Universal Decision Models

Sridhar Mahadevan

Adobe Research, U.Mass, Amherst
www.cics.umass.edu/∼mahadeva

2021 年 10 月 23 日

Adobe Research

UMassAmherst
College of Information & Computer Sciences

Universal Decision Model

- Unify Causal inference, Game theory, Reinforcement Learning.
- Define *universal properties* of *information structures* underlying UDMs.
- Paper coming soon on Arxiv and my UMass web page: `www.cics.umass.edu/~mahadeva`

UDMs are based on Witsenhausen's Intrinsic Model

SIAM J. CONTROL
Vol. 9, No. 2, May 1971

ON INFORMATION STRUCTURES, FEEDBACK AND CAUSALITY*

H. S. WITSENHAUSEN†

Abstract. A finite number of decisions, indexed by $\alpha \in A$, are to be taken. Each decision amounts to selecting a point in a measurable space $(U_\alpha, \mathscr{F}_\alpha)$. Each decision is based on some information fed back from the system and characterized by a subfield $\mathscr{I}_\alpha$ of the product space $(\prod_\alpha U_\alpha, \prod_\alpha \mathscr{F}_\alpha)$. The decision function for each α can be any function γ_α measurable from $\mathscr{I}_\alpha$ to $\mathscr{F}_\alpha$.

A property of the $\{\mathscr{I}_\alpha\}_{\alpha \in A}$ is defined which assures that the setup has a causal interpretation. This property implies that for any combination of choices of the γ_α, the closed loop equations have a unique solution.

The converse implication is false, when card $A > 2$.

1. Introduction. In control-oriented works on dynamic games (in particular, stochastic control problems) one usually finds a "dynamic equation" describing the evolution of a "state" in response to decision (control) variables of the players and to random variables. One also finds "output equations" which define output variables for a player as functions of the state, decision and random variables. Then the information structure is defined by allowing each decision variable to be any desired (measurable) function of the output variables generated for that player up to that time.

Such a setup assumes that the time order in which the various decision variables are selected is fixed in advance. It assumes that each player acts as if he had responsibility only for one station. It assumes that this station has perfect memory.

For large complex systems these tacit assumptions are unlikely to hold.

Universal Decision Models

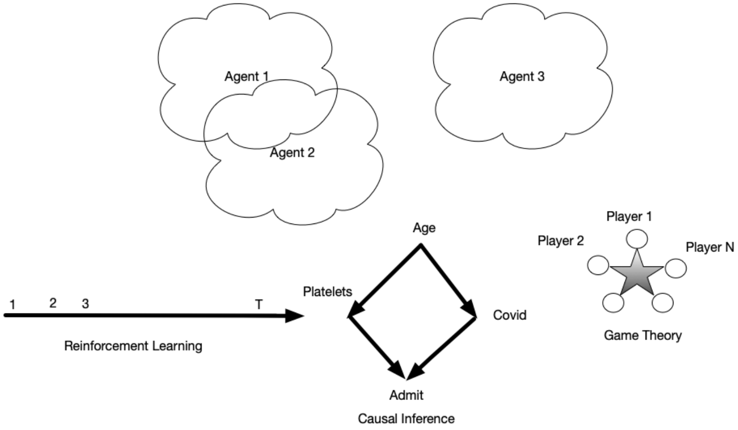

Group Decision Making in Honeybees

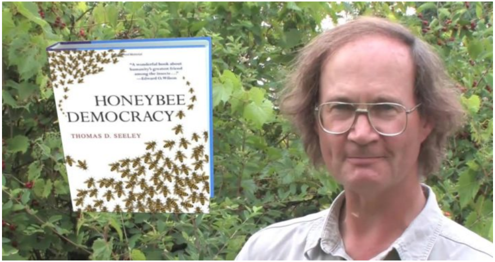

Decentralized Decision Making in Cloud Computing

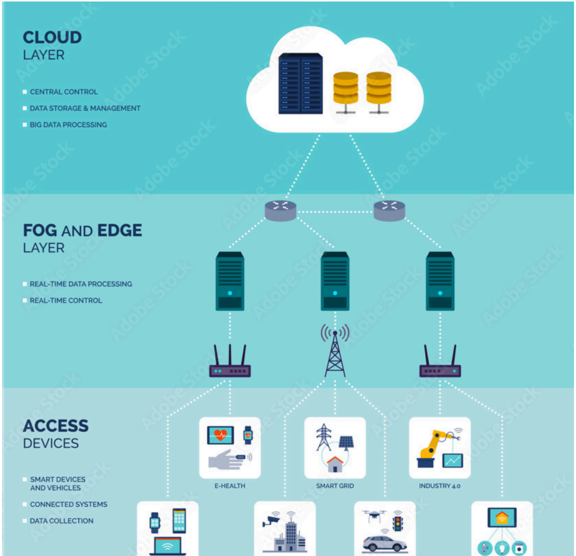

Universal Decision Model

UDM: $\langle A, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_\alpha)_{\alpha \in A} \rangle$:

- A: finite universe of decision points (e.g., agents, exogenous/endogenous variables, states, time)
- $(\Omega, \mathcal{B}, P)$: probability space representing the inherent stochastic state of nature due to randomness
- $(U_\alpha, \mathcal{F}_\alpha)$: measurable space from which a decision $u \in U_\alpha$ is chosen by α, where $\mathcal{F}_\alpha$ is a σ-algebra over U_α
- Product space: $H = \times_{a \in A} U_\alpha$, product field $\mathcal{F} = \otimes_{\alpha \in A} \mathcal{F}_\alpha \otimes \mathcal{B}$
- Policy $\pi_\alpha : \Omega \times H \to U_\alpha$ is a measurable function over $(\mathcal{I}_\alpha, \mathcal{F}_\alpha)$
- Information field $\mathcal{I}_\alpha$: subfield of the overall product field $\mathcal{F}$

Two-Player Game as a UDM

A partial information game $\mathcal{G} = \langle A, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha)_{\alpha \in A} \rangle$:

- Set of players A, with probability space $(\Omega, \mathcal{B}, P)$
- Decision space: $(U_\alpha, \mathcal{F}_\alpha)_{\alpha \in A}$, where $\mathcal{F}_\alpha$ is a partition of Ω.
- Simple two-player game: $A = \alpha, \beta$.
  - State of nature: $\Omega = \{1, 2, \ldots, 9\}$, $\mathcal{B} = 2^\Omega$, $P\{i : i \in \Omega\} = \frac{1}{9}$.
  - Information partition: $\mathcal{F}_\alpha = \{\{1, 2, 3\}, \{4, 5, 6\}, \{7, 8, 9\}\}$.
  - Information partition: $\mathcal{F}_\beta = \{\{1, 2, 3, 4\}, \{5, 6, 7, 8\}, \{9\}\}$.
  - Suppose true state of nature is $\omega \in \Omega = 1$
  - Knowledge of α: $\mathcal{F}_\alpha^1 = \{1, 2, 3\}$
  - Knowledge of β: $\mathcal{F}_\beta^1 = \{1, 2, 3, 4\}$

Causal Inference as a UDM

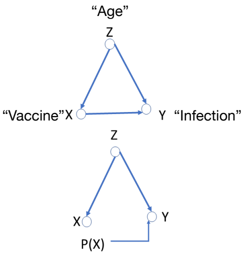

- $A = \{X, Y, Z\}$, $U_X = U_Y = U_Z = \{0, 1\}$.
- σ-algebras: $\mathcal{F}_X = \mathcal{F}_Y = \mathcal{F}_Z = \{\emptyset, \{0\}, \{1\}, \{0, 1\}\}$.
- States of nature: $\Omega = \{0, 1\}^3$, Borel topology $\mathcal{B} = 2^\Omega$.
- Policies:
  $\pi_\alpha^{-1} \subset \mathcal{F}_\alpha \otimes_{b \neq \alpha} \{\emptyset, \Omega_b\} \otimes_{b \in \mathrm{Pa}(\alpha)} \mathcal{F}_b \otimes_{b \notin \mathrm{Pa}(\alpha)} \{\emptyset, U_b\}$

For more details, see [Heymann et al., Arxiv, 2021].

Sequential UDMs

1. Probability space: $(\Omega, \mathcal{B}, P)$
2. Measurable decision spaces $(U_t, \mathcal{F}_t), t = 1, \ldots, T$ at each time point.
3. Information fields $\mathcal{I}_t \subset \mathcal{B} \otimes \mathcal{F}_1 \otimes \ldots \otimes \mathcal{F}_T$
4. Permutation $p : \{1, \ldots, T\} \to \{1, \ldots, T\}$ such that for $t = 1, \ldots T$, the information field $\mathcal{I}_t \subset \mathcal{B} \otimes \mathcal{F}_{p(1)} \otimes \mathcal{F}_{p(2)}, \ldots, \mathcal{F}_{p(t-1)} \otimes \{\emptyset, \mathcal{F}_{p(t)}\} \otimes \ldots \otimes \{\emptyset, \mathcal{F}_{p(T)}\}$.
5. Cost function $c : (\Omega \times U_{1:T}, \mathcal{B} \otimes \mathcal{F}_{1:T}) \to (\mathbb{R}, \mathbb{B})$
6. Objective: minimize cost function $\inf_\pi E[c(\omega, U_1, \ldots, U_T)]$ exactly, or to within ϵ.

See [Nayyar et al., IEEE Trans Aut. Control, 2018]

Solvable UDM

A UDM $\langle A, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_\alpha)_{\alpha \in A} \rangle$ is said to be **solvable** if for every state of nature $\omega \in \Omega$, and every policy $\pi \in \Pi_A$, the below set of simultaneous equations has a fixed point.

$$u_\alpha = \pi_\alpha(h) \equiv \pi_\alpha(\omega, u) \tag{1}$$

Here, π_α can be viewed as a projection from the joint decision h taken by the entire ensemble of decision makers in the UDM.

Causal UDM

An UDM $\langle A, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_\alpha)$ is said to be **causal** if

- There exists $\phi : H \to S$, where S is the set of total orderings of decision makers in A,
- such that for $1 \leq k \leq n$, and any ordered set $(\alpha_1, \ldots, \alpha_k)$ of distinct elements from A, the set $E \subset H$ on which $\phi(h)$ begins with the same ordering $(\alpha_1, \ldots, \alpha_k)$ satisfies the following causality condition:

$$\forall F \in \mathcal{F}_{\alpha_k}, \quad E \cap F \in \mathcal{F}(\{\alpha_1, \ldots, \alpha_{k-1}\}) \tag{2}$$

Classes of UDMs

1. **Monic:** $A = \{\alpha\}$, $\mathcal{I}_\alpha \subset \mathcal{F}(\emptyset)$.
2. **Team:** $\mathcal{I}_\alpha \subset \mathcal{F}(\emptyset)$.
3. **Sequential:** There exists a fixed ordering $\{\alpha_1, \dots, \alpha_n\}$ of decision makers from A such that for any $1 \leq k \leq n$, it holds that $\mathcal{I}_{\alpha_k} \subset \mathcal{F}(\{\alpha_1, \dots, \alpha_{k-1}\})$.
4. **Classical:** A UDM is called classical if it is sequential, and $\mathcal{I}_0 \in \mathcal{F}(\emptyset)$, $\mathcal{I}_{k-1} \subset \mathcal{I}_k$, for all $k = 2 \dots, n$.
5. **Without self-information:** A UDM has no self-information if for all decision makers $\alpha \in A$, it holds that $\mathcal{I}_\alpha \subset \mathcal{F}(A - \{\alpha\})$.

Common Knowledge in UDMs

- The **common knowledge** for the t^{th} decision maker in a sequential UDM is defined as

$$\mathcal{C}_t = \bigcap_{s=t}^{T} \mathcal{I}_s \tag{3}$$

- Coarsening property: $\mathcal{C}_t \subset \mathcal{I}_t$: immediate from definition.
- Nestedness property: $\mathcal{C}_t \subset \mathcal{C}_{t+1}$: immediate from definition.
- Common observations: There exist observations $Z_1, \ldots, Z_T$ with Z_t taking values in a finite measurable space $(\mathcal{Z}_t, 2^{\mathcal{Z}_t})$, and $Z_t = \eta_t(\omega, U_1, \ldots, U_{t-1})$ such that $\sigma(Z_{1:t}) = \mathcal{C}_t$.

See [Nayyar et al., IEEE Trans. Aut. Control, 2018]

Category Theory

- Unifying framework that revolutionized math over the past 50-60 years.
- Instead of describing objects (e.g., sets), characterize their interactions.
- Functors map from one category to another (e.g., $f : \mathbf{Top} \rightarrow \mathbf{Grp}$).
- Universal properties characterize an object uniquely up to isomorphism
- Natural transformations map between two functors
- Yoneda lemma: fully faithful embedding of categorial objects

Categories: Objects and Morphisms

A category $\mathcal{C}$ is

- A *collection* of *objects* $X, Y, \ldots$
- A collection of *morphisms* $f, g, \ldots$, where $f : X \to Y$ is the morphism whose domain is X and co-domain is Y.
- For each pair of morphisms f, g, such that the co-domain of f is the same as the domain of g, there is a composite morphism $g \circ f$, simply defined as the composition of g and f (where f is applied first, followed by g), defined as $gf : X \to Z$.
- Each object X has associated with it an *identity* morphism $1_X : X \to X$, whose composition with any other morphism $f : X \to Y$ is defined as $1_Y f = f = f 1_X = f$.
- *Associativity*, whereby given morphisms $f : X \to Y, g : Y \to Z, h : Z \to W$, the composite morphism $hgf : X \to W$ is associative.

Examples of Categories

- **Set**: Objects are sets, morphisms are mappings on sets.
- **Top:** Topological spaces are objects, and continuous functions as its morphisms.
- **Group:** Groups are its objects, and group homomorphisms as its morphisms.
- **Graph:** Graphs are objects, and graph morphisms (mapping vertices to vertices, preserving adjacency properties) as its morphisms.
- **Poset:** Partially ordered sets as its objects and order-preserving functions as its morphisms.
- **Meas:** Measurable spaces are its objects and measurable functions as its morphisms.

Categories vs. Sets

Set theory	Category theory
set	object
subset	subobject
truth values $\{0, 1\}$	subobject classifier Ω
power set $P(A) = 2^A$	power object $P(A) = \Omega^A$
bijection	isomorphims
injection	monic arrow
surjection	epic arrow
singleton set $\{*\}$	terminal object $\mathbf{1}$
empty set $\emptyset$	initial object $\mathbf{0}$
elements of a set X	morphism $f : \mathbf{1} \to X$
	non-global element $Y \to X$
	functors, nat. transformations,
	limits, colimits, adjunctions

MDPs form a Category

Objects are MDPs: $\langle S, A, \Psi, P, R\rangle$

- S is a discrete set of states
- A is the discrete set of actions
- $\Psi \subset S \times A$ is the set of admissible state-action pairs
- $P : \Psi \times S \rightarrow [0, 1]$ is the transition probability function specifying the one-step dynamics of the model
- $R : \Psi \rightarrow \mathbb{R}$ is the expected reward function

MDP Homomorphisms

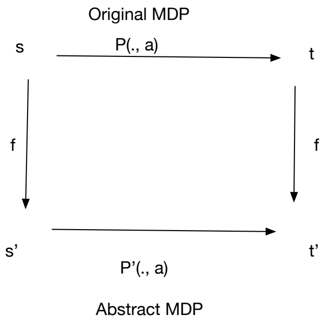

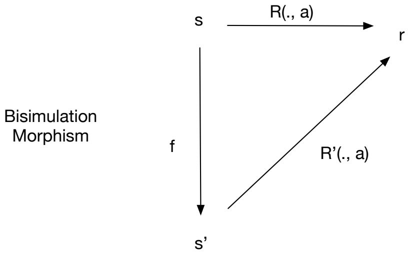

MDP Homomorphisms

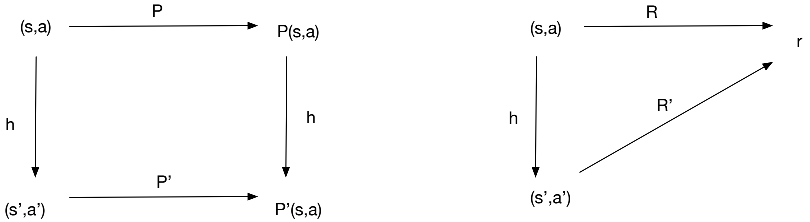

MDP Homomorphisms [Ravindran and Barto]

An MDP homomorphism from MDP $M = \langle S, A, \Psi, P, R \rangle$ to $M' = \langle S', A', \Psi', P', R' \rangle$, denoted $h : M \twoheadrightarrow M'$, is defined by

- A tuple of surjections $\langle f, \{g_s | s \in S\} \rangle$
- where $f : S \twoheadrightarrow S', g_s : A_s \twoheadrightarrow A'_{f(s)}$
- $h((s, a)) = \langle f(s), g_s(a) \rangle$, for $s \in S$
- Stochastic substitution property and reward respecting properties below are respected:

$$P'(f(s), g_s(a), f(s')) = \sum_{s" \in [s']_f} P(s, a, s") \tag{4}$$

$$R'(f(s), g_s(a)) = R(s, a) \tag{5}$$

Predictive State Representations

PSR (and earlier models, like multiplicity automata, observer operator models etc.) form categories:

- Finite set of actions A and observations O.
- A *history*: sequence of actions and observations $h = a_1 o_1 \ldots a_k o_k$.
- A *test*: possible sequence of future actions and observations $t = a_1 o_1 \ldots a_n o_n$.
- $P(t|h)$ is a prediction test t will succeed from history h.
- State ψ: a vector of predictions of *core tests* $\{q_1, \ldots, q_k\}$.
- The prediction vector $\psi_h = \langle P(q_1|h) \ldots P(q_k|h) \rangle$ is a sufficient statistic. The entire predictive state of a PSR can be denoted Ψ.

PSR Homomorphisms [Soni et al., AAAI]

A **PSR homomorphism** from a PSR Ψ to another PSR Ψ' is defined as:

- A tuple of surjections $\langle f, v_\psi(a) \rangle$
- where $f : \Psi \to \Psi'$ and $v_\psi : A \to A'$ for all prediction vectors $\psi \in \Psi$
- such that

$$P(\psi' | f(\psi), v_\psi(a)) = P(f^{-1}(\psi') | \psi, a) \quad (6)$$

for all $\psi' \in \Psi, \psi \in \Psi, a \in A$.

Bisimulation of Linear Dynamical Systems

We are given two linear dynamical systems Σ_i:

$$\begin{aligned} \dot{x}_i &= A_i x_i + B_i u_i + G_i d_i, \quad x_i \in \mathcal{X}_i, u_i \in \mathcal{U}, d_i \in \mathcal{D}_i \\ y_i &= C_i x_i, \quad y_i \in \mathcal{Y}, i = 1, 2 \end{aligned}$$

Bisimulation relation is a subspace $\mathcal{R} \subset \mathcal{X}_1 \times \mathcal{X}_2$ s.t. :

- $(x_1(0), x_2(0)) \in \mathcal{R}$
- Joint input function $u_1(.) = u_2(.)$
- For every disturbance function $d_1(.)$, there exists a $d_2(.)$ s.t.
- $(x_1(t), x_2(t)) \in \mathcal{R}, \quad \forall t \geq 0$
- $C_1 x_1(t) = C_2 x_2(t), \quad \forall t \geq 0$

Category of UDMs

- Objects are UDMs $\langle A, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_\alpha)_{\alpha \in A} \rangle$.
- Morphisms are **bisimulation** relationship between two UDMs $M = \langle A, (\Omega, \mathcal{B}, (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_\alpha)_{\alpha \in A}$ and $M' = \langle A', (\Omega', \mathcal{B}', (U'_\alpha, \mathcal{F}'_\alpha, \mathcal{I}'_\alpha)_{\alpha \in A'}$, denoted as $M \twoheadrightarrow M'$, is defined as is defined by a tuple of surjections as follows:
  - A surjection $f : A \twoheadrightarrow A'$ that maps decision points in A to corresponding points in A'.
  - As f is surjective, it induces an equivalence class in A such that $x \sim y, x, y \in A$ if and only if $f(x) = f(y)$.
  - A surjection $g : H \twoheadrightarrow H'$, where $H = \Omega \times \prod_{\alpha \in A} U_\alpha$, with the product σ-algebra $\mathcal{B} \otimes \prod_{\alpha \in A} \mathcal{F}_\alpha$, and $H' = \Omega' \otimes \prod_{\alpha \in A'} U'_\alpha$, with the corresponding σ-algebra $\mathcal{B}' \otimes \prod_{\alpha \in A'} \mathcal{F}'_\alpha$.
  - The **quotient information field** of a collection of agents $[\alpha]_f$ is defined as the join of the information fields of each agent:

$$\mathcal{I}_{[\alpha]} = \bigvee_{\beta \in [\alpha]_f} \mathcal{I}_\alpha \qquad (7)$$

Sub UDMs form a Topology

- A subset of decision makers $B \subset A$ form a **sub-UDM** if for all $\alpha \in B$, $\mathcal{I}_\alpha \subset \mathcal{F}(B)$.
- The sub-UDM $\langle B, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_{\alpha B})_{\alpha \in B} \rangle$ has an induced information subfield $\mathcal{I}_{\alpha B}$, which is the canonical projection of $\mathcal{I}_B$ upon H_B.
- The **closure** of a decision maker $\alpha \in A$ in a UDM $\langle A, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_\alpha)_{\alpha \in A} \rangle$ is the smallest sub-UDM containing α, denoted by $\overline{\{\alpha\}}$.
- The **preorder** relationship between decision makers, denoted $\alpha \leftarrow \beta$ is defined by the containment between the closure sets, namely $\alpha \leftarrow \beta$ if and only if $\overline{\{\alpha\}} \subset \overline{\{\beta\}}$.

UDM Subsystem Topology

Theorem: Sub-UDMs of a UDM $\langle A, (\Omega, \mathcal{B}, P), (U_\alpha, \mathcal{F}_\alpha, \mathcal{I}_\alpha)_{\alpha \in A} \rangle$ induce a finite space topology on the space A of decision makers.

Proof:

- Given two subsystems S_1 and S_2, if $\alpha \in S_1 \cup S_2$, then either $\mathcal{I}_\alpha \subset \mathcal{F}(S_1)$ or $\mathcal{I}_\alpha \subset \mathcal{F}(S_2)$. It follows that $\mathcal{I}_\alpha \subset \mathcal{F}(S_1) \cup \mathcal{F}(S_2) = \mathcal{F}(S_1 \cup S_2)$.
- The proof for closure under intersection is similar. $\square$

Covariant UDM Functors

A **covariant functor** $F : \mathcal{C} \to \mathcal{D}$ from UDM category $\mathcal{C}$ to category $\mathcal{D}$ is defined as the following:

- An object $\mathcal{F}X$ of the category $\mathcal{D}$ for each UDM object X in category $\mathcal{C}$.
- A morphism $\mathcal{F}f : \mathcal{F}X \to \mathcal{F}Y$ in category $\mathcal{D}$ for every bisimulation morphism $f : X \to Y$ in category $\mathcal{C}$.
- The preservation of identity and composition: $\mathcal{F}\ id_X = id_{\mathcal{F}X}$ and $(\mathcal{F}g)(\mathcal{F}f) = \mathcal{F}(gf)$ for any composable morphisms $f : X \to Y, g : Y \to Z$.

Examples of Covariant UDM Functors

- The "forgetful" functor $F : \mathcal{C}_{\text{MDP}} \rightarrow \textbf{Set}$ that maps an MDP into its set of states S.
- The "PVF" functor $F : \mathcal{C}_{\text{MDP}} \rightarrow \textbf{Graph}$ that maps an MDP into an undirected graph over states S, with an undirected edge between actual transitions.
- The "Top" functor $F : \mathcal{C}_{\text{MDP}} \rightarrow \textbf{Top}$ that maps an MDP into the category of topological spaces.
- The "Grp" functor $F : \mathcal{C}_{\text{MDP}} \rightarrow \textbf{Group}$ that maps an MDP into the category of groups.
- The "Mod" functor $F : \mathcal{C}_{\text{MDP}} \rightarrow \textbf{Modules}$ that maps an MDP into the category of modules.

Contravariant UDM Functors

- A **contravariant UDM functor** $F : \mathcal{C} \to \mathcal{D}$ from UDM category $\mathcal{C}$ to category $\mathcal{D}$ is defined exactly like the covariant UDM functor, except all the mappings are reversed.
- Contravariant functor $\mathcal{F} : C^{\mathrm{op}} \to D$, every bisimulation morphism $f : X \to Y$ is assigned the reverse morphism $\mathcal{F}f : \mathcal{F}Y \to \mathcal{F}X$ in category $\mathcal{D}$.

Functorial Representations of UDMs

- For every UDM object X in UDM category $\mathcal{C}$, there exists a covariant functor $\mathcal{C}(X, -) : \mathcal{C} \to \mathbf{Set}$ that assigns to each UDM object Z in $\mathcal{C}$ the set of bisimulation morphisms $\mathcal{C}(X, Z)$, and to each bisimulation morphism $f : Y \to Z$, the pushforward mapping $f_* : \mathcal{C}(X, Y) \to \mathcal{C}(X, Z)$.
- For every object X in UDM category $\mathcal{C}$, there exists a contravariant functor $\mathcal{C}(-, X) : \mathcal{C}^{\text{op}} \to \mathbf{Set}$ that assigns to each UDM object Z in $\mathcal{C}$ the set of bisimulation morphisms $\mathcal{C}(X, Z)$, and to each bisimulation morphism $f : Y \to Z$, the pullback mapping $f^* : \mathcal{C}(Z, X) \to \mathcal{C}(Y, X)$.

Fully Faithful UDM Representations

Let $\mathcal{F} : \mathcal{C} \to \mathcal{D}$ be a functor from UDM category $\mathcal{C}$ to category $\mathcal{D}$. If for all UDM objects X and Y in $\mathcal{C}$, the map $\mathcal{C}(X, Y) \to \mathcal{D}(\mathcal{F}X, \mathcal{F}Y)$, denoted as $f \mapsto \mathcal{F}f$ is

- injective, then the functor $\mathcal{F}$ is defined to be **faithful**.
- surjective, then the functor $\mathcal{F}$ is defined to be **full**.
- bijective, then the functor $\mathcal{F}$ is defined to be **fully faithful**.

Natural Transformations in UDMs

- Given two functors $\mathcal{F}, \mathcal{G} : \mathcal{C} \to \mathcal{D}$ that map from UDM category $\mathcal{C}$ to category $\mathcal{D}$, a *natural transformation* $\eta : \mathcal{F} \implies \mathcal{G}$ consists of a morphism $\eta_X : \mathcal{F}X \to \mathcal{G}X$ for each object X in $\mathcal{C}$.
- For any two functors $\mathcal{F}, \mathcal{G} : \mathcal{C} \to \mathcal{D}$, let $\mathrm{Nat}(\mathcal{F}, \mathcal{G})$ denote the natural transformations from $\mathcal{F}$ to $\mathcal{G}$. If $\eta_X : \mathcal{F}X \to \mathcal{G}X$ is an isomorphism for each X in category $\mathcal{C}$, then the natural transformation η is called a **natural isomorphism** and $\mathcal{F}$ and $\mathcal{G}$ are naturally isomorphic, denoted as $\mathcal{F} \cong \mathcal{G}$.

Quotients and Bisimulation morphisms

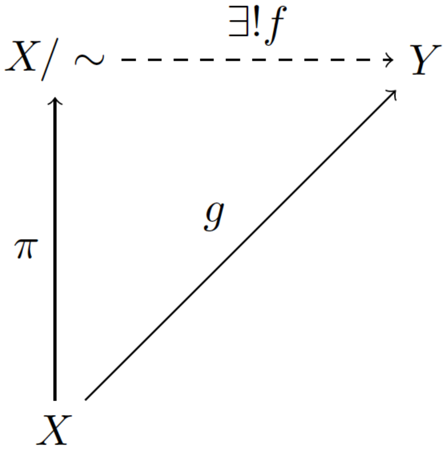

Products and Limits

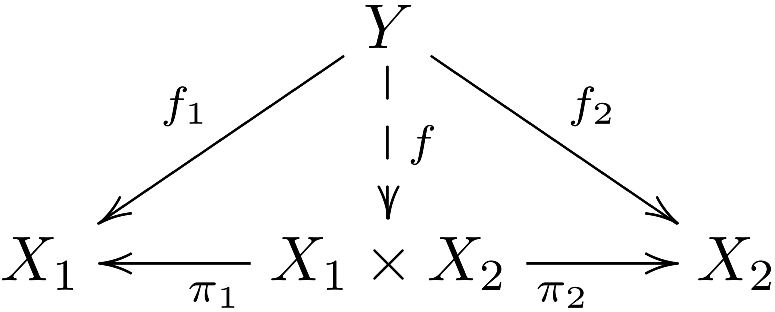

Co-Products and Co-Limits

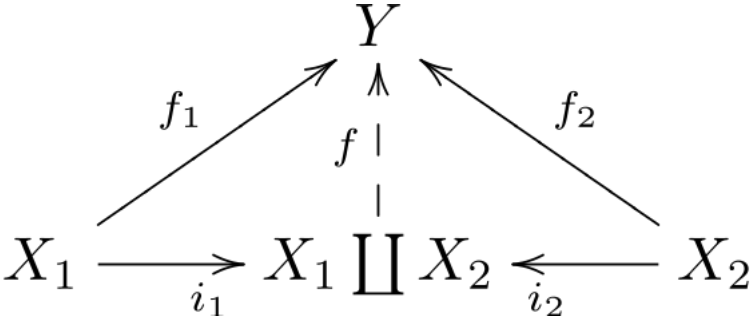

Yoneda Lemma, Pre-Sheaf, and Topoi

UDM Yoneda Lemma: For every object X in UDM category $\mathcal{C}$, and every contravariant functor $\mathcal{F} : \mathcal{C}^{\text{Op}} \to \mathbf{Set}$, the set of natural transformations from $\mathcal{C}(-, X)$ to $\mathcal{F}$ is isomorphic to $\mathcal{F}X$.

- One of the deepest results in category theory.
- A pair of UDM objects are isomorphic $X \cong Y$ if and only if the corresponding contravariant functors are isomorphic, namely $\mathcal{C}(-, X) \cong \mathcal{C}(-, Y)$.
- Given any two categories $\mathcal{C}, \mathcal{D}$, we can always define the new category $\mathcal{D}^{\mathcal{C}}$, whose objects are functors $\mathcal{C} \to \mathcal{D}$, and whose morphisms are natural transformations.
- If we take $\mathcal{D} = \mathbf{Set}$, and consider the contravariant version $\mathbf{Set}^{\mathcal{C}^{\text{Op}}}$, we obtain a category whose objects are presheafs.
- Presheafs have some very nice properties, which makes them a *topos*.

Kan Extensions of UDMs

A **left Kan extension** of a functor $H : \mathcal{C} \to \mathcal{E}$ along F, another functor $F : \mathcal{C} \to \mathcal{D}$, is a functor $\mathrm{Lan}_F H : \mathcal{D} \to \mathcal{E}$ with a natural transformation $\eta : H \implies \mathrm{Lan}_F H \circ F$ such that for any other such pair $(G : \mathcal{D} \to \mathcal{E}, \gamma : H \implies G \circ F)$, γ factors uniquely through η.

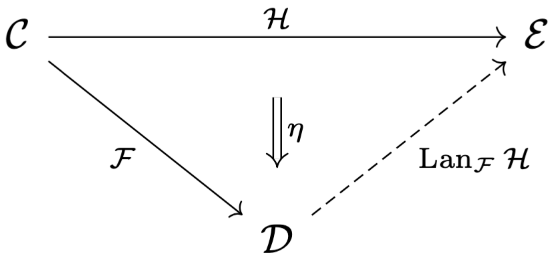

UDM: Optimization to Equilibration

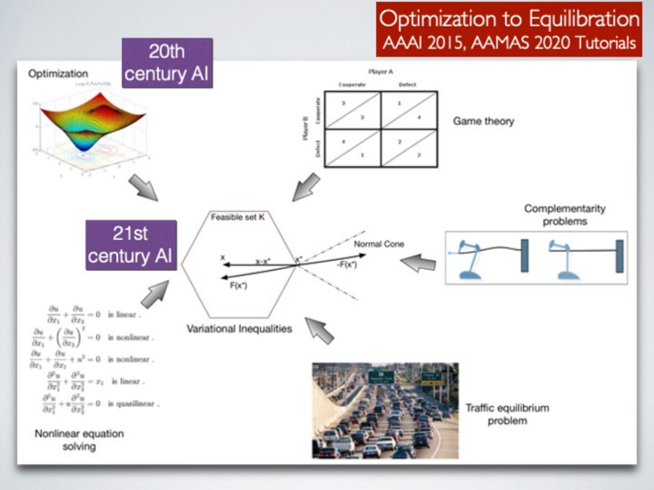

Summary

- We proposed Universal Decision Model (UDM)
  - Categorial generalization of Witsenhausen's intrinsic model
  - Universal decision making objects: n-player games, MDPs, PSRs, intrinsic models, ...
  - Morphisms: bisimulations across UDMs
  - Functors: Probe a UDM category by mapping it into a different category
  - Yoneda lemma shows how to construct fully faithful UDM embeddings
- Draft paper coming soon on Arxiv and my UMass web page (www.cics.umass.edu/ mahadeva)