PL formulae




ShadowAdjudicator


- A nascent automated reasoner for generating and adjudicating arguments
- Builds upon ShadowProver
 - Uses ShadowProver for sub-proofs of modal/FOL/PL formulae
 - Implements an algorithm and inference schemata for generating arguments with strength factors



Simulation



Simulation



Conference Paper Under Review

- 5th International Conference on Robot Ethics and Standards

ICRES 2020: 5th International Conference on Robot Ethics and Standards, Taipei, Taiwan, 28-29 September 2020.

ETHICAL REASONING FOR AUTONOMOUS AGENTS UNDER UNCERTAINTY

MICHAEL GIANCOLA* and SELMER BRINGSJORD[†] and NAVEEN SUNDAR GOVINDARAJULU* and CARLOS VARELA[‡]

Rensselaer AI & Reasoning (RAIR) Lab[†]; Worldwide Computing Lab (WCL)[‡]
Department of Computer Science*^{†‡}; Department of Cognitive Science[†]
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
www.rpi.edu*

E-mail: { mike.j.giancola, selmer.bringsjord, naveen.sundar.g } @gmail.com, cvarela@cs.rpi.edu


Autonomous (and partially autonomous) agents are beginning to play significant roles in safety-critical and privacy-critical domains, such as driving and healthcare. When humans operate in these spaces, not only are there regulations and laws dictating proper behavior, but crucially, neurobiologically normal humans can be expected to comprehend how to reason with certain principles to ensure that their actions are legally/ethically/prudentially correct (whether or not these humans choose to abide by the principles in question). It seems reasonable that we should hold autonomous agents to, minimally, the same standard we hold humans to. In this paper, we present a framework for autonomous aircraft piloting agents to reason about ethical problems in the context of emergency landings. In particular, we are concerned with ethical problems in which every option is equally unethical with regard to the ethical principles the options violate; and the only distinguishing factor is the *likelihood* that a plan will violate an ethical principle. We conclude by discussing why, in general, we find an inference-theoretic approach to ethical reasoning to be superior to the model-theoretic approach of prior work.

Keywords: Ethical reasoning; Reasoning under uncertainty; Modal logic


1. Introduction

Autonomous (and partially autonomous) agents are beginning to play significant roles in safety-critical and privacy-critical domains, such as driving and healthcare. When humans operate in these spaces, not only are there regulations and laws dictating proper behavior, but crucially, neurobiologically normal humans can be expected to comprehend how to reason with regulations to ensure that their actions remain within the law (regardless of whether or not they choose to abide by the law). It seems reasonable that we should hold autonomous agents to, minimally, the same standard we hold humans to. That is, autonomous agents must be aware of relevant ethical constraints (those actions which are e.g. *obligatory*, *permissible*, *forbidden*, *supererogatory*, etc.; see [1] for a full discussion of these concepts in deontic logic), and be able to reason with them to determine actions which will satisfy those constraints. In addition, we require of our autonomous agents to not only verify their behavior to be ethical, but to output a proof which can be inspected by a human. By our lights, this requires an inference-theoretic approach to ethical reasoning. We will discuss in greater detail why we find an inference-theoretic approach to be superior to a model-theoretic approach in §4.

Formally Shown: “Uber Fatality Avoided”




Formally Shown: “Uber Fatality Avoided”






Michael Giancola
Graduate Research Assistant (PhD)







Michael Giancola
Graduate Research Assistant (PhD)




At the start of the event, the vehicle was in autonomous mode in the rightmost of four lanes traveling in the same direction, and the pedestrian was walking her bicycle across the street starting on the leftmost side of the roadway.




The vehicle's radar first detected the pedestrian **5.6 seconds** before the fatal collision. Less than half a second later, the lidar detected the pedestrian but classified her as “**Other**”.




For the next 2.5 seconds, the lidar re-classified her several times, alternating between “Vehicle” and “Other”. The vehicle’s automated-driving system (ADS) attempted to predict her direction of travel several times, but **discarded any previous information about her trajectory every time it reclassified her.**



With 2.6 seconds until collision, the lidar classified her as a bicycle but, as it was yet again changing her classification, discarded any past trajectory information, and hence determined that she was not moving. Up to this point, **the car had not taken any evasive or corrective action.**




With 1.5 seconds left, the lidar re-classified her yet again, this time as “Unknown”. **The system once again loses all of its tracking history.** However, since at this point the pedestrian had entered the vehicle’s lane, the ADS generated a plan to turn the car to the right to avoid her.




Three hundred milliseconds later, the lidar re-classified her as a bicycle, and determined that it would be impossible at this point to maneuver around her. With just 200 ms until collision, the ADS began braking the vehicle, pitifully too late to stop in time.

Now, how our AI
would've handled it...



Three hundred milliseconds later, the lidar re-classified her as a bicycle, and determined that it would be impossible at this point to maneuver around her. With just 200 ms until collision, the ADS began braking the vehicle, pitifully too late to stop in time.



Label the time point at which the radar first detected the pedestrian $\mathbf{t_0}$,
and the location of the pedestrian at that time ℓ_1 .

Argument 1

Argument 2




Label the time point at which the radar first detected the pedestrian $\mathbf{t}_0$,
and the location of the pedestrian at that time ℓ_1 .

Argument 1

$P(r_1, t_0, At(o^*, t_0, \ell_1))$

Argument 2




Label the time point at which the radar first detected the pedestrian t_0 ,
and the location of the pedestrian at that time ℓ_1 .

Argument 1

$$\mathbf{P}(\mathbf{r}_1, t_0, \text{At}(o^*, t_0, \ell_1))$$

Argument 2

$$\mathbf{P}(\mathbf{r}_2, t_0, \text{At}(o^*, t_0, \ell_1))$$



Label the time point at which the radar first detected the pedestrian $\mathbf{t}_0$,
and the location of the pedestrian at that time ℓ_1 .

Argument 1

$P(r_1, t_0, At(o^*, t_0, \ell_1))$

Argument 2

$P(r_2, t_0, At(o^*, t_0, \ell_1))$



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t}_1$,
and new location ℓ_2 .


Argument 1

$P(r_1, t_0, At(o^*, t_0, \ell_1))$

...

Argument 2

$P(r_2, t_0, At(o^*, t_0, \ell_1))$



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t}_1$,
and new location ℓ_2 .

Argument 1

$\mathbf{P}(\mathbf{r}_1, t_0, \text{At}(o^*, t_0, \ell_1))$

...

$\mathbf{B}(\mathbf{r}_1, t_1, \neg \text{Moving}(o^*))$

$\succ_t^a$

$\mathbf{B}(\mathbf{r}_1, t_1, \text{Moving}(o^*))$

Argument 2

$\mathbf{P}(\mathbf{r}_2, t_0, \text{At}(o^*, t_0, \ell_1))$



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t}_1$,
and new location ℓ_2 .

Argument 1

$\mathbf{P}(\mathbf{r}_1, t_0, \text{At}(o^*, t_0, \ell_1))$

...

$\mathbf{B}(\mathbf{r}_1, t_1, \neg \text{Moving}(o^*))$

$\succ_t^a$


$\mathbf{B}(\mathbf{r}_1, t_1, \text{Moving}(o^*))$

$\therefore \mathbf{B}^2(\mathbf{r}_1, t_1, \neg \text{Moving}(o^*))$

$\therefore \mathbf{B}^2(\mathbf{r}_1, t_1, \neg \text{NeedToBrake}(c))$

Argument 2

$\mathbf{P}(\mathbf{r}_2, t_0, \text{At}(o^*, t_0, \ell_1))$



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t}_1$,
and new location ℓ_2 .

Argument 1

$\mathbf{P}(\mathbf{r}_1, t_0, \text{At}(o^*, t_0, \ell_1))$

...

$\mathbf{B}(\mathbf{r}_1, t_1, \neg \text{Moving}(o^*))$

$\succ_t^a$

$\mathbf{B}(\mathbf{r}_1, t_1, \text{Moving}(o^*))$

$\therefore \mathbf{B}^2(\mathbf{r}_1, t_1, \neg \text{Moving}(o^*))$


$\therefore \mathbf{B}^2(\mathbf{r}_1, t_1, \neg \text{NeedToBrake}(c))$

Argument 2

$\mathbf{P}(\mathbf{r}_2, t_0, \text{At}(o^*, t_0, \ell_1))$

$\mathbf{P}(\mathbf{r}_2, t_1, \text{At}(o^*, t_0, \ell_2))$

$\mathbf{P}(\mathbf{r}_2, t_1, \ell_1 = \ell_2)$



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t}_1$,
and new location ℓ_2 .

Argument 1

$\mathbf{P}(\mathbf{r}_1, t_0, \text{At}(o^*, t_0, \ell_1))$

...

$\mathbf{B}(\mathbf{r}_1, t_1, \neg \text{Moving}(o^*))$

$\succ_t^a$

$\mathbf{B}(\mathbf{r}_1, t_1, \text{Moving}(o^*))$

$\therefore \mathbf{B}^2(\mathbf{r}_1, t_1, \neg \text{Moving}(o^*))$

$\therefore \mathbf{B}^2(\mathbf{r}_1, t_1, \neg \text{NeedToBrake}(c))$


Argument 2

$\mathbf{P}(\mathbf{r}_2, t_0, \text{At}(o^*, t_0, \ell_1))$

$\mathbf{P}(\mathbf{r}_2, t_1, \text{At}(o^*, t_0, \ell_2))$

$\mathbf{P}(\mathbf{r}_2, t_1, \ell_1 = \ell_2)$

...



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t}_1$,
and new location ℓ_2 .

Argument 1

$P(r_1, t_0, At(o^*, t_0, \ell_1))$

...

$B(r_1, t_1, \neg Moving(o^*))$

$\succ_t^a$

$B(r_1, t_1, Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg NeedToBrake(c))$

Argument 2

$P(r_2, t_0, At(o^*, t_0, \ell_1))$


$P(r_2, t_1, At(o^*, t_0, \ell_2))$

$P(r_2, t_1, \ell_1 = \ell_2)$

...

$\therefore B^5(r_2, t_1, Moving(o^*))$

$\therefore B^5(r_2, t_1, NeedToBrake(c))$



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t_1}$,
and new location ℓ_2 .

Argument 1

$P(r_1, t_0, At(o^*, t_0, \ell_1))$

...

$B(r_1, t_1, \neg Moving(o^*))$

$\succ_t^a$

$B(r_1, t_1, Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg NeedToBrake(c))$

Argument 2

$P(r_2, t_0, At(o^*, t_0, \ell_1))$

$P(r_2, t_1, At(o^*, t_0, \ell_2))$


$P(r_2, t_1, \ell_1 = \ell_2)$

...

$\therefore B^5(r_2, t_1, Moving(o^*))$

$\therefore B^5(r_2, t_1, NeedToBrake(c))$

Adjudicator



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t_1}$,
and new location ℓ_2 .

Argument 1

$P(r_1, t_0, At(o^*, t_0, \ell_1))$

...

$B(r_1, t_1, \neg Moving(o^*))$

$\succ_t^a$

$B(r_1, t_1, Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg NeedToBrake(c))$

Argument 2

$P(r_2, t_0, At(o^*, t_0, \ell_1))$

$P(r_2, t_1, At(o^*, t_0, \ell_2))$

$P(r_2, t_1, \ell_1 = \ell_2)$


...

$\therefore B^5(r_2, t_1, Moving(o^*))$

$\therefore B^5(r_2, t_1, NeedToBrake(c))$

Adjudicator

$B(a, t_1, NeedToBrake(c))$



Then, the next time it detects and re-classifies the pedestrian as $\mathbf{t_1}$,
and new location ℓ_2 .

Argument 1

$P(r_1, t_0, At(o^*, t_0, \ell_1))$

...

$B(r_1, t_1, \neg Moving(o^*))$

$\succ_t^a$

$B(r_1, t_1, Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg Moving(o^*))$

$\therefore B^2(r_1, t_1, \neg NeedToBrake(c))$

Argument 2

$P(r_2, t_0, At(o^*, t_0, \ell_1))$

$P(r_2, t_1, At(o^*, t_0, \ell_2))$

$P(r_2, t_1, \ell_1 = \ell_2)$


...

$\therefore B^5(r_2, t_1, Moving(o^*))$


$\therefore B^5(r_2, t_1, NeedToBrake(c))$

Adjudicator


$B(a, t_1, NeedToBrake(c))$




Pollock's (Polyanna) Loop




Pollock's (Polyanna) Loop




The Learning *Ex Nihilo* Loop




The Learning *Ex Nihilo* Loop




The Learning *Ex Nihilo* Loop




The Learning *Ex Nihilo* Loop



The Learning *Ex Nihilo* Loop



The Learning *Ex Nihilo* Loop



13-Valued Likelihood Continuum

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not

Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

13-Valued Likelihood Continuum

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

..... Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

13-Valued Likelihood Continuum

Epistemically Positive

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

13-Valued Likelihood Continuum

Epistemically Positive

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

Epistemically Negative

13-Valued Likelihood Continuum

Epistemically Positive

Certain

Evident

Overwhelmingly Likely
Beyond Reasonable Doubt

Likely

More Likely Than Not

Counterbalanced

More Unlikely Than Not
Unlikely

Overwhelmingly Unlikely
Beyond Reasonable Belief

Evidently False

Certainly False

Epistemically Negative


13-Valued Likelihood Continuum

Epistemically Positive

- (6) Certain
- (5) Evident
- (4) Overwhelmingly Likely
- (3) Beyond Reasonable Doubt
- (2) Likely
- (1) More Likely Than Not
- (0) Counterbalanced
- (-1) More Unlikely Than Not
- (-2) Unlikely
- (-3) Overwhelmingly Unlikely
- (-4) Beyond Reasonable Belief
- (-5) Evidently False
- (-6) Certainly False

Epistemically Negative

13-Valued Likelihood Continuum



Formalizing & implementing this ...

DCEC Signature

$$\begin{aligned}
 S &::= \text{Agent} \mid \text{ActionType} \mid \text{Action} \sqsubseteq \text{Event} \mid \text{Moment} \mid \text{Fluent} \\
 f &::= \begin{cases} \text{action} : \text{Agent} \times \text{ActionType} \rightarrow \text{Action} \\ \text{initially} : \text{Fluent} \rightarrow \text{Formula} \\ \text{holds} : \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{happens} : \text{Event} \times \text{Moment} \rightarrow \text{Formula} \\ \text{clipped} : \text{Moment} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{initiates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{terminates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{prior} : \text{Moment} \times \text{Moment} \rightarrow \text{Formula} \end{cases} \\
 t &::= x : S \mid c : S \mid f(t_1, \dots, t_n) \\
 \phi &::= \begin{cases} q : \text{Formula} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \forall x : \phi(x) \mid \exists x : \phi(x) \\ \mathbf{P}(a, t, \phi) \mid \mathbf{K}(a, t, \phi) \mid \mathbf{S}(a, b, t, \phi) \mid \mathbf{S}(a, t, \phi) \\ \mathbf{C}(t, \phi) \mid \mathbf{B}(a, t, \phi) \mid \mathbf{D}(a, t, \phi) \mid \mathbf{I}(a, t, \phi) \\ \mathbf{O}(a, t, \phi, (\neg)\text{happens}(\text{action}(a^*, \alpha), t')) \end{cases}
 \end{aligned}$$

Modal Operator Descriptors:
Perceives, **K**nows, **S**ays, **C**ommon-knowledge
Believes, **D**esires, **I**ntends, **O**ught-to

DCEC Inference Schemata


$$\begin{aligned}
 &\frac{\mathbf{K}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{K}(a, t_2, \phi)} [I_K] \quad \frac{\mathbf{B}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{B}(a, t_2, \phi)} [I_B] \\
 &\frac{}{\mathbf{C}(t, \mathbf{P}(a, t, \phi) \rightarrow \mathbf{K}(a, t, \phi))} [I_1] \quad \frac{}{\mathbf{C}(t, \mathbf{K}(a, t, \phi) \rightarrow \mathbf{B}(a, t, \phi))} [I_2] \\
 &\frac{\mathbf{C}(t, \phi) t \leq t_1 \dots t \leq t_n}{\mathbf{K}(a_1, t_1, \dots \mathbf{K}(a_n, t_n, \phi) \dots)} [I_3] \quad \frac{\mathbf{K}(a, t, \phi)}{\phi} [I_4] \\
 &\frac{}{\mathbf{C}(t, \mathbf{K}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{K}(a, t_2, \phi_1) \rightarrow \mathbf{K}(a, t_3, \phi_2)} [I_5] \\
 &\frac{}{\mathbf{C}(t, \mathbf{B}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a, t_2, \phi_1) \rightarrow \mathbf{B}(a, t_3, \phi_2)} [I_6] \\
 &\frac{}{\mathbf{C}(t, \mathbf{C}(t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2, \phi_1) \rightarrow \mathbf{C}(t_3, \phi_2)} [I_7] \\
 &\frac{}{\mathbf{C}(t, \forall x. \phi \rightarrow \phi[x \mapsto t])} [I_8] \quad \frac{}{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)} [I_9] \\
 &\frac{}{\mathbf{C}(t, [\phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi] \rightarrow [\phi_1 \rightarrow \dots \rightarrow \phi_n \rightarrow \phi])} [I_{10}] \\
 &\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \phi \rightarrow \psi)}{\mathbf{B}(a, t, \psi)} [I_{11a}] \quad \frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \psi)}{\mathbf{B}(a, t, \phi \wedge \psi)} [I_{11b}] \\
 &\frac{\mathbf{S}(s, h, t, \phi)}{\mathbf{B}(h, t, \mathbf{B}(s, t, \phi))} [I_{12}] \quad \frac{\mathbf{I}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))}{\mathbf{P}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))} [I_{13}] \\
 &\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \mathbf{O}(a, t, \phi, \chi)) \mathbf{O}(a, t, \phi, \chi)}{\mathbf{K}(a, t, \mathbf{I}(a, t, \chi))} [I_{14}]
 \end{aligned}$$

Formalizing & implementing this ...

DCEC Signature

$$\begin{aligned}
 S &::= \text{Agent} \mid \text{ActionType} \mid \text{Action} \sqsubseteq \text{Event} \mid \text{Moment} \mid \text{Fluent} \\
 f &::= \begin{cases} \text{action} : \text{Agent} \times \text{ActionType} \rightarrow \text{Action} \\ \text{initially} : \text{Fluent} \rightarrow \text{Formula} \\ \text{holds} : \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{happens} : \text{Event} \times \text{Moment} \rightarrow \text{Formula} \\ \text{clipped} : \text{Moment} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{initiates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{terminates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{prior} : \text{Moment} \times \text{Moment} \rightarrow \text{Formula} \end{cases} \\
 t &::= x : S \mid c : S \mid f(t_1, \dots, t_n) \\
 \phi &::= \begin{cases} q : \text{Formula} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \forall x : \phi(x) \mid \exists x : \phi(x) \\ \mathbf{P}(a, t, \phi) \mid \mathbf{K}(a, t, \phi) \mid \mathbf{S}(a, b, t, \phi) \mid \mathbf{S}(a, t, \phi) \\ \mathbf{C}(t, \phi) \mid \mathbf{B}(a, t, \phi) \mid \mathbf{D}(a, t, \phi) \mid \mathbf{I}(a, t, \phi) \\ \mathbf{O}(a, t, \phi, (\neg)\text{happens}(\text{action}(a^*, \alpha), t')) \end{cases}
 \end{aligned}$$

Modal Operator Descriptors:
Perceives, **K**nows, **S**ays, **C**ommon-knowledge
Believes, **D**esires, **I**ntends, **O**ught-to



DCEC Inference Schemata

$$\begin{aligned}
 &\frac{\mathbf{K}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{K}(a, t_2, \phi)} [I_K] \quad \frac{\mathbf{B}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{B}(a, t_2, \phi)} [I_B] \\
 &\frac{\mathbf{C}(t, \mathbf{P}(a, t, \phi) \rightarrow \mathbf{K}(a, t, \phi))}{\mathbf{C}(t, \phi) t \leq t_1 \dots t \leq t_n} [I_1] \quad \frac{\mathbf{C}(t, \mathbf{K}(a, t, \phi) \rightarrow \mathbf{B}(a, t, \phi))}{\mathbf{K}(a_1, t_1, \dots \mathbf{K}(a_n, t_n, \phi) \dots)} [I_2] \\
 &\frac{\mathbf{C}(t, \phi) t \leq t_1 \dots t \leq t_n}{\mathbf{K}(a_1, t_1, \dots \mathbf{K}(a_n, t_n, \phi) \dots)} [I_3] \quad \frac{\mathbf{K}(a, t, \phi)}{\phi} [I_4] \\
 &\frac{\mathbf{C}(t, \mathbf{K}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{K}(a, t_2, \phi_1) \rightarrow \mathbf{K}(a, t_3, \phi_2)}{\mathbf{C}(t, \mathbf{B}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a, t_2, \phi_1) \rightarrow \mathbf{B}(a, t_3, \phi_2)} [I_5] \\
 &\frac{\mathbf{C}(t, \mathbf{B}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a, t_2, \phi_1) \rightarrow \mathbf{B}(a, t_3, \phi_2)}{\mathbf{C}(t, \mathbf{C}(t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2, \phi_1) \rightarrow \mathbf{C}(t_3, \phi_2)} [I_6] \\
 &\frac{\mathbf{C}(t, \mathbf{C}(t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2, \phi_1) \rightarrow \mathbf{C}(t_3, \phi_2)}{\mathbf{C}(t, \forall x. \phi \rightarrow \phi[x \mapsto t])} [I_7] \quad \frac{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)}{\mathbf{C}(t, \forall x. \phi \rightarrow \phi[x \mapsto t])} [I_8] \\
 &\frac{\mathbf{C}(t, [\phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi] \rightarrow [\phi_1 \rightarrow \dots \rightarrow \phi_n \rightarrow \phi])}{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)} [I_9] \\
 &\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \phi \rightarrow \psi)}{\mathbf{B}(a, t, \psi)} [I_{10}] \quad \frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \psi)}{\mathbf{B}(a, t, \phi \wedge \psi)} [I_{11a}] \quad \frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \psi)}{\mathbf{B}(a, t, \phi \wedge \psi)} [I_{11b}] \\
 &\frac{\mathbf{S}(s, h, t, \phi)}{\mathbf{B}(h, t, \mathbf{B}(s, t, \phi))} [I_{12}] \quad \frac{\mathbf{I}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))}{\mathbf{P}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))} [I_{13}] \\
 &\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \mathbf{O}(a, t, \phi, \chi)) \mathbf{O}(a, t, \phi, \chi)}{\mathbf{K}(a, t, \mathbf{I}(a, t, \chi))} [I_{14}]
 \end{aligned}$$

Formalizing & implementing this ...

DCEC Signature

$S ::= \text{Agent} \mid \text{ActionType} \mid \text{Action} \sqsubseteq \text{Event} \mid \text{Moment} \mid \text{Fluent}$
 $f ::= \begin{cases} \text{action} : \text{Agent} \times \text{ActionType} \rightarrow \text{Action} \\ \text{initially} : \text{Fluent} \rightarrow \text{Formula} \\ \text{holds} : \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{happens} : \text{Event} \times \text{Moment} \rightarrow \text{Formula} \\ \text{clipped} : \text{Moment} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{initiates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{terminates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{prior} : \text{Moment} \times \text{Moment} \rightarrow \text{Formula} \end{cases}$
 $t ::= x : S \mid c : S \mid f(t_1, \dots, t_n)$
 $\phi ::= \begin{cases} q : \text{Formula} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \forall x : \phi(x) \mid \exists x : \phi(x) \\ \mathbf{P}(a, t, \phi) \mid \mathbf{K}(a, t, \phi) \mid \mathbf{S}(a, b, t, \phi) \mid \mathbf{S}(a, t, \phi) \\ \mathbf{C}(t, \phi) \mid \mathbf{B}(a, t, \phi) \mid \mathbf{D}(a, t, \phi) \mid \mathbf{I}(a, t, \phi) \\ \mathbf{O}(a, t, \phi, (\neg)\text{happens}(\text{action}(a^*, \alpha), t')) \end{cases}$

Modal Operator Descriptors:
Perceives, Knows, Says, Common-knowledge
Believes, Desires, Intends, Ought-to

DCEC Inference Schemata

$\frac{\mathbf{K}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{K}(a, t_2, \phi)} [I_K] \quad \frac{\mathbf{B}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{B}(a, t_2, \phi)} [I_B]$
 $\frac{\mathbf{C}(t, \mathbf{P}(a, t, \phi) \rightarrow \mathbf{K}(a, t, \phi))}{\mathbf{C}(t, \phi) t \leq t_1 \dots t \leq t_n} [I_1] \quad \frac{\mathbf{C}(t, \mathbf{K}(a, t, \phi) \rightarrow \mathbf{B}(a, t, \phi))}{\mathbf{C}(t, \phi) t \leq t_1 \dots t \leq t_n} [I_2]$
 $\frac{\mathbf{K}(a_1, t_1, \dots, \mathbf{K}(a_n, t_n, \phi) \dots)}{\mathbf{K}(a, t, \phi)} [I_3] \quad \frac{\mathbf{K}(a, t, \phi)}{\phi} [I_4]$
 $\frac{\mathbf{C}(t, \mathbf{K}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{K}(a, t_2, \phi_1) \rightarrow \mathbf{K}(a, t_3, \phi_2)}{\mathbf{C}(t, \mathbf{K}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{K}(a, t_2, \phi_1) \rightarrow \mathbf{K}(a, t_3, \phi_2)} [I_5]$
 $\frac{\mathbf{C}(t, \mathbf{B}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a, t_2, \phi_1) \rightarrow \mathbf{B}(a, t_3, \phi_2)}{\mathbf{C}(t, \mathbf{B}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a, t_2, \phi_1) \rightarrow \mathbf{B}(a, t_3, \phi_2)} [I_6]$
 $\frac{\mathbf{C}(t, \mathbf{C}(t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2, \phi_1) \rightarrow \mathbf{C}(t_3, \phi_2)}{\mathbf{C}(t, \mathbf{C}(t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2, \phi_1) \rightarrow \mathbf{C}(t_3, \phi_2)} [I_7]$
 $\frac{\mathbf{C}(t, \forall x. \phi \rightarrow \phi[x \mapsto t])}{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)} [I_8] \quad \frac{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)}{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)} [I_9]$
 $\frac{\mathbf{C}(t, [\phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi] \rightarrow [\phi_1 \rightarrow \dots \rightarrow \phi_n \rightarrow \phi])}{\mathbf{C}(t, [\phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi] \rightarrow [\phi_1 \rightarrow \dots \rightarrow \phi_n \rightarrow \phi])} [I_{10}]$
 $\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \phi \rightarrow \psi)}{\mathbf{B}(a, t, \psi)} [I_{11a}] \quad \frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \psi)}{\mathbf{B}(a, t, \phi \wedge \psi)} [I_{11b}]$
 $\frac{\mathbf{S}(s, h, t, \phi)}{\mathbf{B}(h, t, \mathbf{B}(s, t, \phi))} [I_{12}] \quad \frac{\mathbf{I}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))}{\mathbf{P}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))} [I_{13}]$
 $\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \mathbf{O}(a, t, \phi, \chi)) \mathbf{O}(a, t, \phi, \chi)}{\mathbf{K}(a, t, \mathbf{I}(a, t, \chi))} [I_{14}]$

2.2 Inductive Deontic Cognitive Event Calculus

DCEC employs no uncertainty system (e.g., probability measures, *strength factors*, or likelihood measures) and hence is purely deductive. Therefore, as we wish to enable our agents to reason about situations involving uncertainty, we must ultimately utilize the *Inductive DCEC*: *IDCEC*.

In general, to go from a deductive to an inductive cognitive calculus, we require two components: (1) an uncertainty system, and (2) inference schemata that delineate the methods by which inferences linking formulae and other information can be used to build formally valid arguments.

The particular uncertainty system we use herein is discussed in §2.3. The inference schemata of *IDCEC* consist of the union of the set presented in §2.1 with that in the box titled **Additional Inference Schemata for *IDCEC***. Likewise, the signature of *IDCEC* subsumes that of the deductive *DCEC*; the syntax of *IDCEC* also includes the forms given in the box titled **Additional Syntax for *IDCEC***.

Additional Syntax for *IDCEC*

$\phi ::= \{ \mathbf{B}^\sigma(a, t, \phi) \}$
 where $\sigma \in [-5, -4, \dots, 4, 5]$

Additional Inference Schemata for *IDCEC*

$\frac{\mathbf{P}(a, t_1, \phi_1), \Gamma \vdash t_1 < t_2}{\mathbf{B}^4(a, t_2, \phi)} [I_P^s]$
 $\frac{\mathbf{B}^{\sigma_1}(a, t_1, \phi_1), \dots, \mathbf{B}^{\sigma_m}(a, t_m, \phi_m), \{\phi_1, \dots, \phi_m\} \vdash \phi, \{\phi_1, \dots, \phi_m\} \not\vdash \zeta, \Gamma \vdash t_i < t}{\mathbf{B}^{\min(\sigma_1, \dots, \sigma_m)}(a, t, \phi)} [I_B^s]$
 where $\sigma \in [0, 1, \dots, 5, 6]$
 $\frac{}{\mathbf{C}(t, \mathbf{B}^{-\sigma}(a, t, \phi) \leftrightarrow \mathbf{B}^\sigma(a, t, \neg\phi))} [I_\neg^s]$

Briefly, $\mathbf{B}^\sigma(a, t, \phi)$ denotes that agent a at time t believes ϕ with uncertainty σ . We justify in the next section the range of values for σ .

The first inference schema allows agents to infer evident beliefs ($\sigma = 4$, as defined in the next section) from what they perceive.⁵ The second schema allows agents to infer a belief that is provable from the beliefs they currently assert, so long as the belief set is not inconsistent. In practice, we usually check that the belief set is consistent by attempting to prove a reserved propositional atom ζ which does not

Formalizing & implementing this ...

DCEC Signature

$S ::= \text{Agent} \mid \text{ActionType} \mid \text{Action} \sqsubseteq \text{Event} \mid \text{Moment} \mid \text{Fluent}$
 $f ::= \begin{cases} \text{action} : \text{Agent} \times \text{ActionType} \rightarrow \text{Action} \\ \text{initially} : \text{Fluent} \rightarrow \text{Formula} \\ \text{holds} : \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{happens} : \text{Event} \times \text{Moment} \rightarrow \text{Formula} \\ \text{clipped} : \text{Moment} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{initiates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{terminates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Formula} \\ \text{prior} : \text{Moment} \times \text{Moment} \rightarrow \text{Formula} \end{cases}$
 $t ::= x : S \mid c : S \mid f(t_1, \dots, t_n)$
 $\phi ::= \begin{cases} q : \text{Formula} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \forall x : \phi(x) \mid \exists x : \phi(x) \\ \mathbf{P}(a, t, \phi) \mid \mathbf{K}(a, t, \phi) \mid \mathbf{S}(a, b, t, \phi) \mid \mathbf{S}(a, t, \phi) \\ \mathbf{C}(t, \phi) \mid \mathbf{B}(a, t, \phi) \mid \mathbf{D}(a, t, \phi) \mid \mathbf{I}(a, t, \phi) \\ \mathbf{O}(a, t, \phi, (\neg)\text{happens}(\text{action}(a^*, \alpha), t')) \end{cases}$

Modal Operator Descriptors:
Perceives, Knows, Says, Common-knowledge
Believes, Desires, Intends, Ought-to

DCEC Inference Schemata

$\frac{\mathbf{K}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{K}(a, t_2, \phi)} [I_K] \quad \frac{\mathbf{B}(a, t_1, \Gamma), \Gamma \vdash \phi, t_1 \leq t_2}{\mathbf{B}(a, t_2, \phi)} [I_B]$
 $\frac{\mathbf{C}(t, \mathbf{P}(a, t, \phi) \rightarrow \mathbf{K}(a, t, \phi))}{\mathbf{C}(t, \phi) t \leq t_1 \dots t \leq t_n} [I_1] \quad \frac{\mathbf{C}(t, \mathbf{K}(a, t, \phi) \rightarrow \mathbf{B}(a, t, \phi))}{\mathbf{C}(t, \phi) t \leq t_1 \dots t \leq t_n} [I_2]$
 $\frac{\mathbf{K}(a_1, t_1, \dots, \mathbf{K}(a_n, t_n, \phi) \dots)}{\mathbf{K}(a, t, \phi)} [I_3] \quad \frac{\mathbf{K}(a, t, \phi)}{\phi} [I_4]$
 $\frac{\mathbf{C}(t, \mathbf{K}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{K}(a, t_2, \phi_1) \rightarrow \mathbf{K}(a, t_3, \phi_2)}{\mathbf{C}(t, \mathbf{K}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{K}(a, t_2, \phi_1) \rightarrow \mathbf{K}(a, t_3, \phi_2)} [I_5]$
 $\frac{\mathbf{C}(t, \mathbf{B}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a, t_2, \phi_1) \rightarrow \mathbf{B}(a, t_3, \phi_2)}{\mathbf{C}(t, \mathbf{B}(a, t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a, t_2, \phi_1) \rightarrow \mathbf{B}(a, t_3, \phi_2)} [I_6]$
 $\frac{\mathbf{C}(t, \mathbf{C}(t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2, \phi_1) \rightarrow \mathbf{C}(t_3, \phi_2)}{\mathbf{C}(t, \mathbf{C}(t_1, \phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2, \phi_1) \rightarrow \mathbf{C}(t_3, \phi_2)} [I_7]$
 $\frac{\mathbf{C}(t, \forall x. \phi \rightarrow \phi[x \mapsto t])}{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)} [I_8] \quad \frac{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)}{\mathbf{C}(t, \phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)} [I_9]$
 $\frac{\mathbf{C}(t, [\phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi] \rightarrow [\phi_1 \rightarrow \dots \rightarrow \phi_n \rightarrow \phi])}{\mathbf{C}(t, [\phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi] \rightarrow [\phi_1 \rightarrow \dots \rightarrow \phi_n \rightarrow \phi])} [I_{10}]$
 $\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \phi \rightarrow \psi)}{\mathbf{B}(a, t, \psi)} [I_{11a}] \quad \frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \psi)}{\mathbf{B}(a, t, \phi \wedge \psi)} [I_{11b}]$
 $\frac{\mathbf{S}(s, h, t, \phi)}{\mathbf{B}(h, t, \mathbf{B}(s, t, \phi))} [I_{12}] \quad \frac{\mathbf{I}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))}{\mathbf{P}(a, t, \text{happens}(\text{action}(a^*, \alpha), t'))} [I_{13}]$
 $\frac{\mathbf{B}(a, t, \phi) \mathbf{B}(a, t, \mathbf{O}(a, t, \phi, \chi)) \mathbf{O}(a, t, \phi, \chi)}{\mathbf{K}(a, t, \mathbf{I}(a, t, \chi))} [I_{14}]$

2.2 Inductive Deontic Cognitive Event Calculus

DCEC employs no uncertainty system (e.g., probability measures, *strength factors*, or likelihood measures) and hence is purely deductive. Therefore, as we wish to enable our agents to reason about situations involving uncertainty, we must ultimately utilize the *Inductive DCEC*: *IDCEC*.

In general, to go from a deductive to an inductive cognitive calculus, we require two components: (1) an uncertainty system, and (2) inference schemata that delineate the methods by which inferences linking formulae and other information can be used to build formally valid arguments.

The particular uncertainty system we use herein is discussed in §2.3. The inference schemata of *IDCEC* consist of the union of the set presented in §2.1 with that in the box titled **Additional Inference Schemata for *IDCEC***. Likewise, the signature of *IDCEC* subsumes that of the deductive *DCEC*; the syntax of *IDCEC* also includes the forms given in the box titled **Additional Syntax for *IDCEC***.

Additional Syntax for *IDCEC*

$\phi ::= \{ \mathbf{B}^\sigma(a, t, \phi) \}$
 where $\sigma \in [-5, -4, \dots, 4, 5]$

Additional Inference Schemata for *IDCEC*

$\frac{\mathbf{P}(a, t_1, \phi_1), \Gamma \vdash t_1 < t_2}{\mathbf{B}^4(a, t_2, \phi)} [I_P^s]$
 $\frac{\mathbf{B}^{\sigma_1}(a, t_1, \phi_1), \dots, \mathbf{B}^{\sigma_m}(a, t_m, \phi_m), \{\phi_1, \dots, \phi_m\} \vdash \phi, \{\phi_1, \dots, \phi_m\} \not\vdash \zeta, \Gamma \vdash t_i < t}{\mathbf{B}^{\min(\sigma_1, \dots, \sigma_m)}(a, t, \phi)} [I_B^s]$
 where $\sigma \in [0, 1, \dots, 5, 6]$
 $\frac{}{\mathbf{C}(t, \mathbf{B}^{-\sigma}(a, t, \phi) \leftrightarrow \mathbf{B}^\sigma(a, t, \neg\phi))} [I_\neg^s]$

Briefly, $\mathbf{B}^\sigma(a, t, \phi)$ denotes that agent a at time t believes ϕ with uncertainty σ . We justify in the next section the range of values for σ .

The first inference schema allows agents to infer evident beliefs ($\sigma = 4$, as defined in the next section) from what they perceive.⁵ The second schema allows agents to infer a belief that is provable from the beliefs they currently assert, so long as the belief set is not inconsistent. In practice, we usually check that the belief set is consistent by attempting to prove a reserved propositional atom ζ which does not

Loggik kan hjelpe deg å leve for alltid.