Applying Learning Systems Theory to Model Cognitive Unmanned Aerial Vehicles

2023 IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW)

Tyler Cody, Ph.D.
Peter Beling, Ph.D.
Virginia Tech National Security Institute

Academic Mission

- I am a systems theorist and machine learning engineer.
- I develop and apply systems theory to bridge systems engineering and artificial intelligence.

My mission is to develop rigorous first-principles for engineering AI-heavy systems.

"In order to improve your game, you must study the endgame before everything else, for whereas the endings can be studied and mastered by themselves, the middle and the opening must be studied in relation to the endgame."

- Jose Raul Capablanca

Scope

$$
\begin{aligned}
\textbf{minimize:} \quad & f(\mathbf{x}) \\
& \mathbf{x} \in \Re^n \\
\textbf{subject to:} \quad & \mathbf{g}_L \le \mathbf{g}(\mathbf{x}) \le \mathbf{g}_U \\
& \mathbf{h}(\mathbf{x}) = \mathbf{h}_t \\
& \mathbf{a}_L \le \mathbf{A}_i\mathbf{x} \le \mathbf{a}_U \\
& \mathbf{A}_e\mathbf{x} = \mathbf{a}_t \\
& \mathbf{x}_L \le \mathbf{x} \le \mathbf{x}_U
\end{aligned}
$$

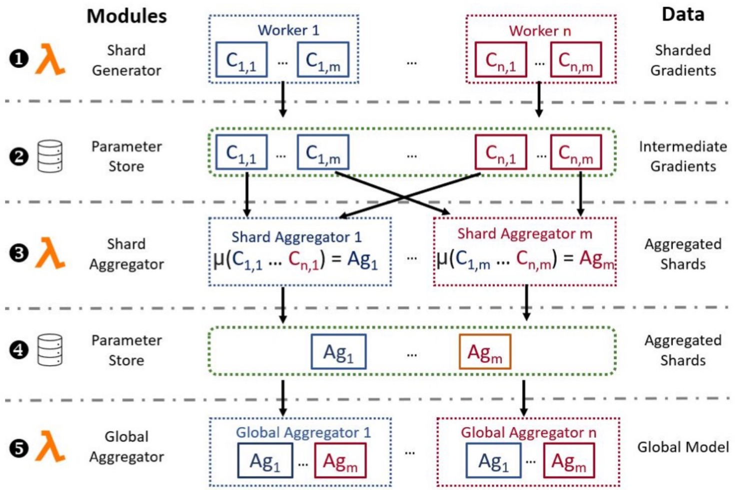

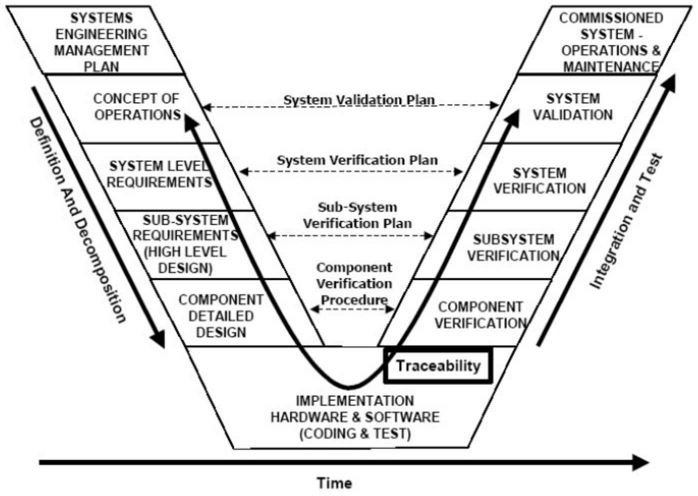

AI
Research

AI Engineering
Research

AI Systems Engineering
Research

Mesarovician Abstract Systems Theory

A system is defined as a relation on sets

$$S \subset \times\{V_i : i \in I\},$$

and theory is developed by specifying that relation and those sets to formalize phenomena of interest.

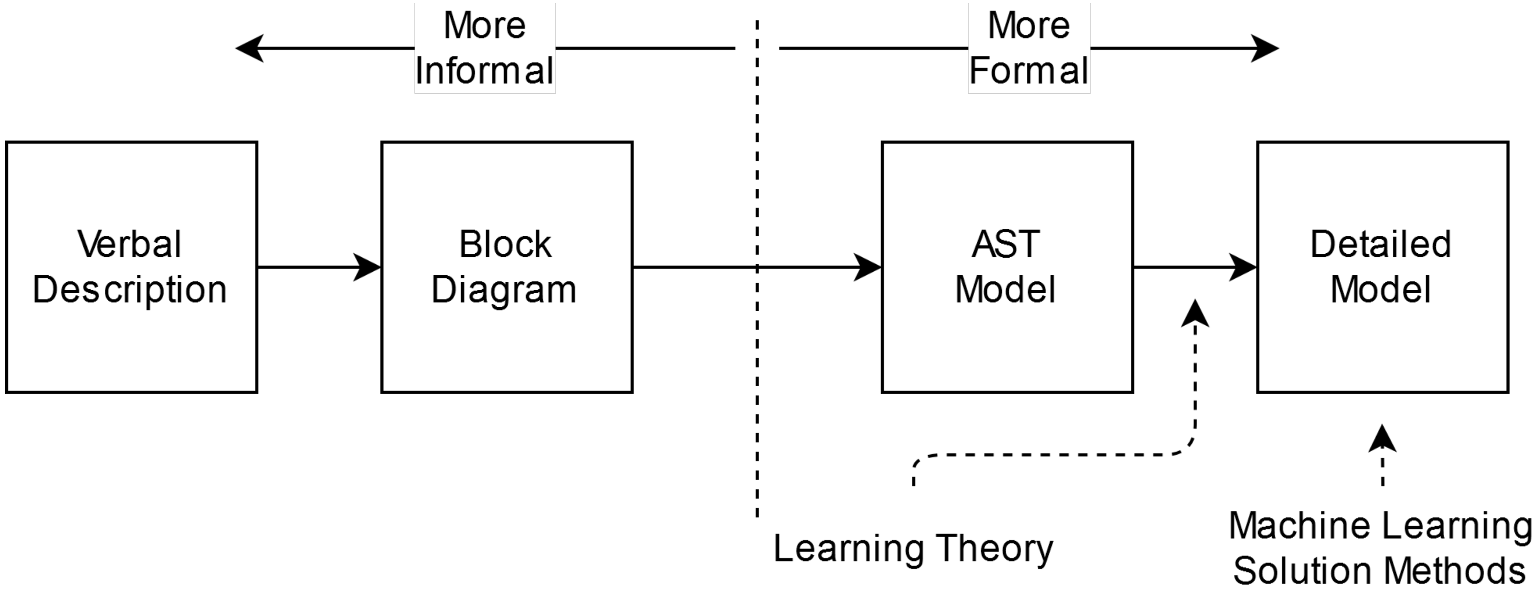

Top-Down Formalization of Learning Systems

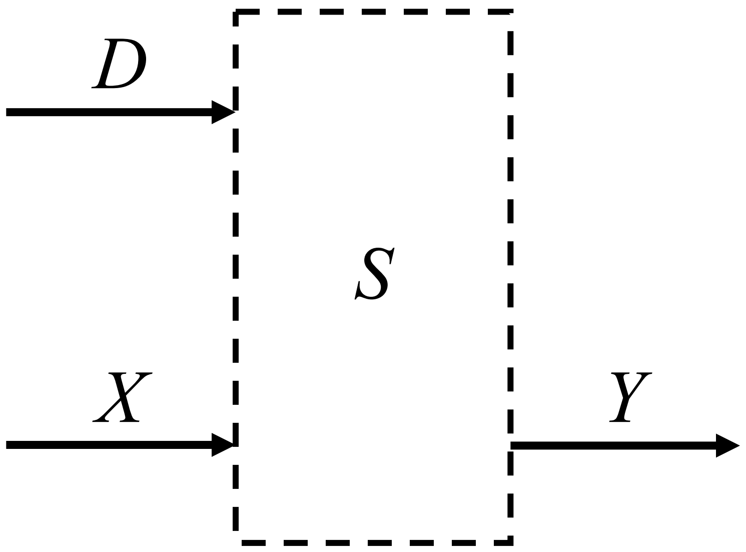

At a high level, learning systems are **input-output** systems:

$$S: D \times X \to Y.$$

Cody, Tyler. "Mesarovician abstract learning systems." *Artificial General Intelligence: 14th International Conference, AGI 2021, Palo Alto, CA, USA, October 15–18, 2021, Proceedings 14*. Springer International Publishing, 2022.

Top-Down Formalization of Learning Systems

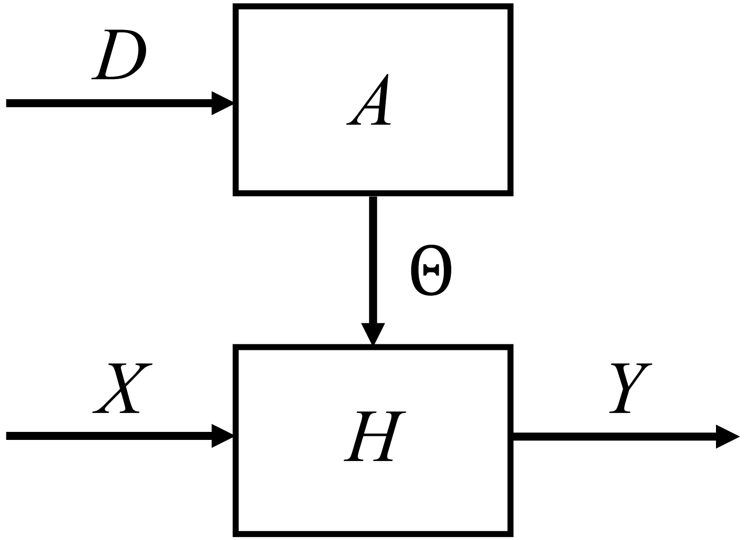

At a high level, learning systems are **input-output** systems:

$$S: D \times X \rightarrow Y.$$

A level deeper, are a **coupling of input-output systems** A and H.

Cody, Tyler. "Mesarovician abstract learning systems." *Artificial General Intelligence: 14th International Conference, AGI 2021, Palo Alto, CA, USA, October 15–18, 2021, Proceedings 14*. Springer International Publishing, 2022.

Top-Down Formalization of Learning Systems

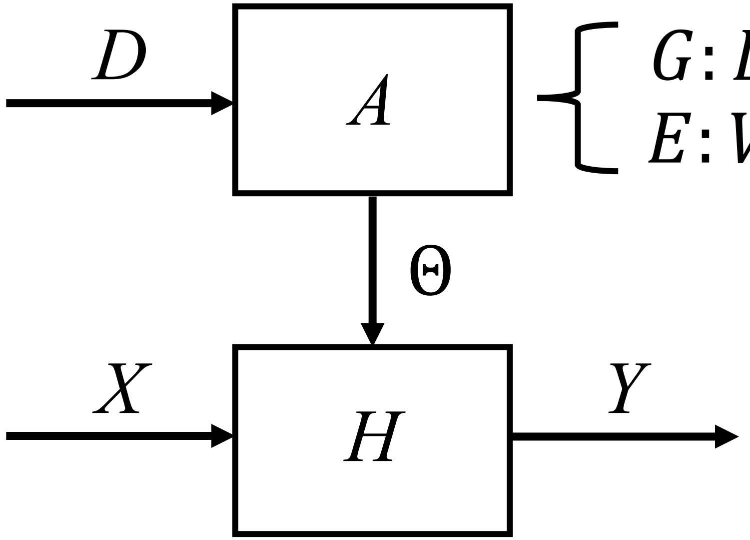

$$\begin{cases} G: D \times \Theta \to V \\ E: V \times D \to \Theta \end{cases}$$

At a high level, learning systems are **input-output** systems:

$$S: D \times X \to Y.$$

A level deeper, are a **coupling of input-output systems** A and H.

A level deeper still, we find A is not a composition, but specified by a **closed subsystem** formed by G and E.

Cody, Tyler. "Mesarovician abstract learning systems." *Artificial General Intelligence: 14th International Conference, AGI 2021, Palo Alto, CA, USA, October 15–18, 2021, Proceedings 14*. Springer International Publishing, 2022.

Top-Down Formalization of Learning Systems

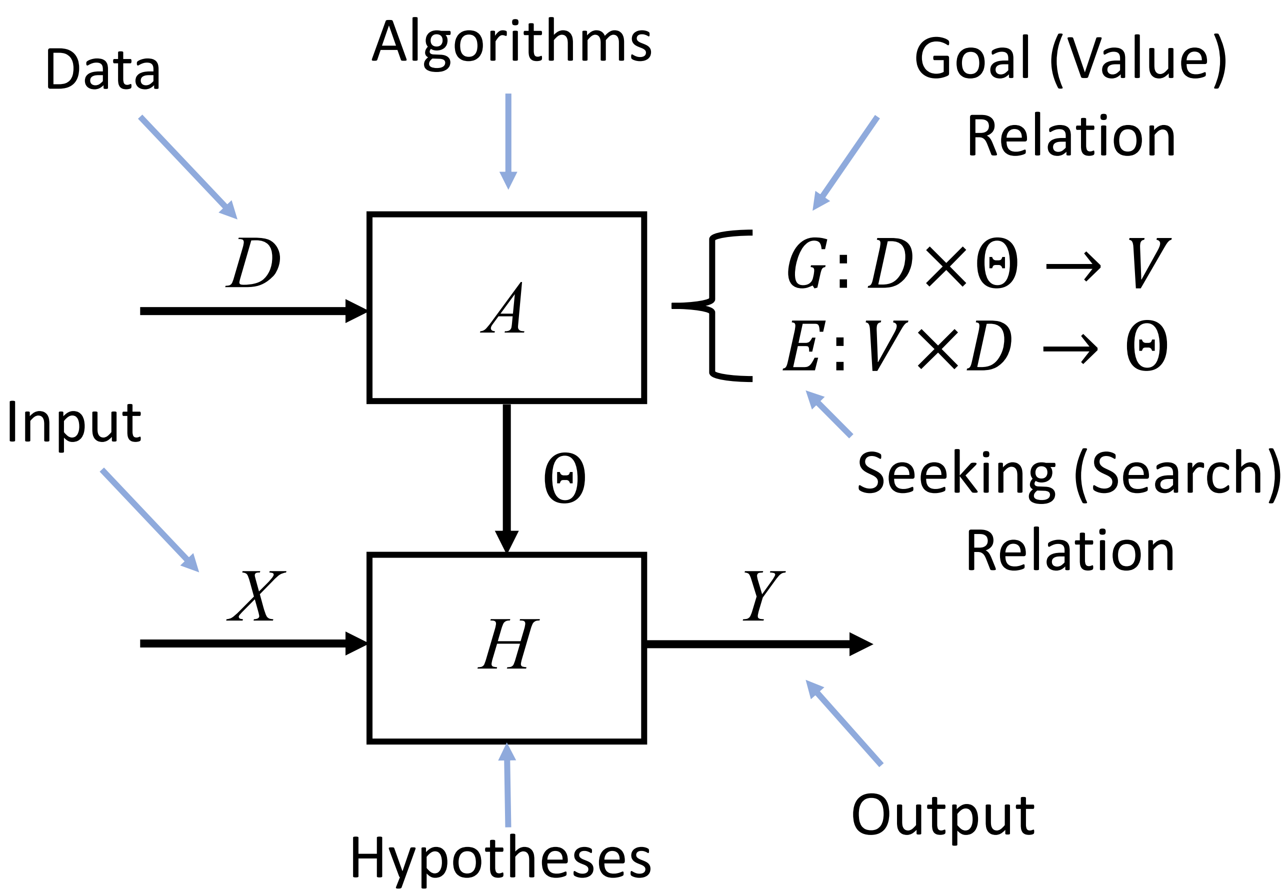

At a high level, learning systems are **input-output** systems:

$$S: D \times X \to Y.$$

A level deeper, are a **coupling of input-output systems** A and H.

A level deeper still, we find A is not a composition, but specified by a **closed subsystem** formed by G and E.

Cody, Tyler. "Mesarovician abstract learning systems." *Artificial General Intelligence: 14th International Conference, AGI 2021, Palo Alto, CA, USA, October 15–18, 2021, Proceedings 14*. Springer International Publishing, 2022.

Levels of Abstraction in Learning Systems

- Elementary Level
- Cascade Level
- Goal-Seeking Level

ALST Level | AST Diagram

Elementary Level

$S \subset \times\{D, \mathcal{X}, \mathcal{Y}\}$

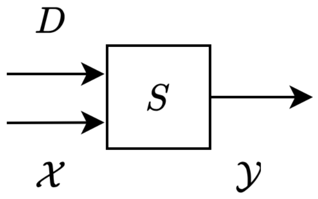

Cascade Level

$S \subset \times\{A, D, \Theta, H, \mathcal{X}, \mathcal{Y}\}$

$S_I \subset \times\{A, D, \Theta\}$
$S_F \subset \times\{\Theta, H, \mathcal{X}, \mathcal{Y}\}$

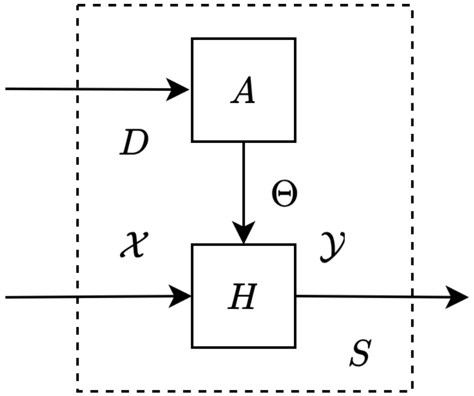

Goal-Seeking Level

$S \subset \times\{A, D, \Theta, G, E, H, \mathcal{X}, \mathcal{Y}\}$

$S_I \subset \times\{A, D, \Theta, G, E\}$
$S_F \subset \times\{\Theta, H, \mathcal{X}, \mathcal{Y}\}$

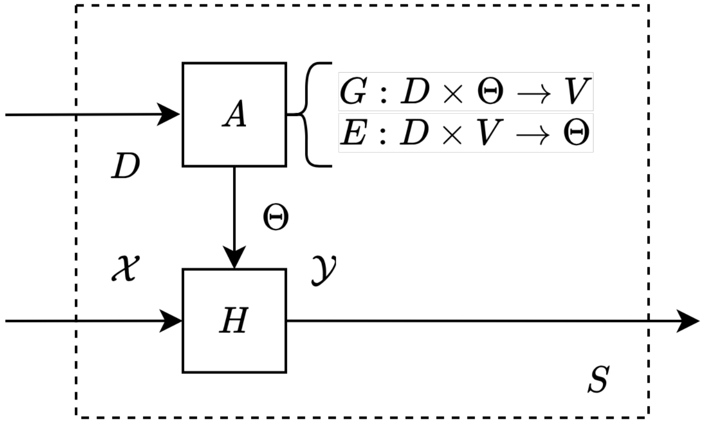

Projecting Onto Learning Theory

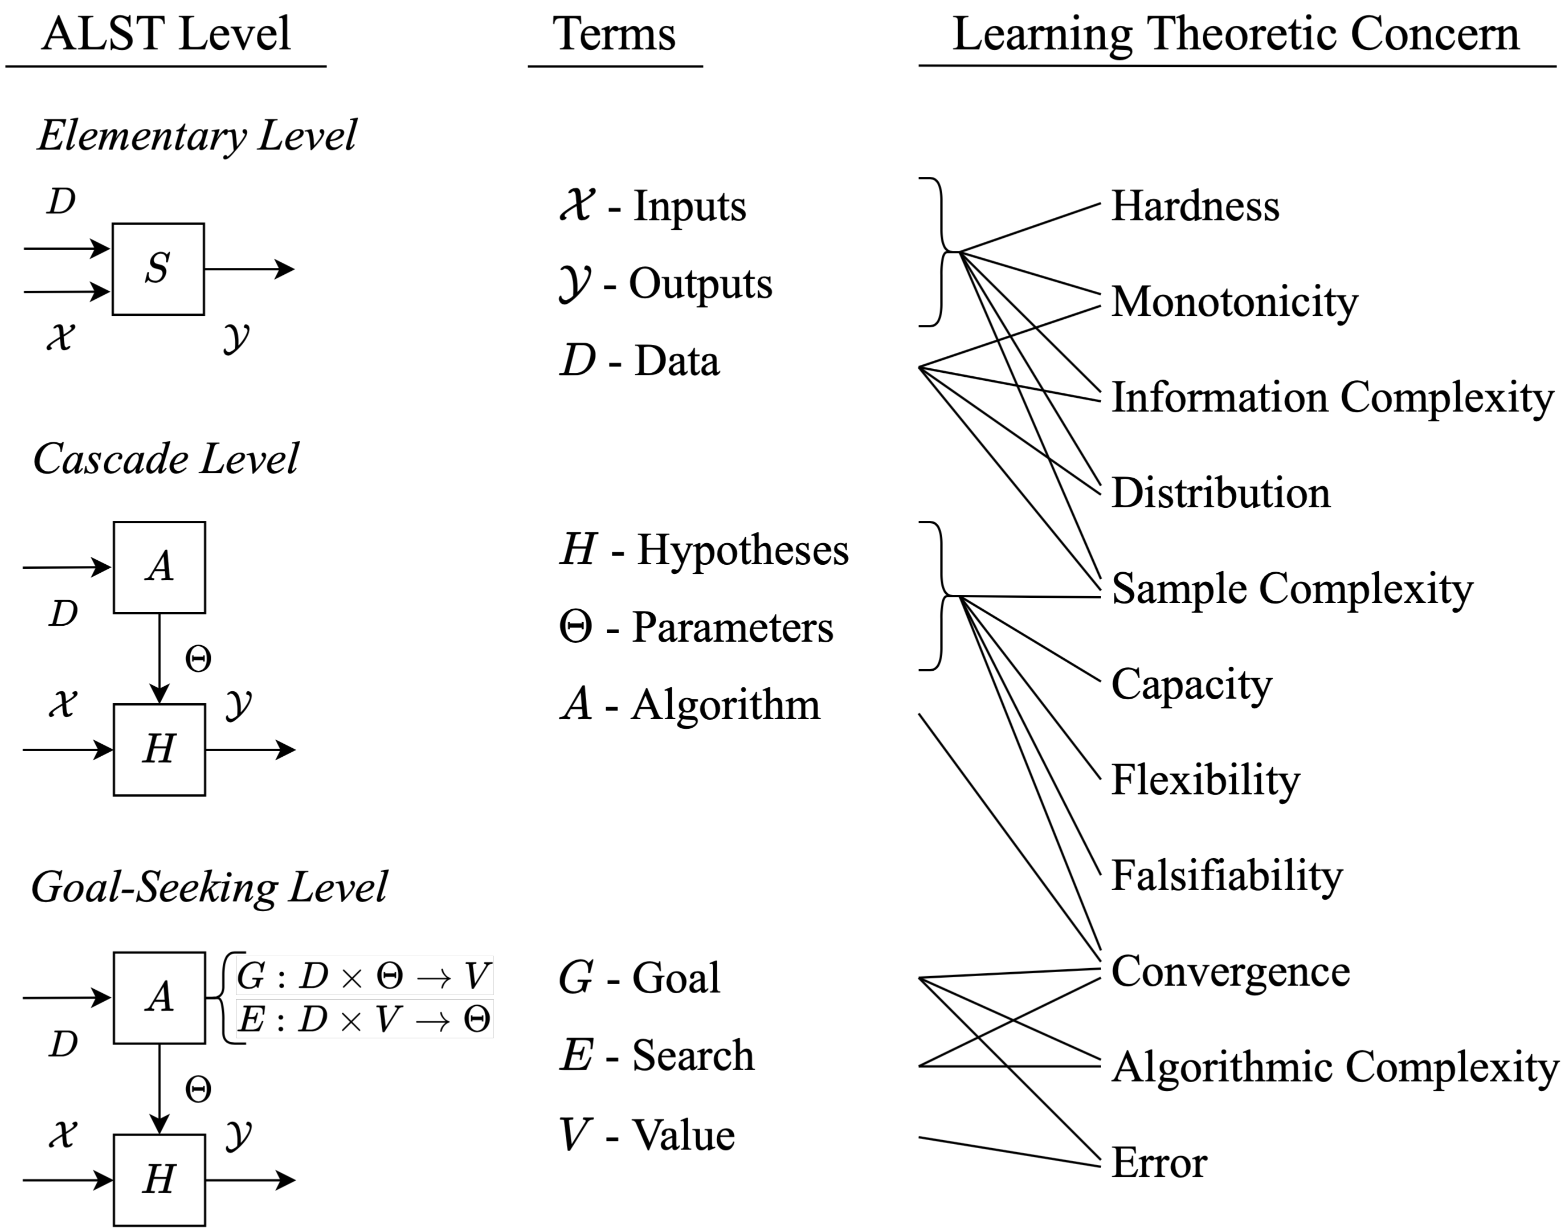

Connecting Learning Components to Their Systems

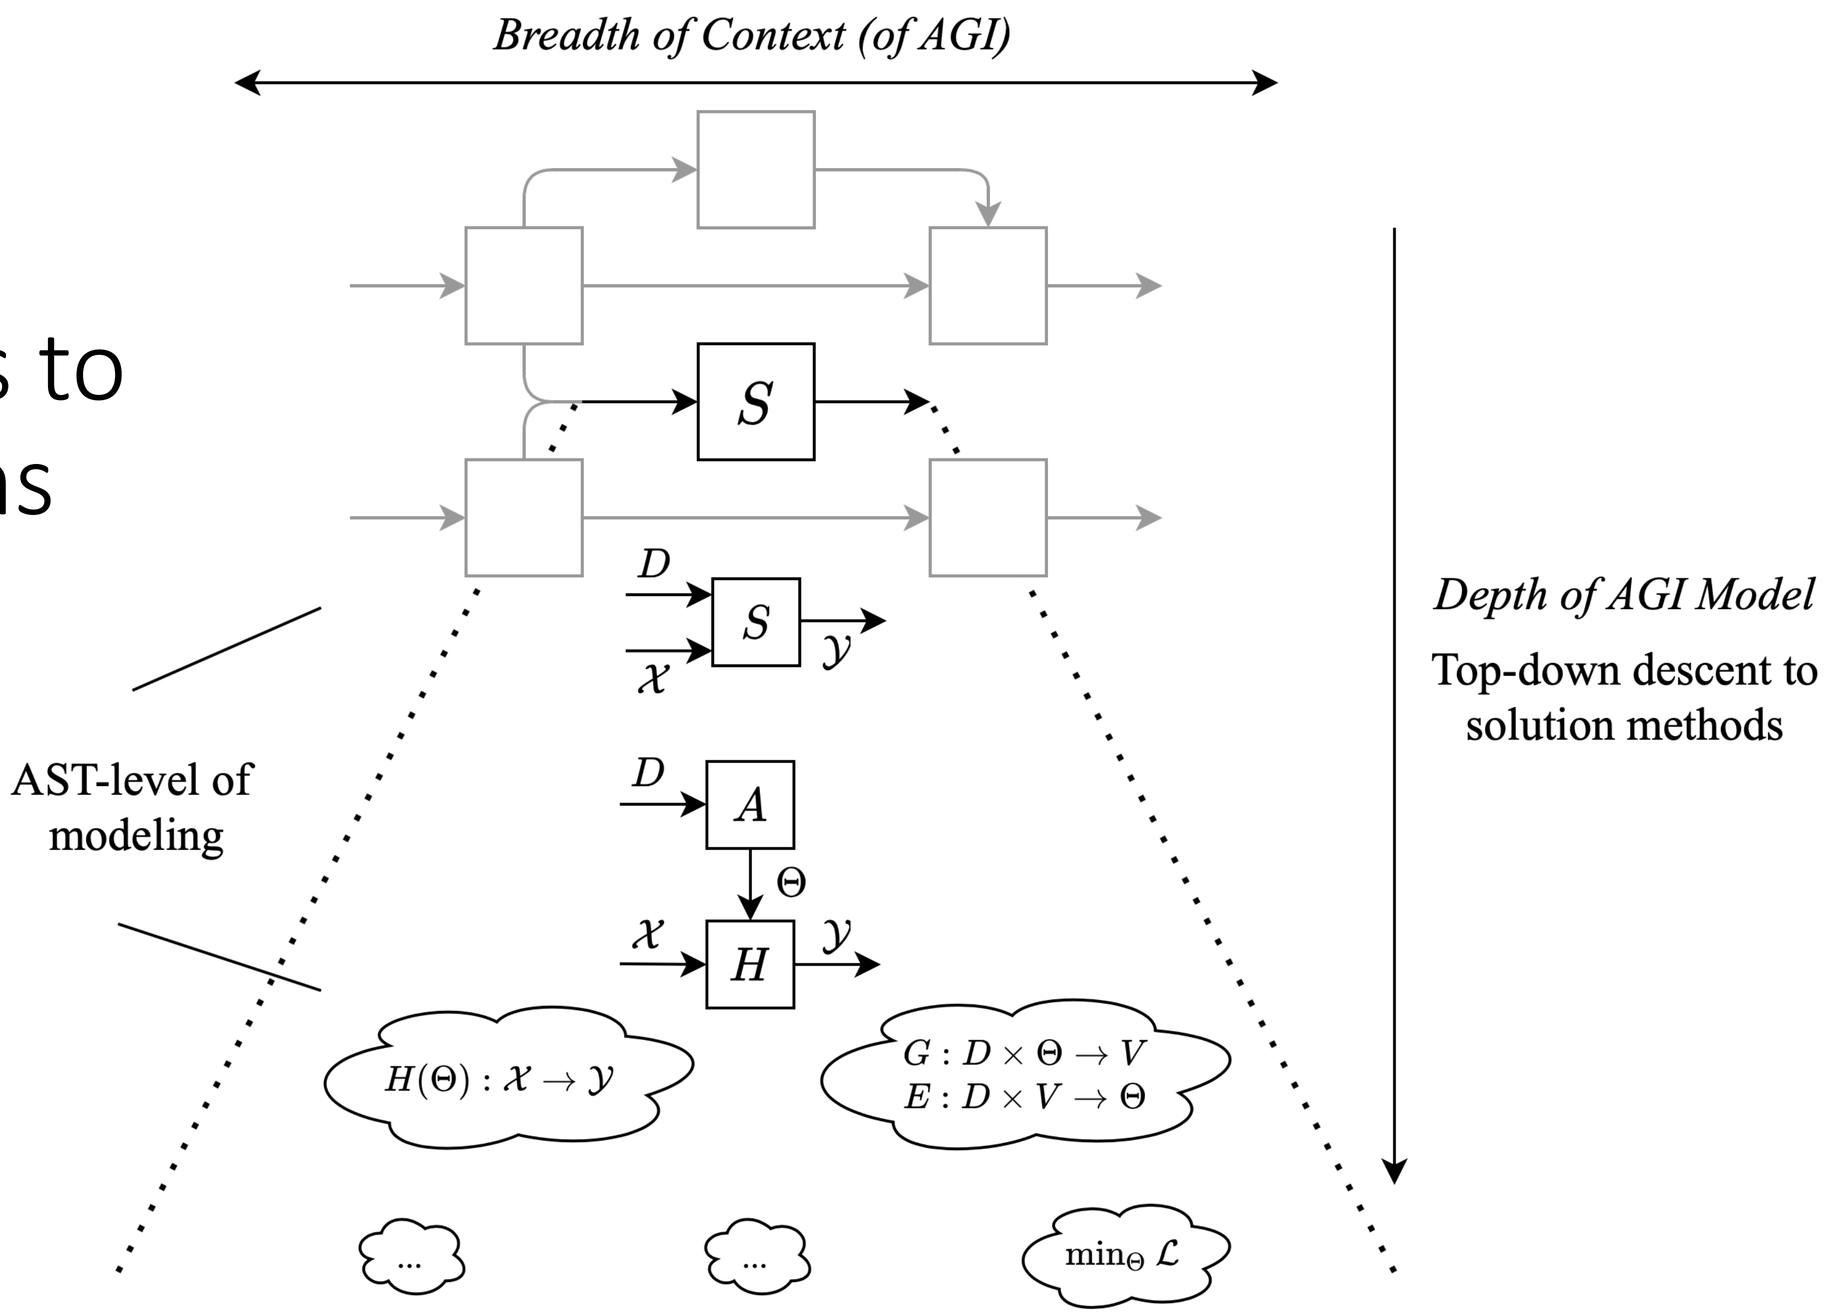

He et. al’s Method

He, Lei, Nabil Aouf, and Bifeng Song. "Explainable Deep Reinforcement Learning for UAV autonomous path planning." *Aerospace science and technology* 118 (2021): 107052.

- He et al. provide an explainable deep reinforcement learning (RL) method for autonomous path planning
- They contribute
  - (1) a deep neural network (DNN) based reactive controller that can be used for small UAVs with limited computational resources and scenarios requiring rapid reaction to changes in the environment
  - (2) a novel explanation framework to explain the DNN-based controller
- They evaluate in the real-world, thereby providing a case study in explainable deep RL for UAV path planning with real-world experiments

He et. al's Formulation

- Path planning in an unknown environment is treated as a sequential decision making problem
- At each time step, only the current sensor information is used to generate the control signal (Markov assumption)
- Assume UAV has a 3D departure position and target destination
- The state at each time step has both the raw depth image from the UAV's camera as well as UAV state features for current velocities and relative position
- To get a smooth control command, they use Twin Delayed DDPG (TD3)
- They use a combination of AirSim and Gazebo for training
- They use Shapely values to extend existing explanation methods CAM and Grad-CAM

Elementary Level

- In RL, we know that data $\mathcal{D}$ consists of states, actions, and rewards
- Inputs $\mathcal{X}$ are states S and outputs $\mathcal{Y}$ are actions A

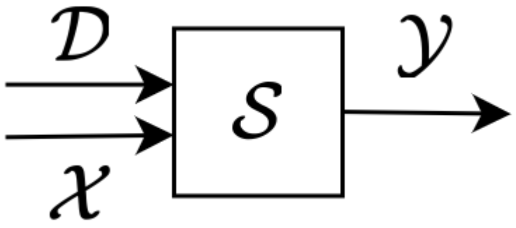

$$\mathcal{D} \subset \times\{S, A, R\}$$

$$\mathcal{X} = \{\text{depth image, velocities, relative position}\},$$

$$\mathcal{Y} = \{\text{forward speed, climb speed, steering speed}\}$$

Elementary Level

Cascade Level

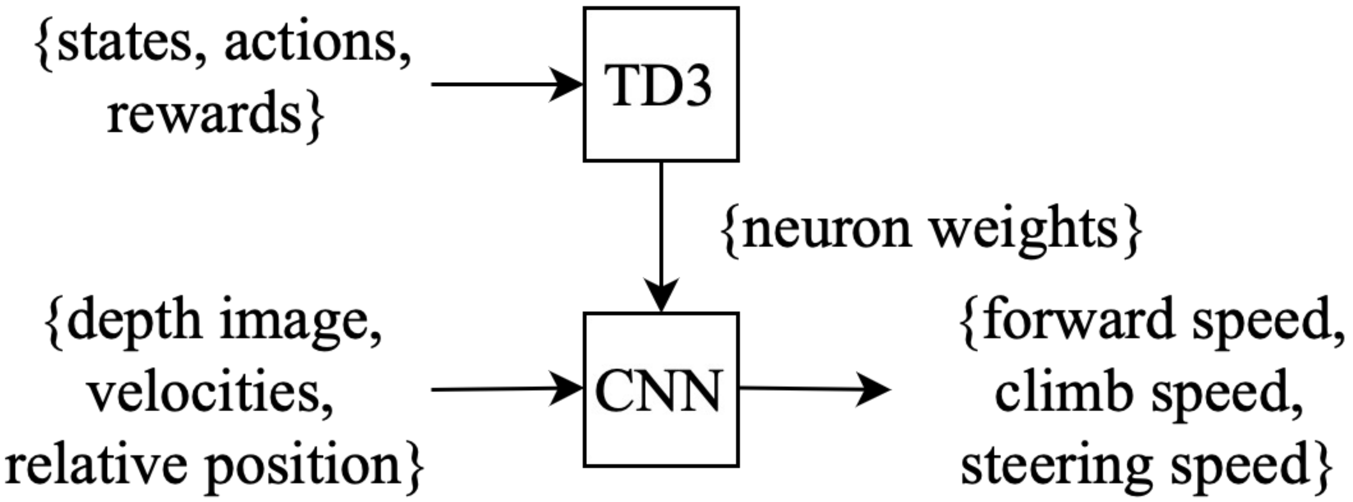

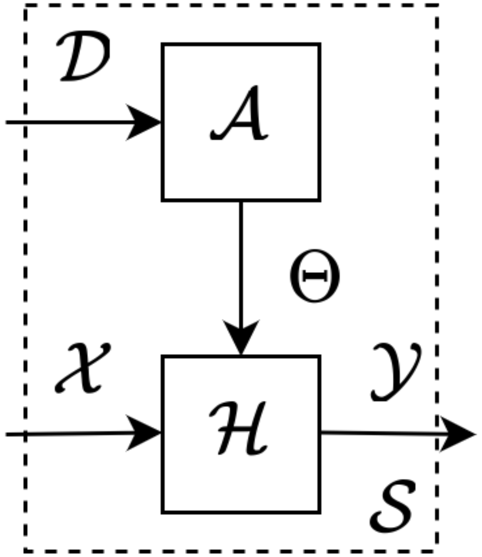

Cascade Level

Goal-Seeking Level

- In RL, the Q-value is the maximum expected future reward for a given state and action
- For TD3, $\mathcal{E}$ is simply a maximization of the expected Q-values over Θ
- $\mathcal{G}$ is more complex, and includes clipped double-Q learning, delayed policy updates, and target policy smoothing
- Adding detail to $\mathcal{G}$ and $\mathcal{E}$ bring the model near the level of the RL solution method

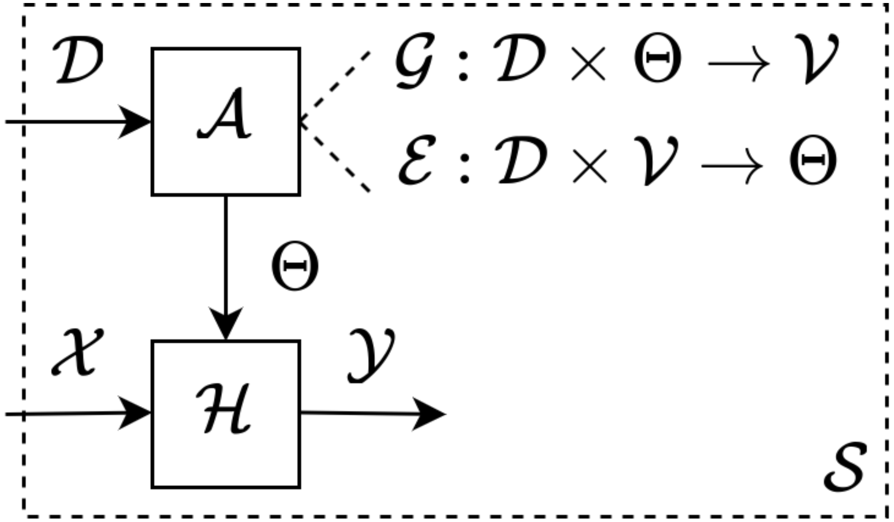

Goal-Seeking Level

He et. al's Explainability Method

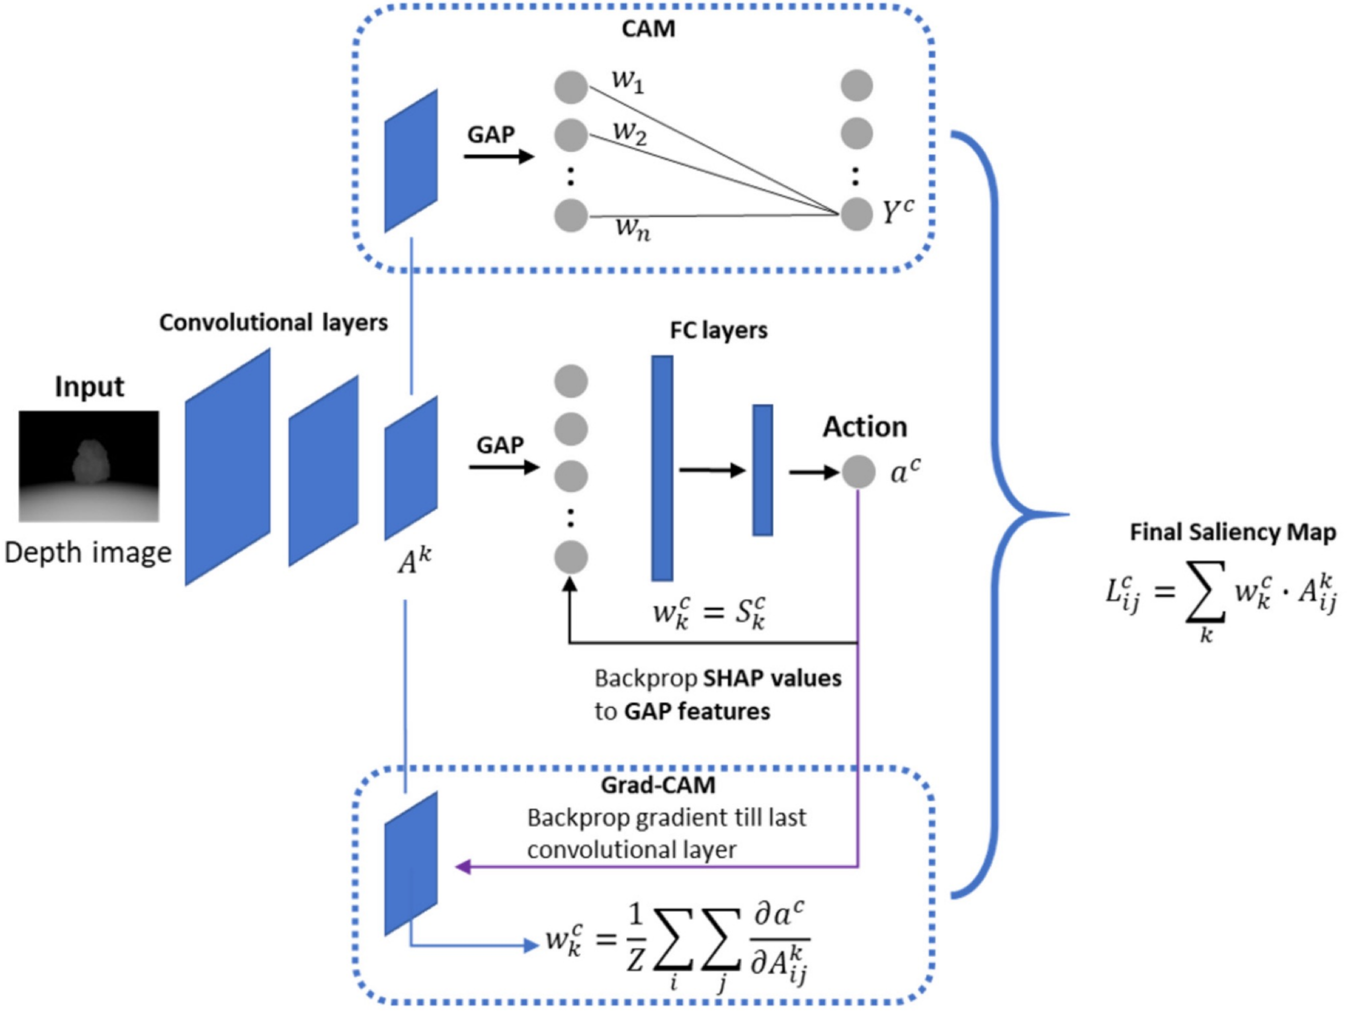

Auxiliary (Explainability) Functions

$$M_{\text{SHAP}} : \mathcal{H}_{\text{FCN}} \times \mathcal{X} \rightarrow \text{Shapely values}$$
$$M_{\text{CAM}} : \mathcal{H}_{\text{CNN}} \times \mathcal{X} \times \text{Shapely values} \rightarrow \text{saliency map}$$
$$M_{\text{Text}} : \text{Shapely values} \rightarrow \text{textual description}$$

Acknowledgements

This material is based upon work supported, in whole or in part, by the U.S. Department of Defense through the Office of the Assistant Secretary of Defense for Research and Engineering (ASD(R&E)) under Contract HQ003419D0003. The Systems Engineering Research Center (SERC) is a federally funded University Affiliated Research Center managed by Stevens Institute of Technology. Any views, opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense nor ASD(R&E).

Contact

Tyler Cody, Ph.D.

Research Assistant Professor

Intelligent Systems Division, Virginia Tech National Security Institute

tcody@vt.edu

LinkedIn, ResearchGate as “Tyler Cody”