# Logistics

- **Assignment #4** is due **April 11** (this Thursday) at 11:59pm
    - Late submissions for 20% deduction until **April 15** at 11:59pm
- **SPOT** (formerly USRI) surveys are now available
    - Available until **April 14** at 11:59pm
    - You should have gotten an email
    - ***Please do fill one out***, even if you weren't here for today's lecture

# Goal Recognition Design

William Yeoh

Computer Science and Engineering
Washington University in St. Louis

# Content

- Goal Recognition
- Goal Recognition Design
- Stochastic Goal Recognition Design
- Partially-Observable Stochastic Goal Recognition Design
- Ongoing Work: Data-driven Goal Recognition Design

# Goal Recognition

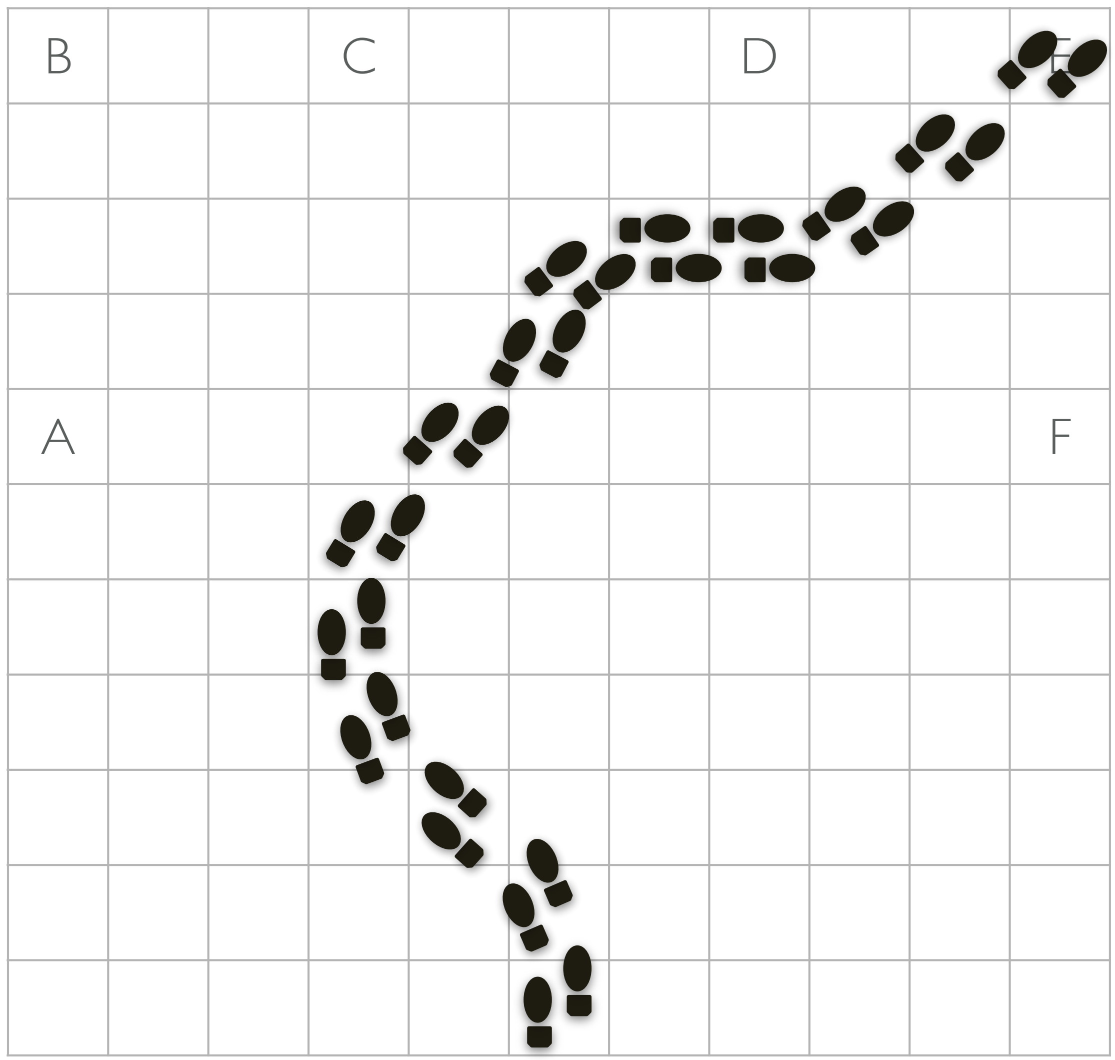

# Goal Recognition

- **A Goal Recognition Model:**
  - Trying to identify the goal  of an agent based on its observations :

$$P(G \mid O) = \alpha P(O \mid G) P(G)$$

  - $P(G)$ is the probability that $G$ is the true goal (assumed to be given)
  - $P(O \mid G)$ is the probability that we observe $O$ given than $G$ is the true goal
    - Based on the cost of the trajectory observed so far
    - The closer its cost to the optimal cost, the larger the probability

*Ramírez and Geffner: Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. AAAI 2010*

# Goal Recognition

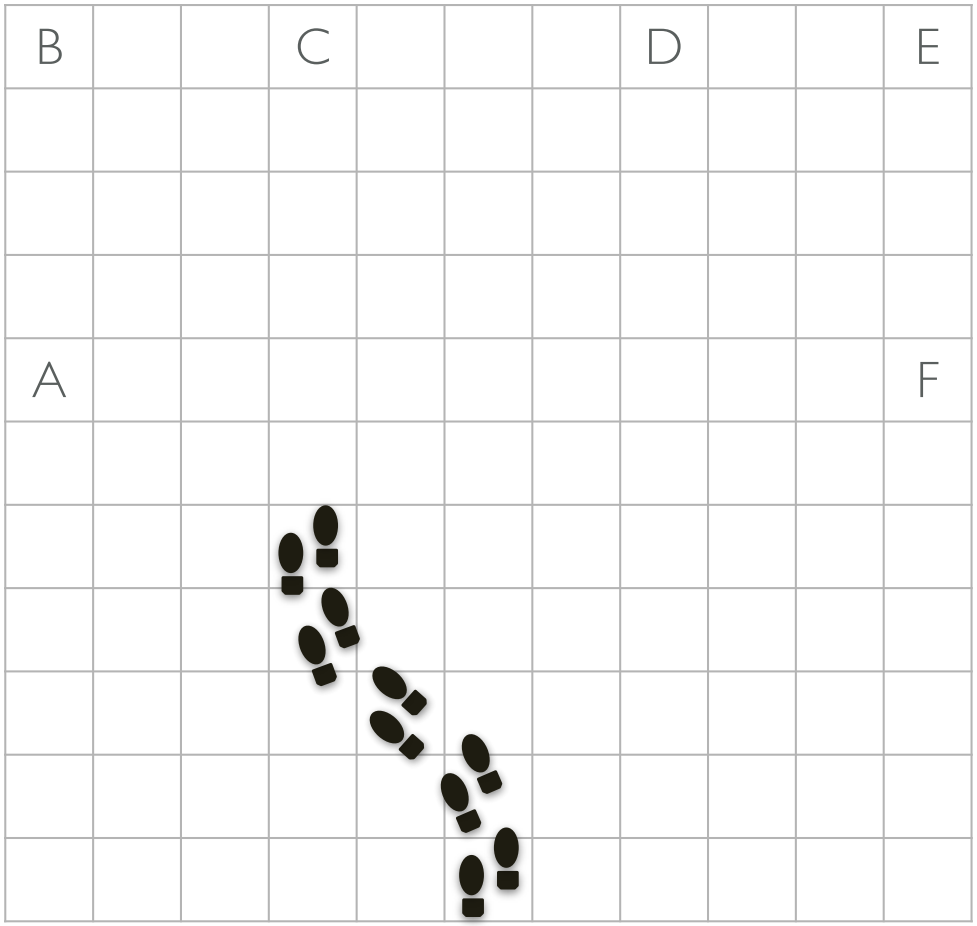


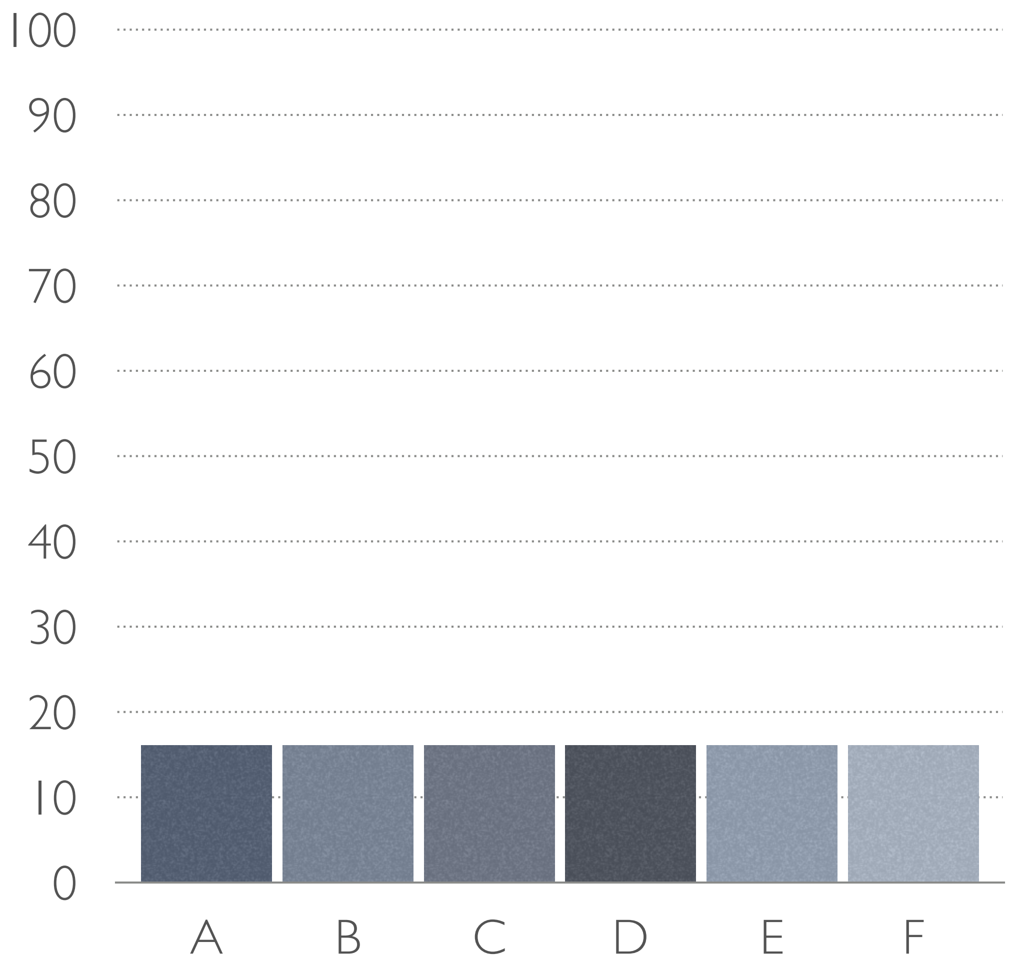




YODAlab

# Goal Recognition

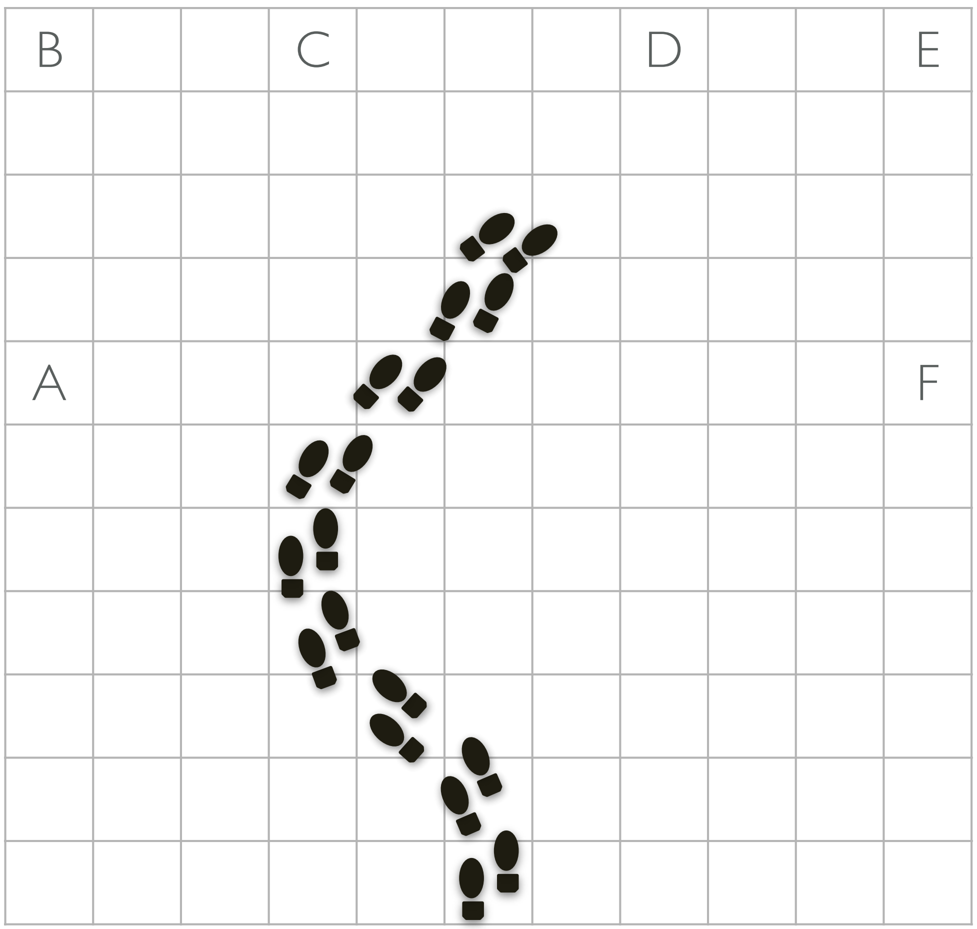


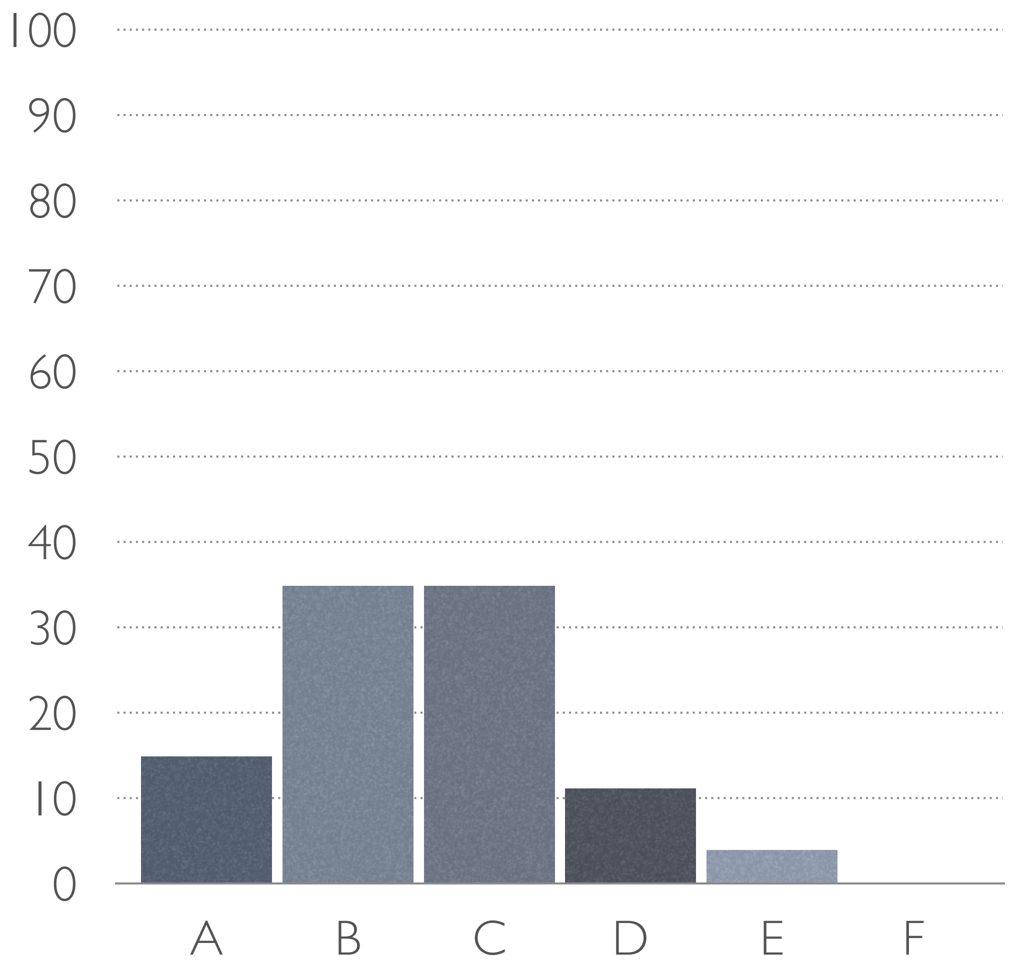




YODAlab

# Goal Recognition

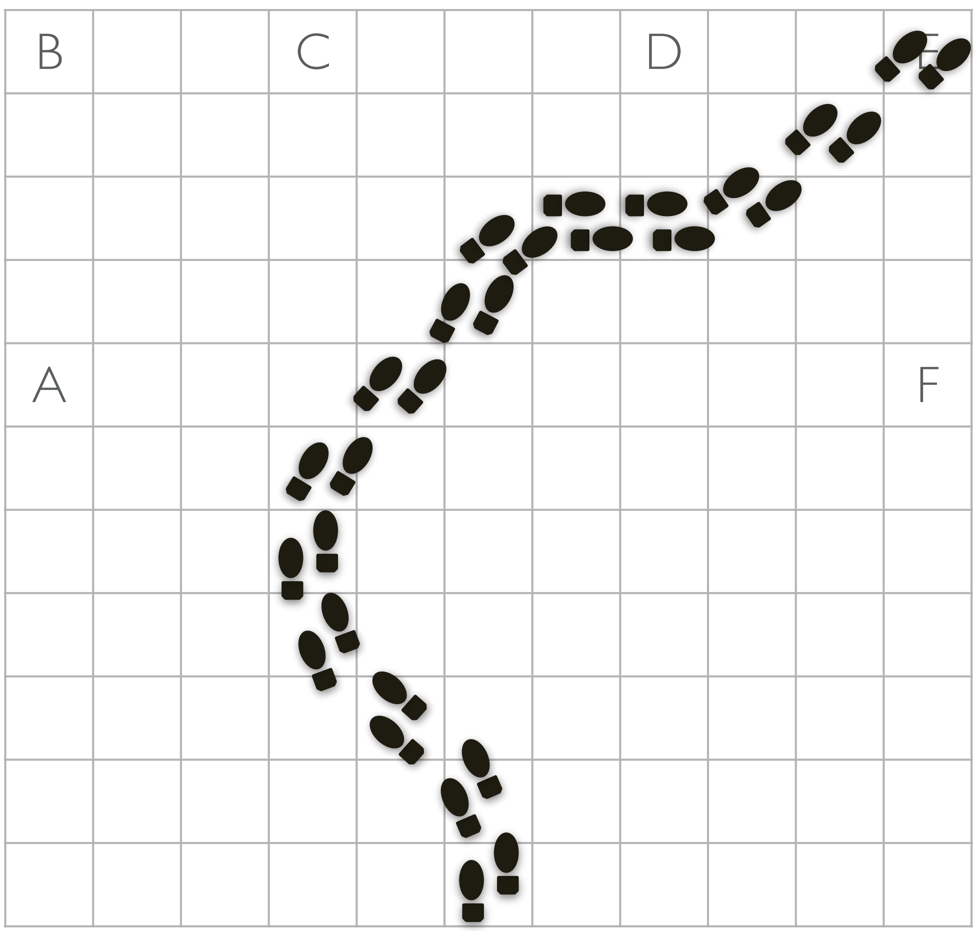


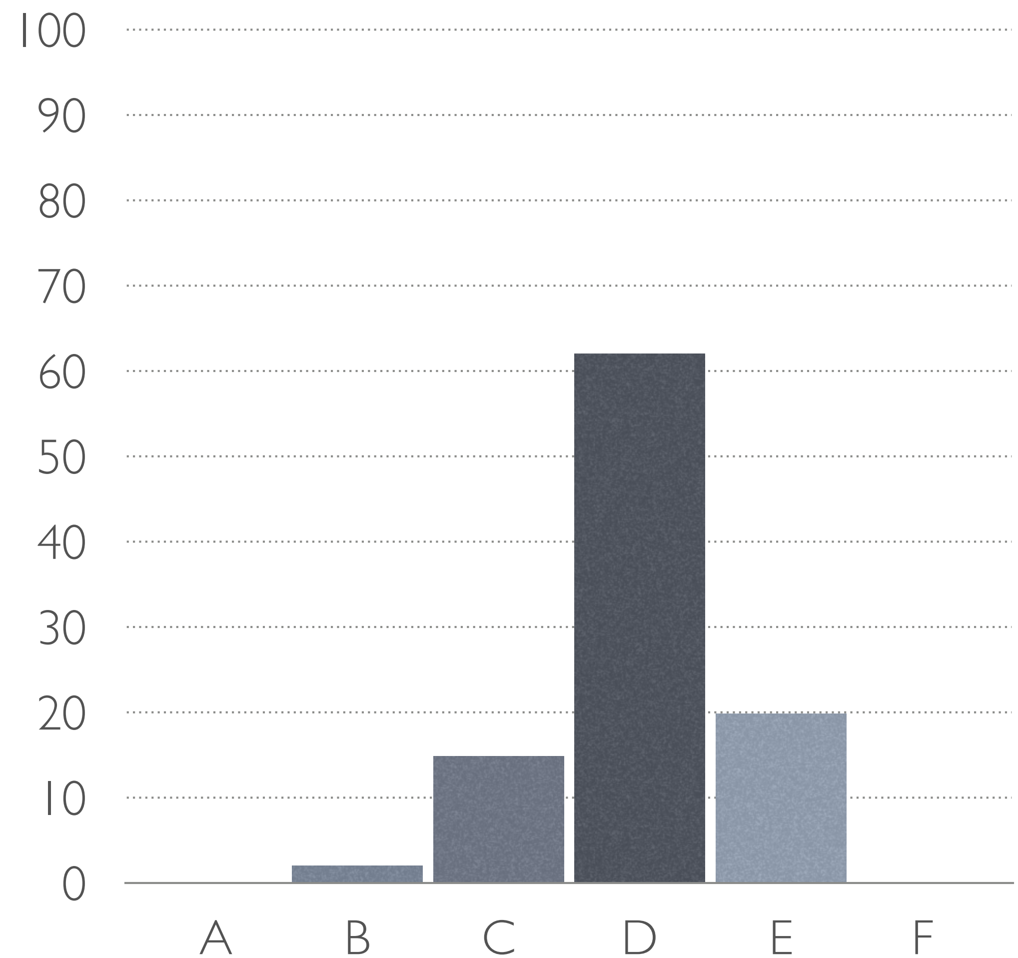




YODAlab

# Goal Recognition

- Applications:
  - Human-robot interactions [Tvakkoli et al., 2007; Kelley et al., 2012]

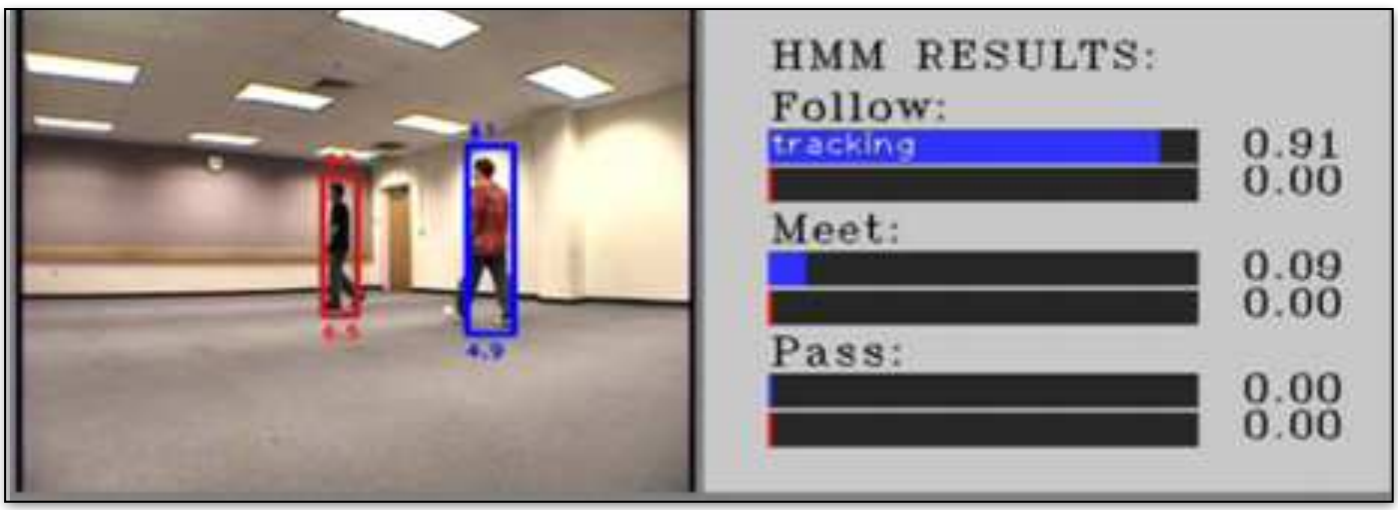


*Tavakkoli, Kelley, King, Nicolescu, Nicolescu, and Bebis. A vision-based architecture for intent recognition. International Symposium on Advances in Visual Computing 2007*

# Goal Recognition

- Applications:
  - Human-robot interactions [Tvakkoli et al., 2007; Kelley et al., 2012]

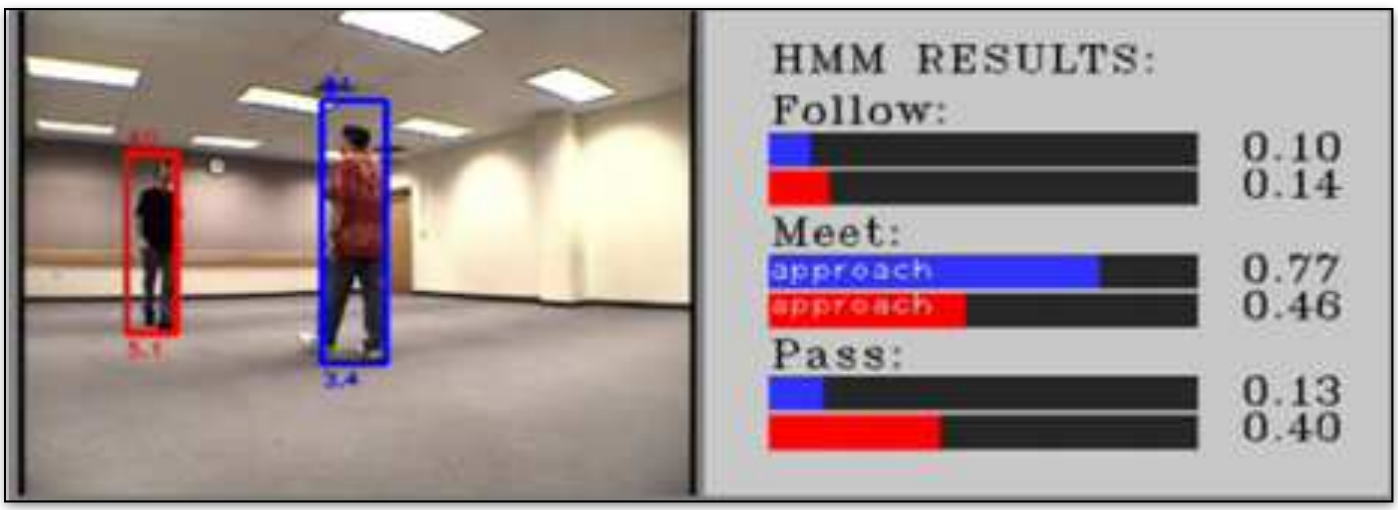


*Tavakkoli, Kelley, King, Nicolescu, Nicolescu, and Bebis. A vision-based architecture for intent recognition. International Symposium on Advances in Visual Computing 2007*

# Goal Recognition

- Applications:
  - Human-robot interactions [Tvakkoli et al., 2007; Kelley et al., 2012]

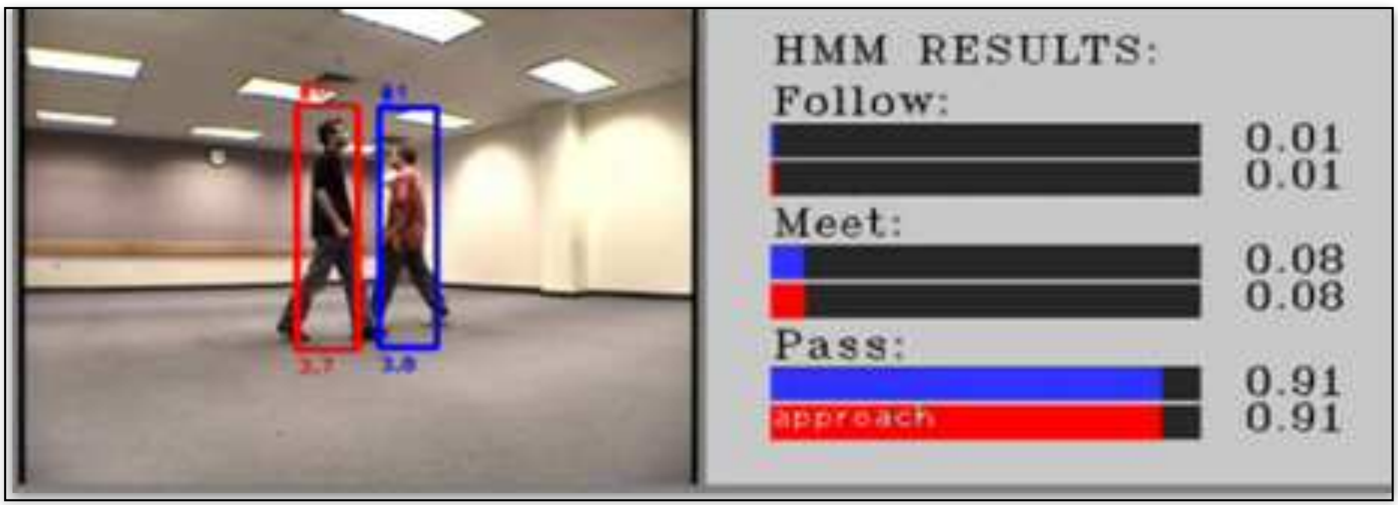


*Tavakkoli, Kelley, King, Nicolescu, Nicolescu, and Bebis. A vision-based architecture for intent recognition. International Symposium on Advances in Visual Computing 2007*

# Goal Recognition

- Applications:
  - Human-robot interactions [Tvakkoli et al., 2007; Kelley et al., 2012]
  - Software personal assistants [Oh et al., 2010, 2011]

*Source: http://assets.pewresearch.org/*



YODAlab

# Goal Recognition

- Applications:
    - Human-robot interactions [Tvakkoli et al., 2007; Kelley et al., 2012]
    - Software personal assistants [Oh et al., 2010, 2011]
    - Intelligent tutoring systems [McQuiggan et al., 2008; Johnson, 2010; Min et al., 2014]

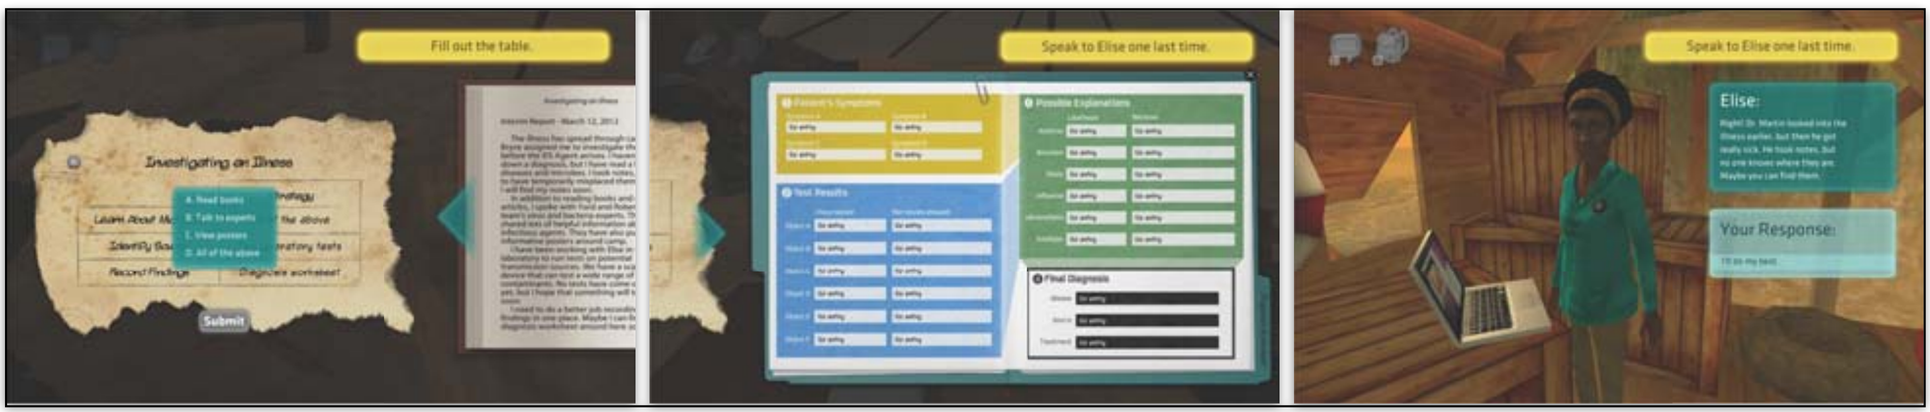


Source: http://projects.intellimedia.ncsu.edu/crystalisland/

# Goal Recognition

- Applications:
  - Human-robot interactions [Tvakkoli et al., 2007; Kelley et al., 2012]
  - Software personal assistants [Oh et al., 2010, 2011]
  - Intelligent tutoring systems [McQuiggan et al., 2008; Johnson, 2010; Min et al., 2014]
  - Security applications [Jarvis et al., 2005]

Source: http://www.netralnews.com/

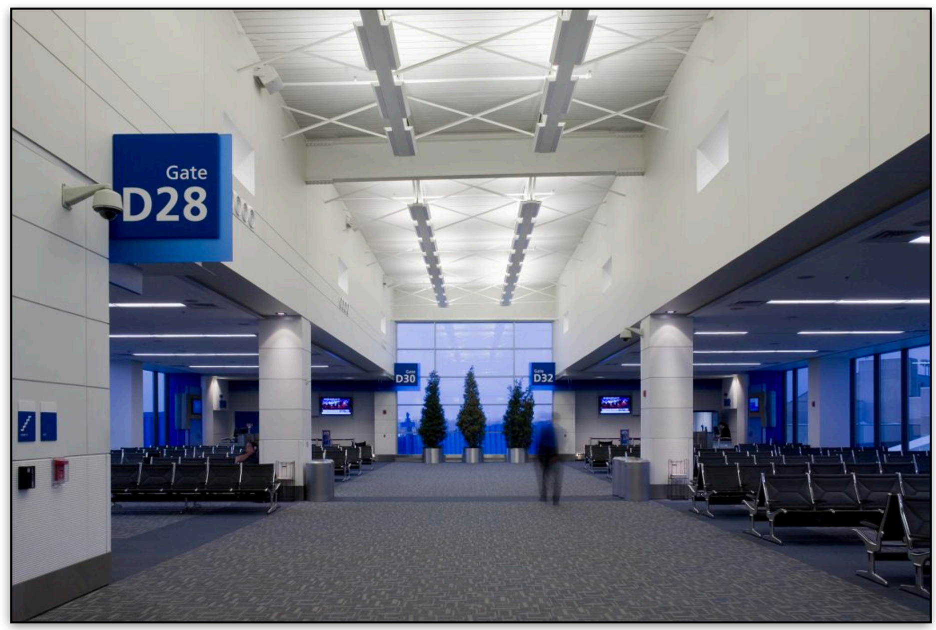


Source: http://www.walbridge.com/

# Content

- Goal Recognition
- Goal Recognition Design
- Stochastic Goal Recognition Design
- Partially-Observable Stochastic Goal Recognition Design
- Ongoing Work: Data-driven Goal Recognition Design



YODAlab

# Goal Recognition

- **Goal Recognition Design (GRD):**
  - Introduced by Sarah Keren, Avigdor Gal, Erez Karpas at ICAPS 2014
  - How to ***modify/design*** the underlying environment to improve goal recognition?
  - Orthogonal to goal recognition; advances made will complement advances in goal recognition

# Goal Recognition Design

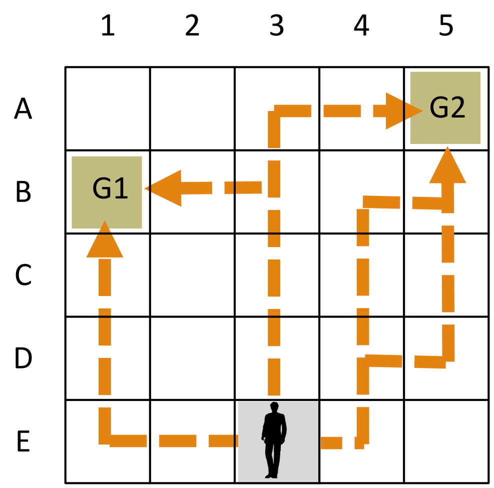


- Assumptions:
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are deterministic

*Keren, Gal, and Karpas: Goal Recognition Design. ICAPS 2014*

# Goal Recognition Design

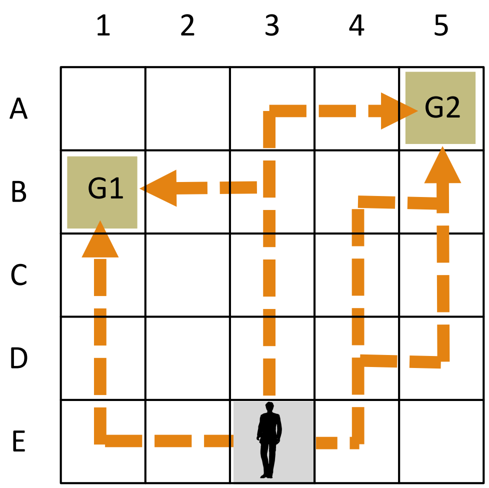


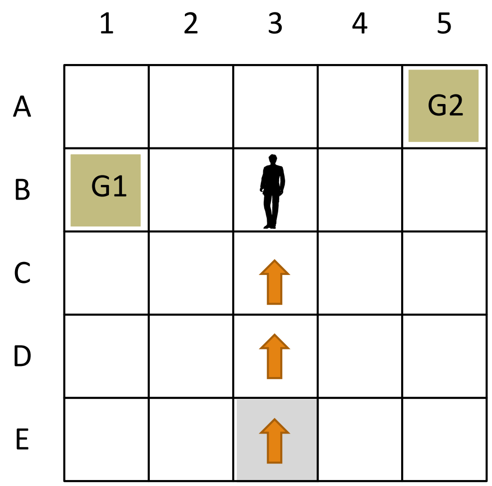


Maximum ambiguous steps: 3

- Assumptions:
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are deterministic

*Keren, Gal, and Karpas: Goal Recognition Design. ICAPS 2014*

# Goal Recognition Design

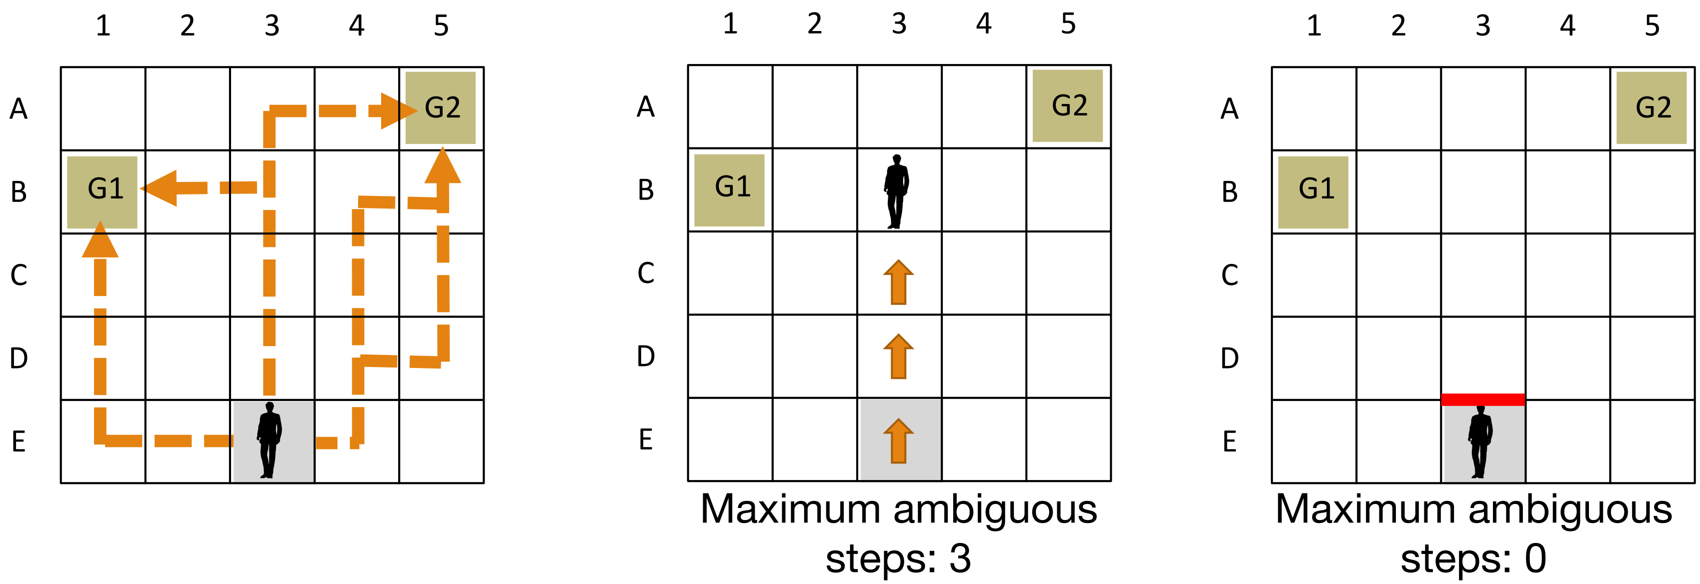


- Assumptions:
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are deterministic

*Keren, Gal, and Karpas: Goal Recognition Design. ICAPS 2014*

# Goal Recognition Design

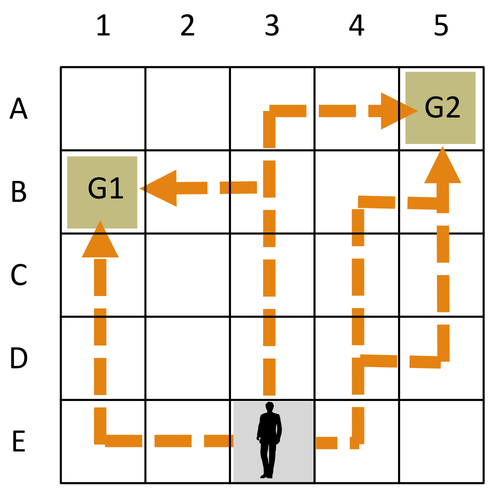


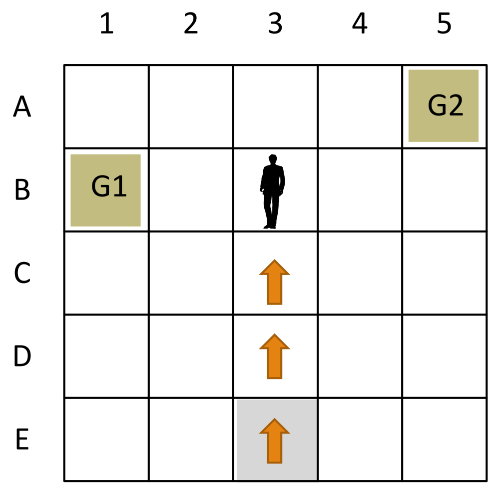


Maximum ambiguous steps: 3

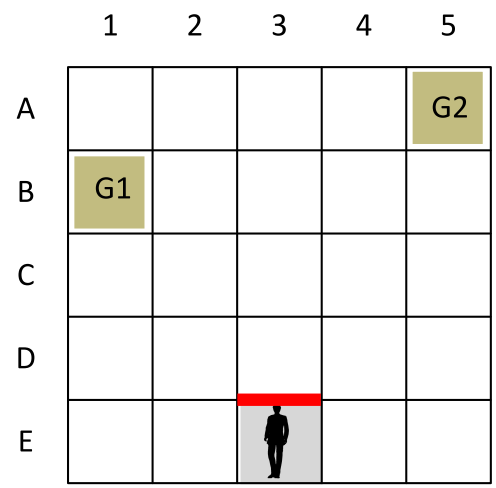


Maximum ambiguous steps: 0

- High-level idea:
  - Assess the difficulty of the problem using a *metric* called worst-case distinctiveness (*wcd*)
  - Find minimal modification to the environment that minimizes *wcd*
    - subject to requirement that optimal cost to each goal remains unchanged

*Keren, Gal, and Karpas: Goal Recognition Design. ICAPS 2014*

# Content

- Goal Recognition
- Goal Recognition Design
- Stochastic Goal Recognition Design
- Partially-Observable Stochastic Goal Recognition Design
- Ongoing Work: Data-driven Goal Recognition Design

# Stochastic GRD

- **Goal Recognition Design (GRD):**
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are **deterministic**
- **Stochastic GRD (S-GRD):**
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are **stochastic**
    - Important in some applications (e.g., robotic, cybersecurity, etc.)

*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*

# Stochastic GRD

- **Goal Recognition Design (GRD):**
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are **deterministic**

- **Stochastic GRD (S-GRD):**
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are **stochastic**
    - Important in some applications (e.g., robotic, cybersecurity, etc.)
    - … *and in some wizarding worlds!!*

*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*

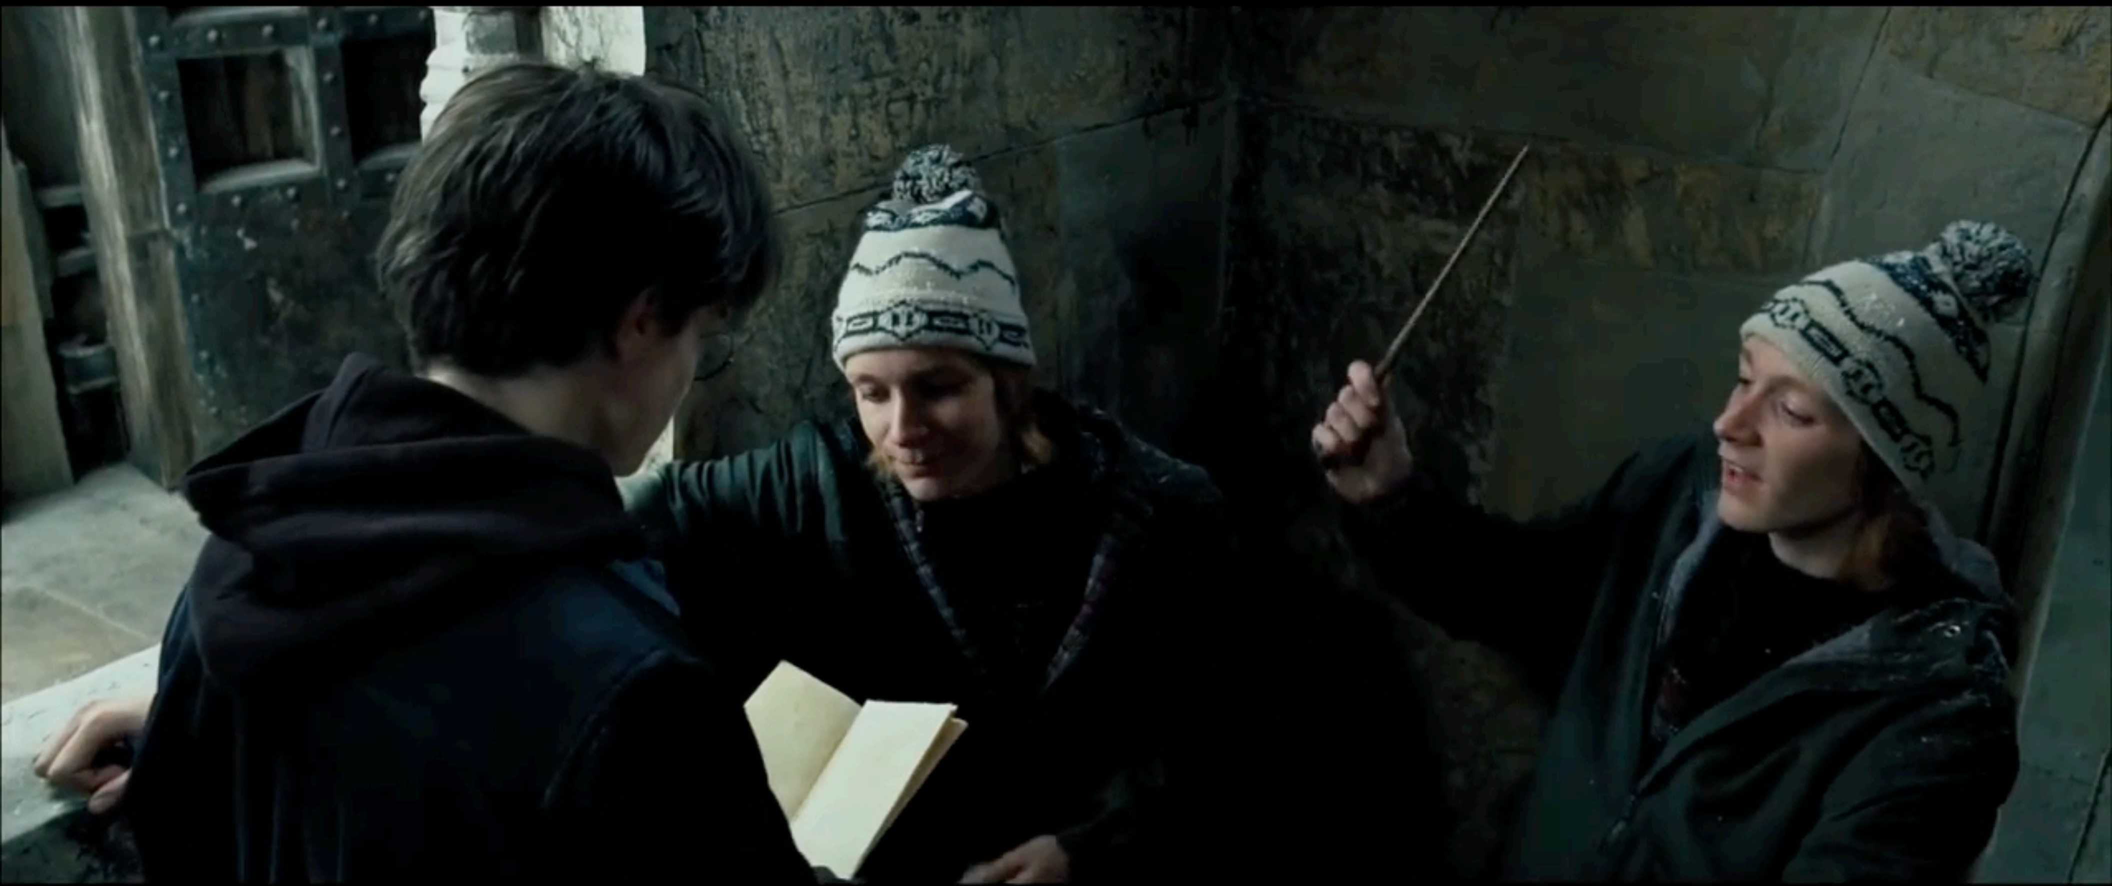

Source: https://www.youtube.com/watch?v=vNc43oKqQzg

Source: https://www.youtube.com/watch?v=uFvizAQHJz8

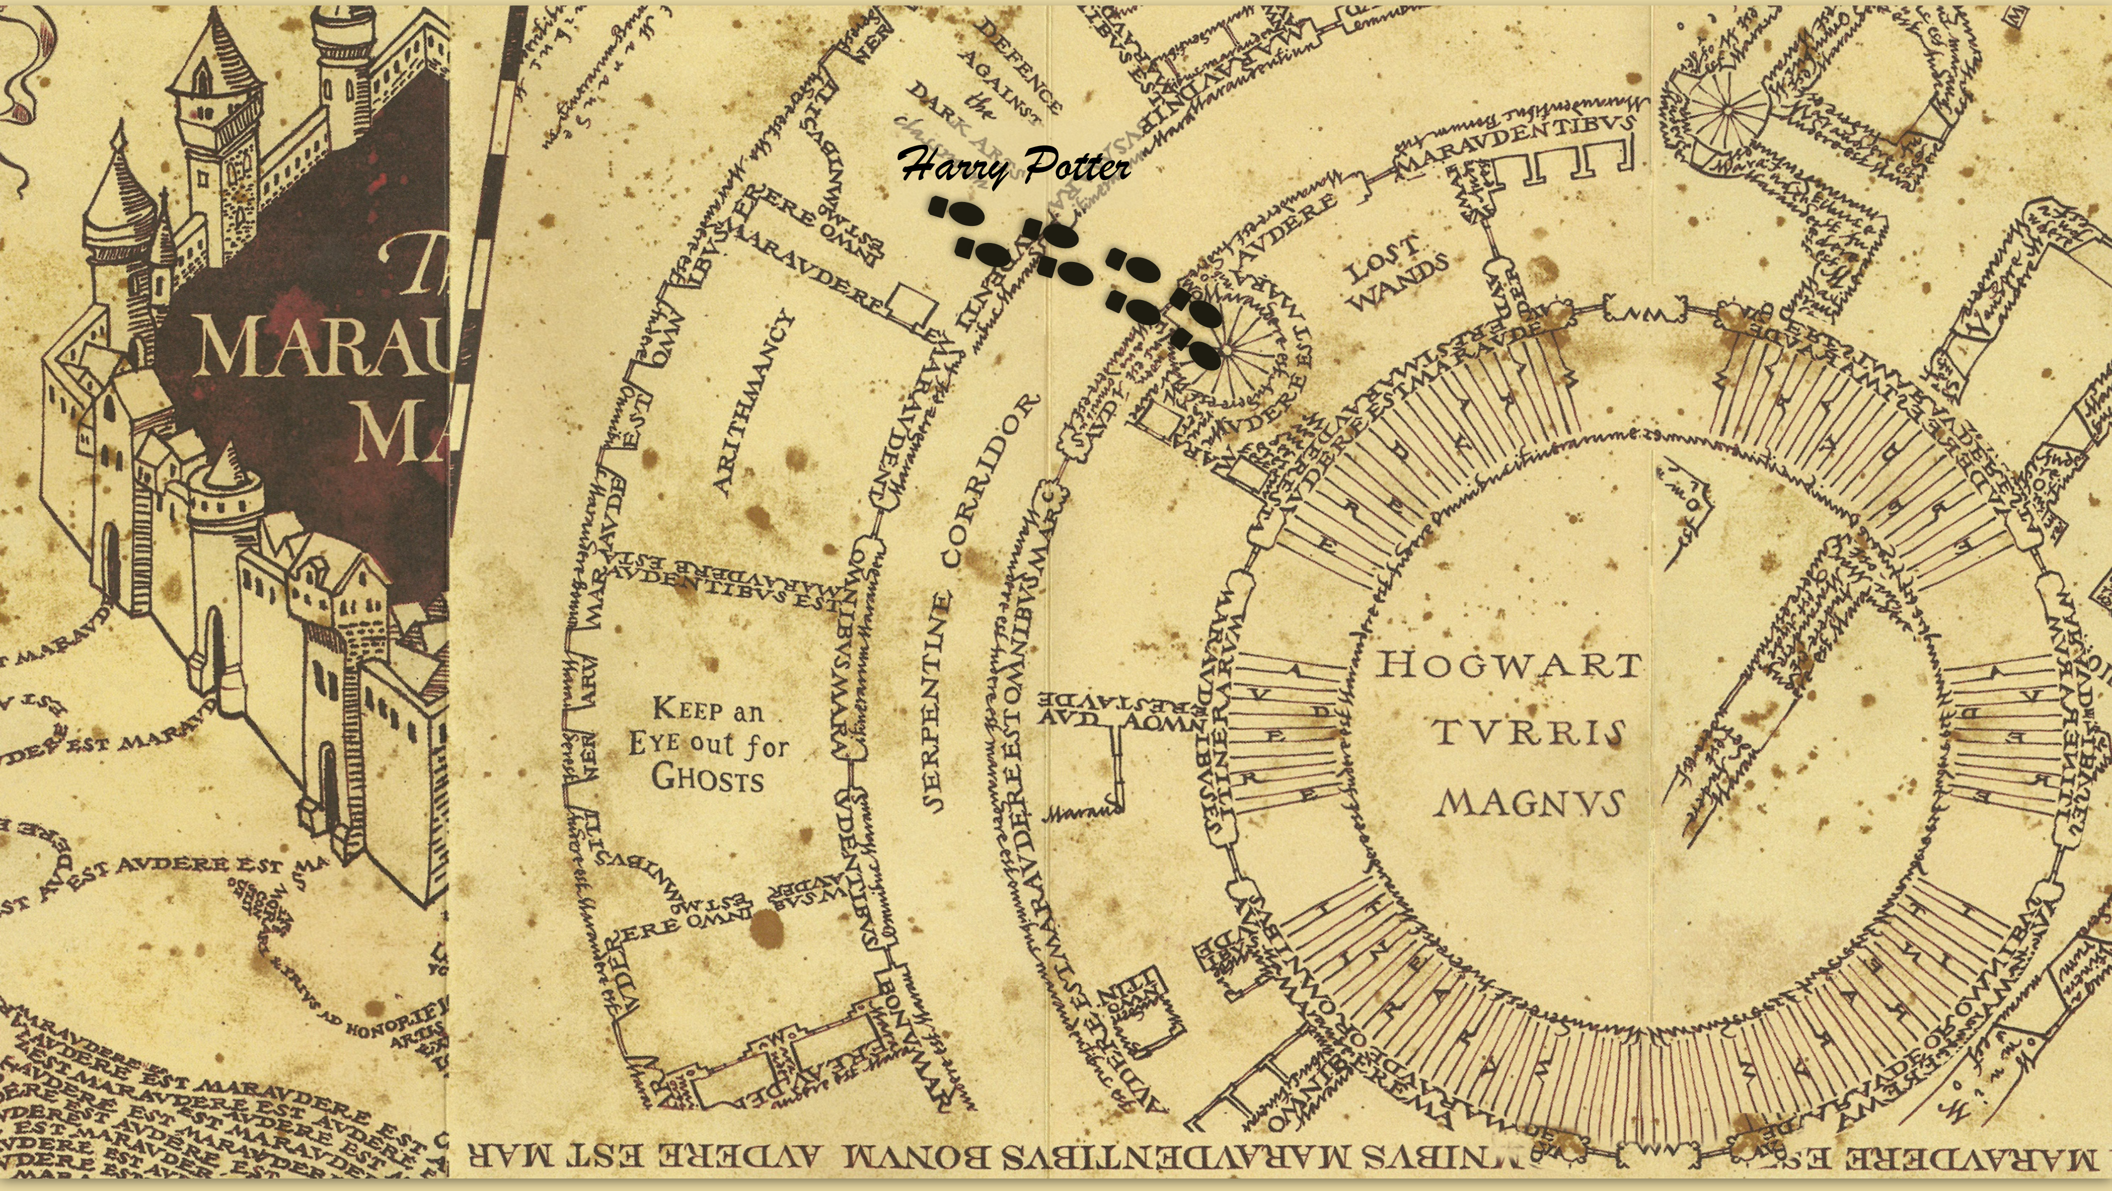


Source: https://littlefallingstar.deviantart.com/art/Marauders-Map-page-1-379091499

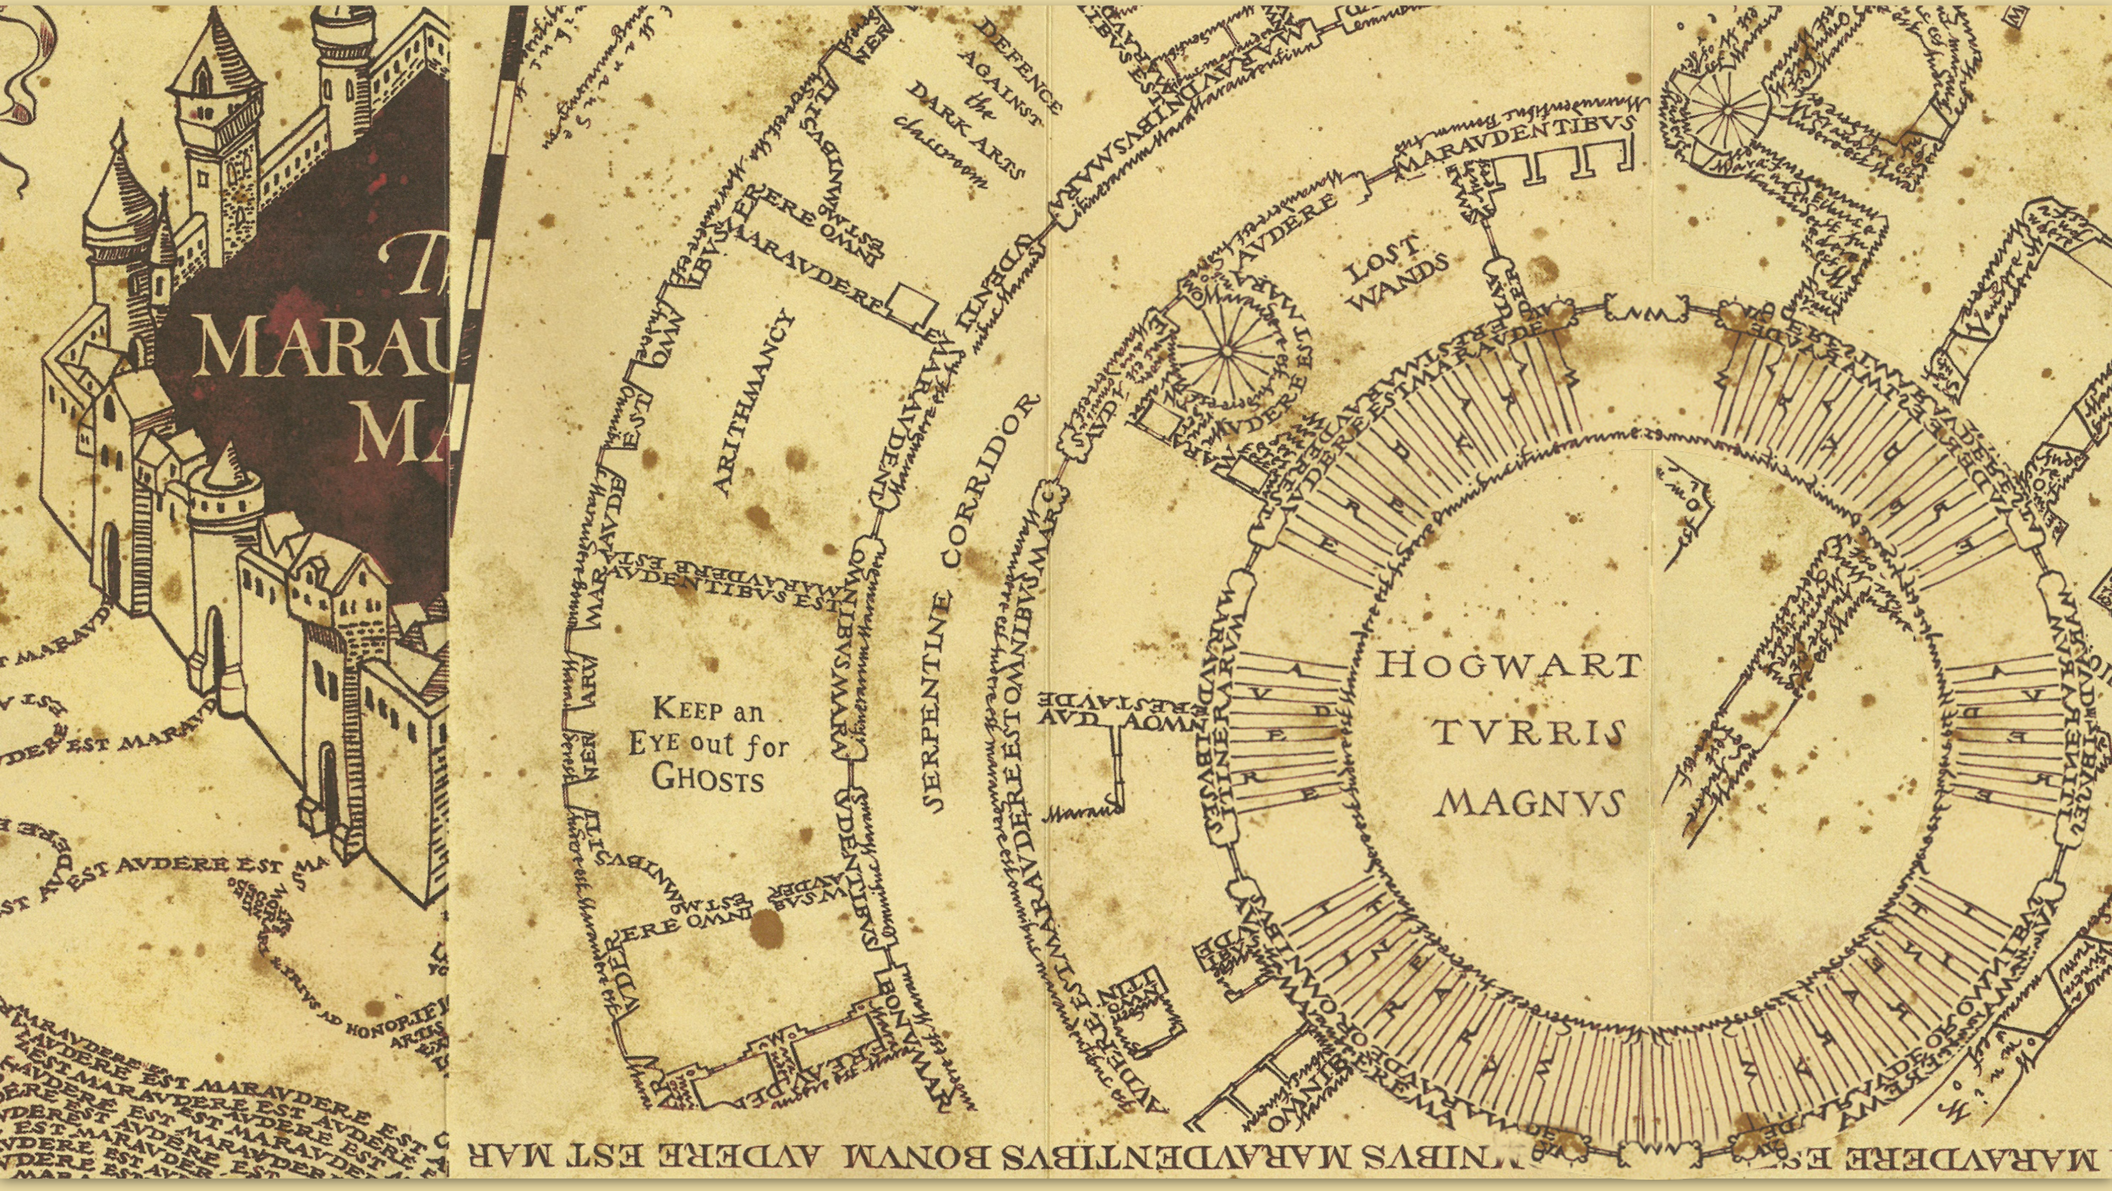


Source: https://littlefallingstar.deviantart.com/art/Marauders-Map-page-1-379091499

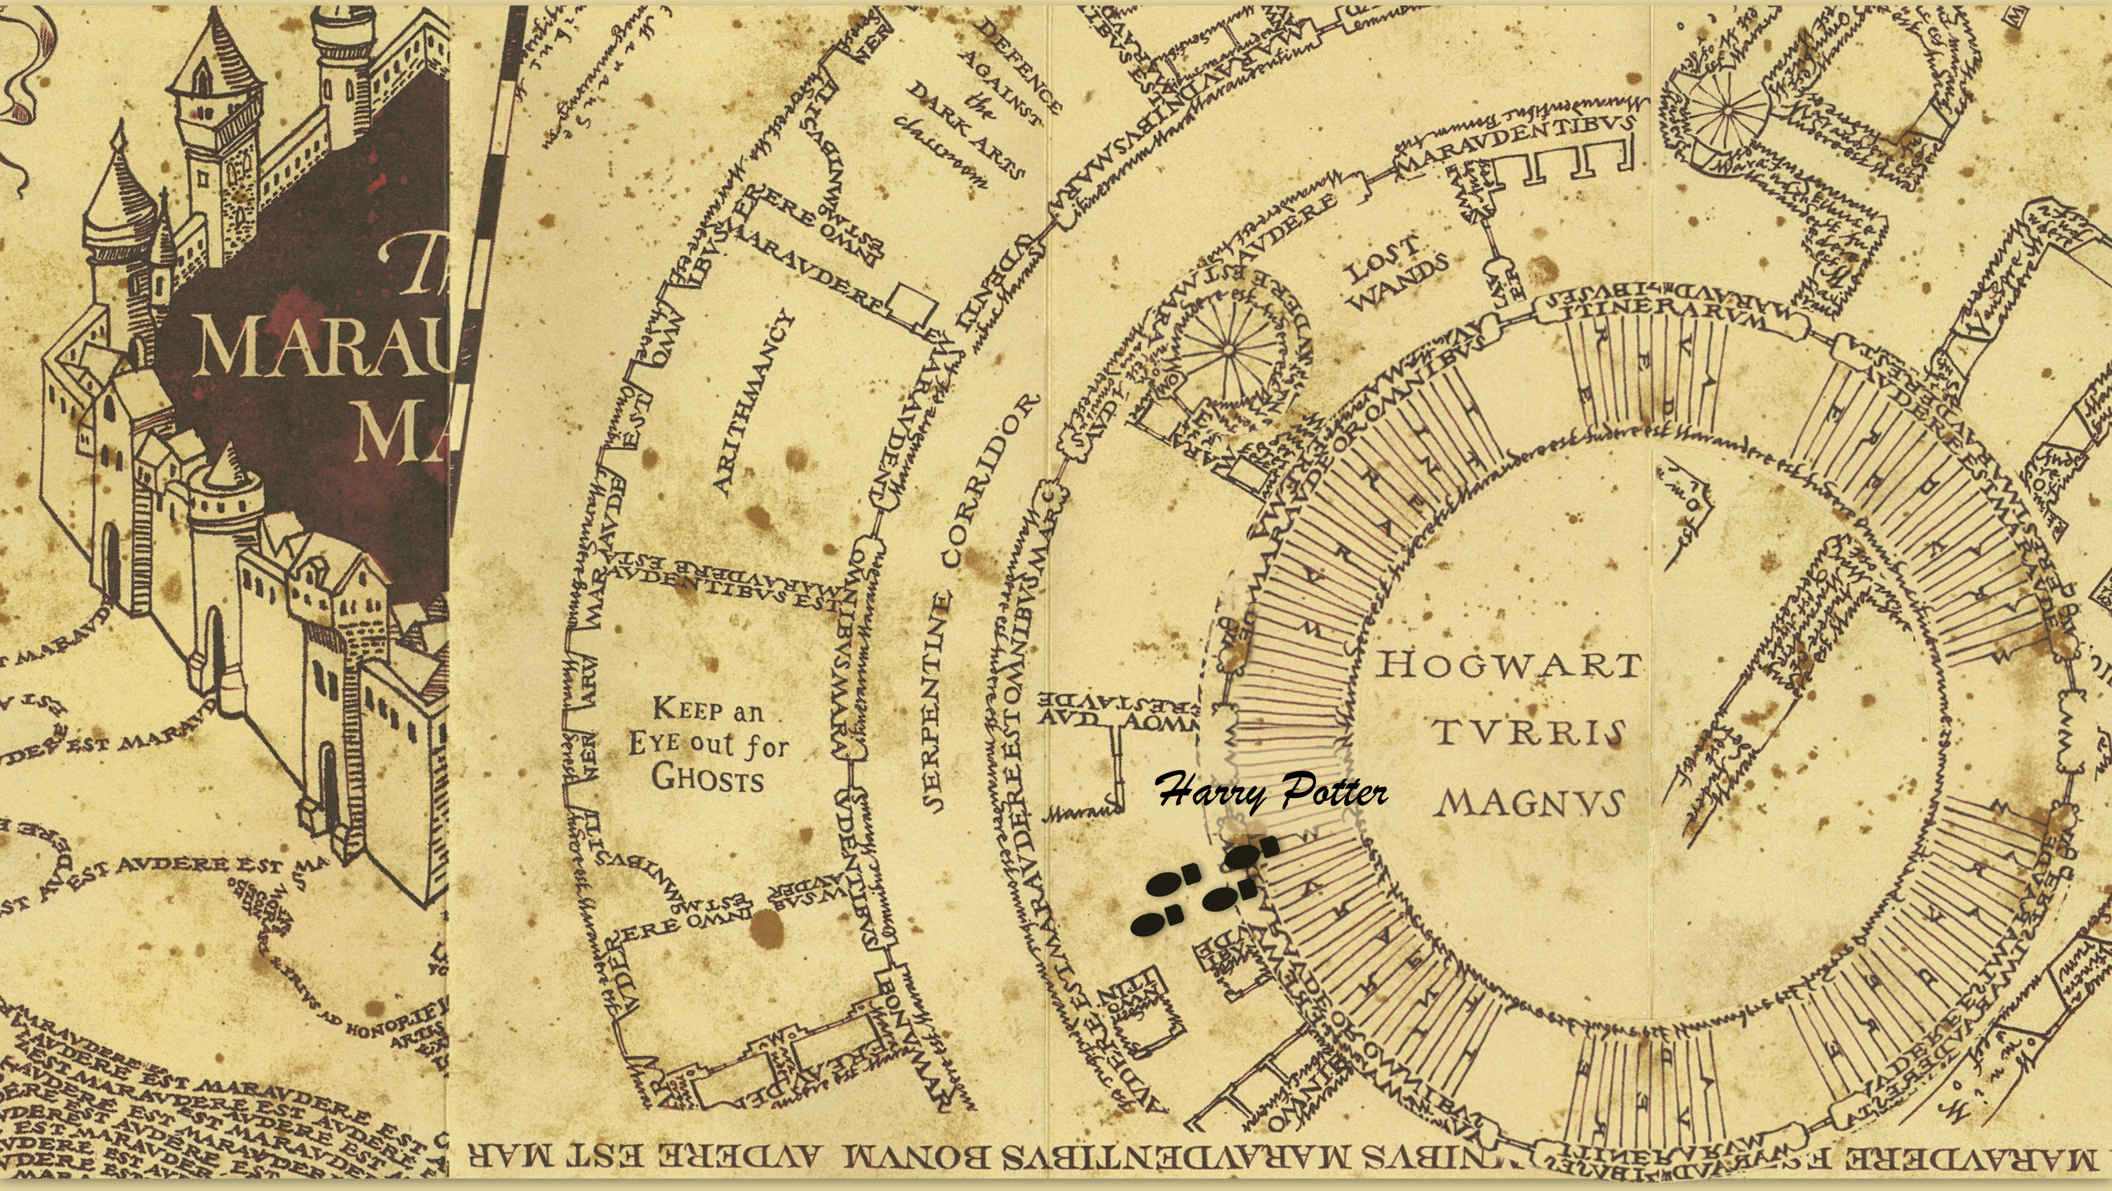


Source: https://littlefallingstar.deviantart.com/art/Marauders-Map-page-1-379091499

# Stochastic GRD

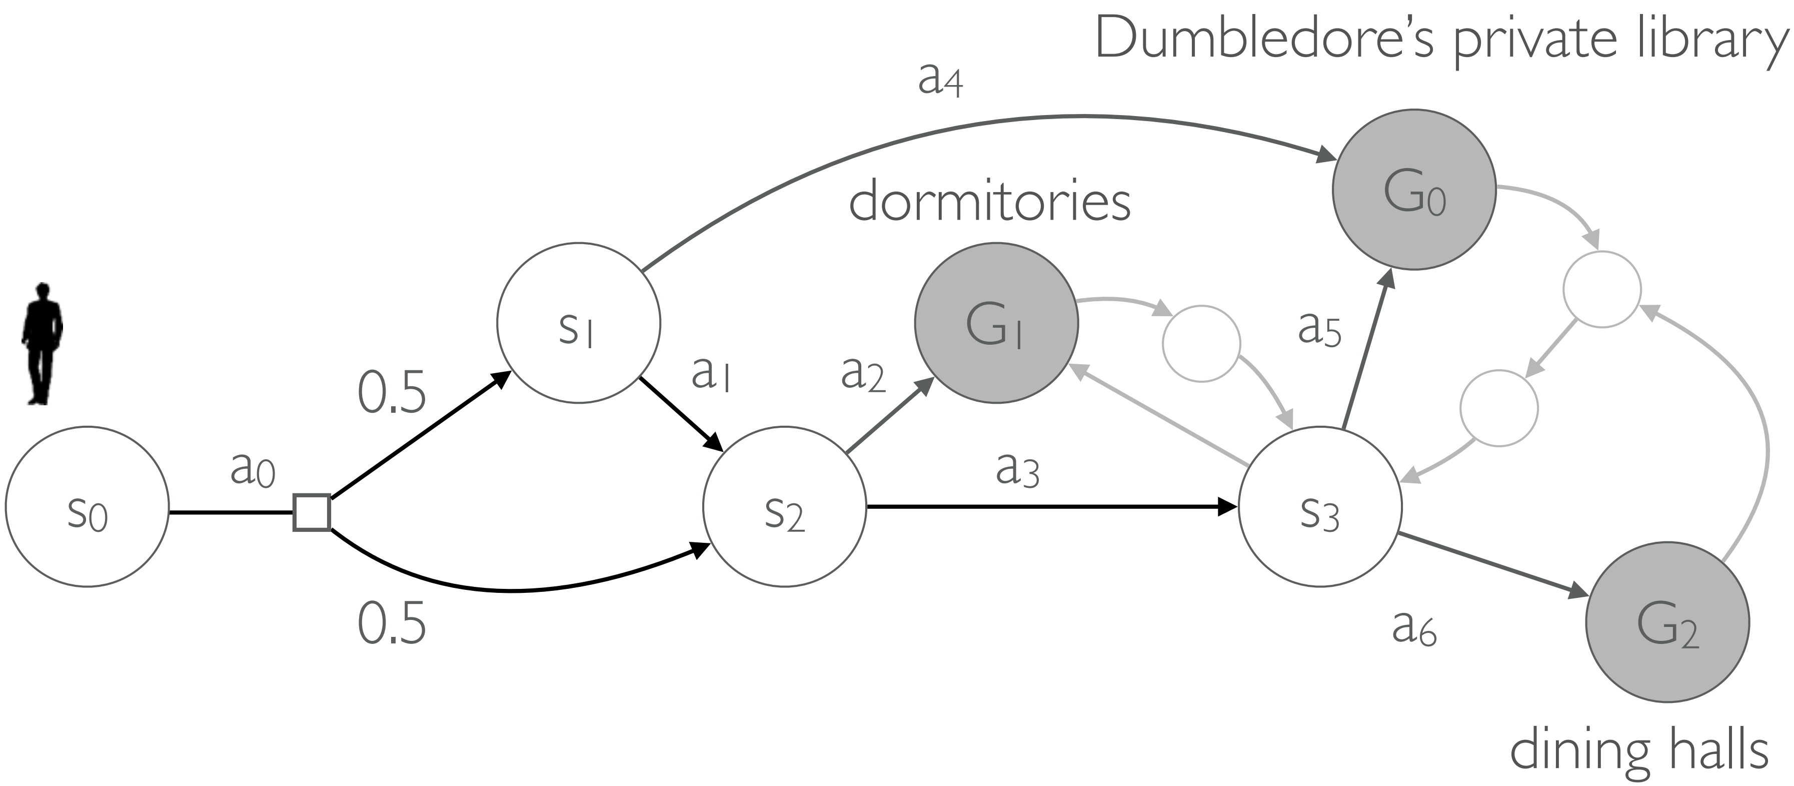


*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*

# Stochastic GRD

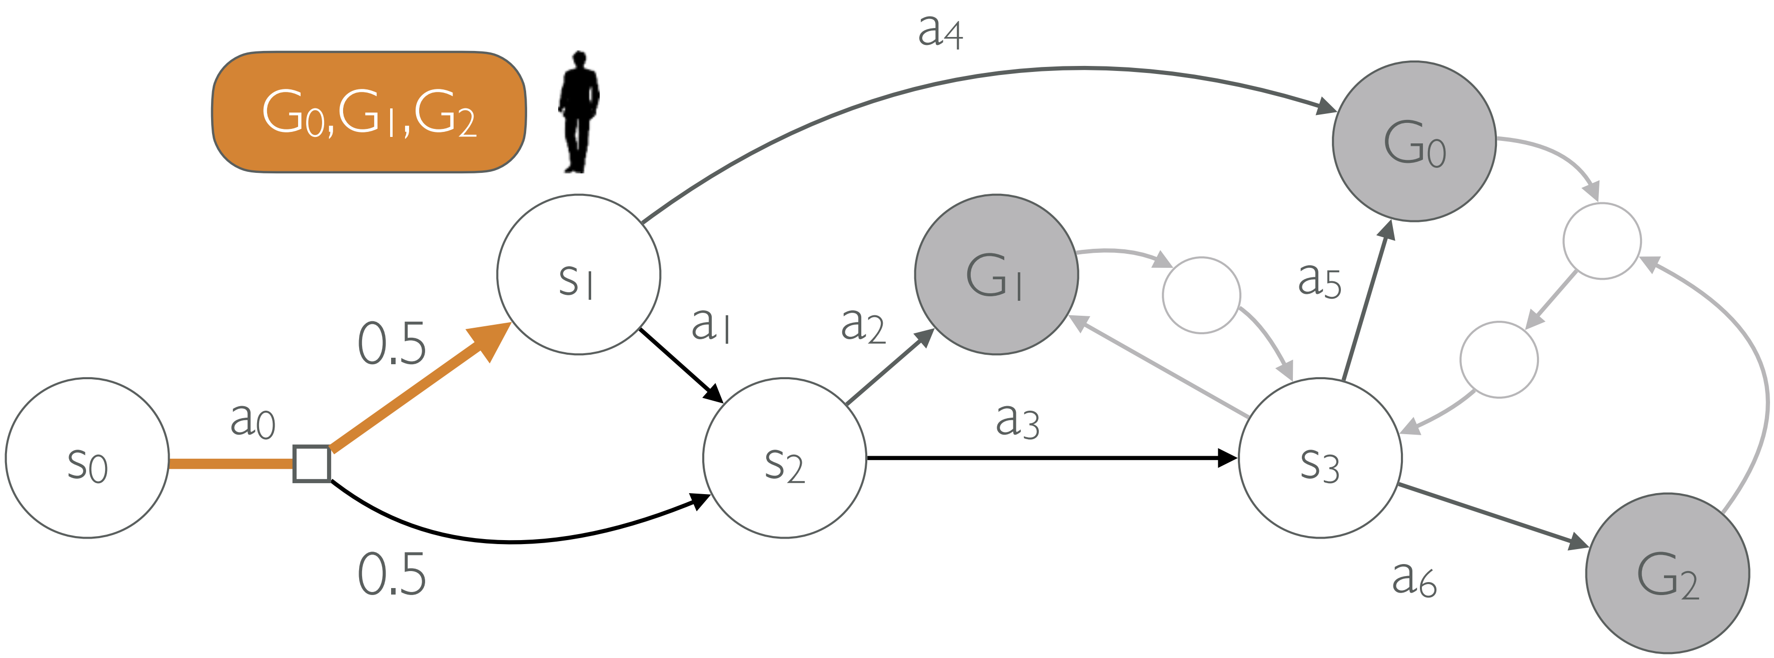


*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*

# Stochastic GRD

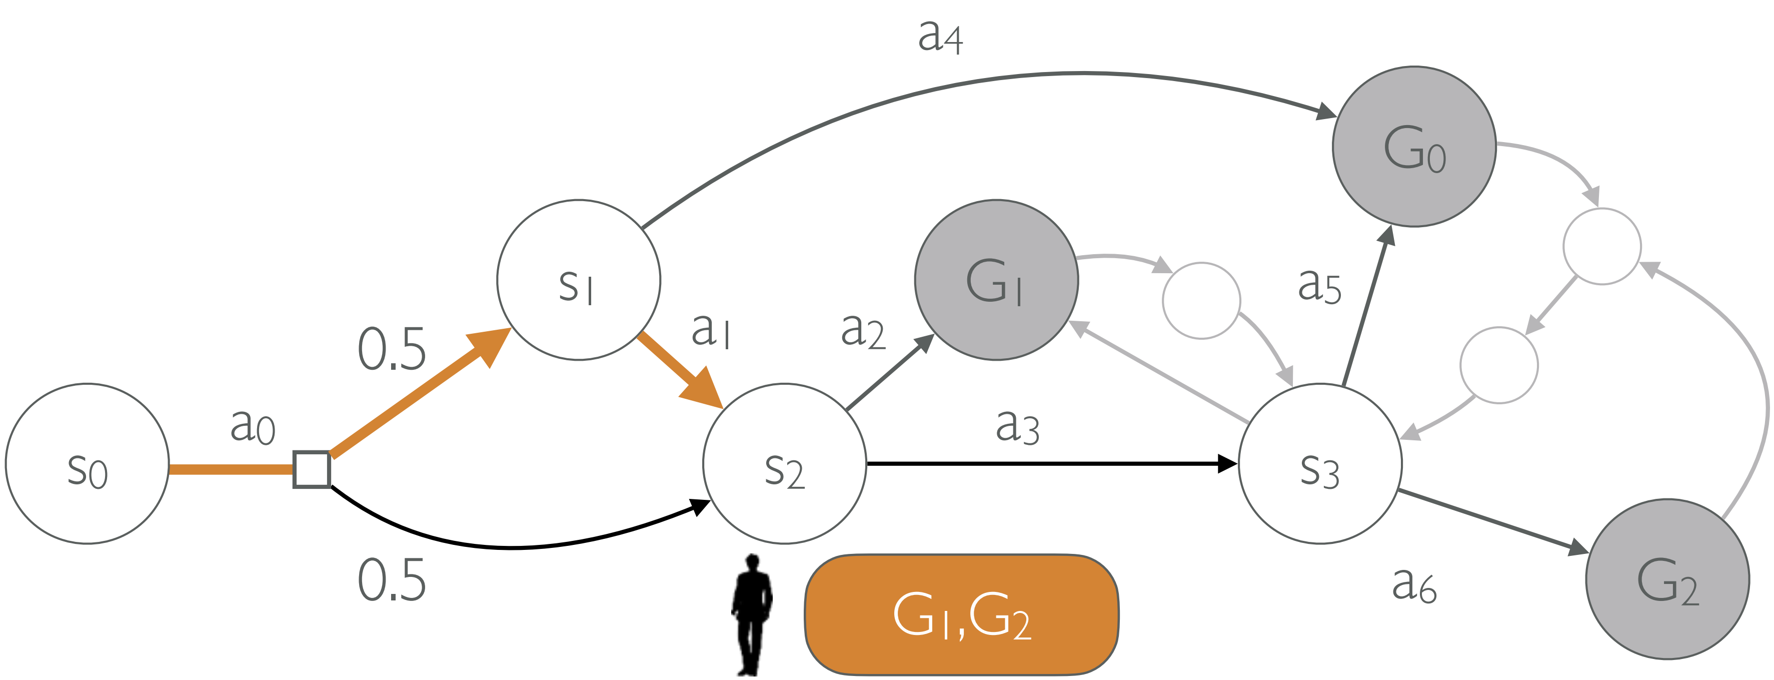


*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*

# Stochastic GRD

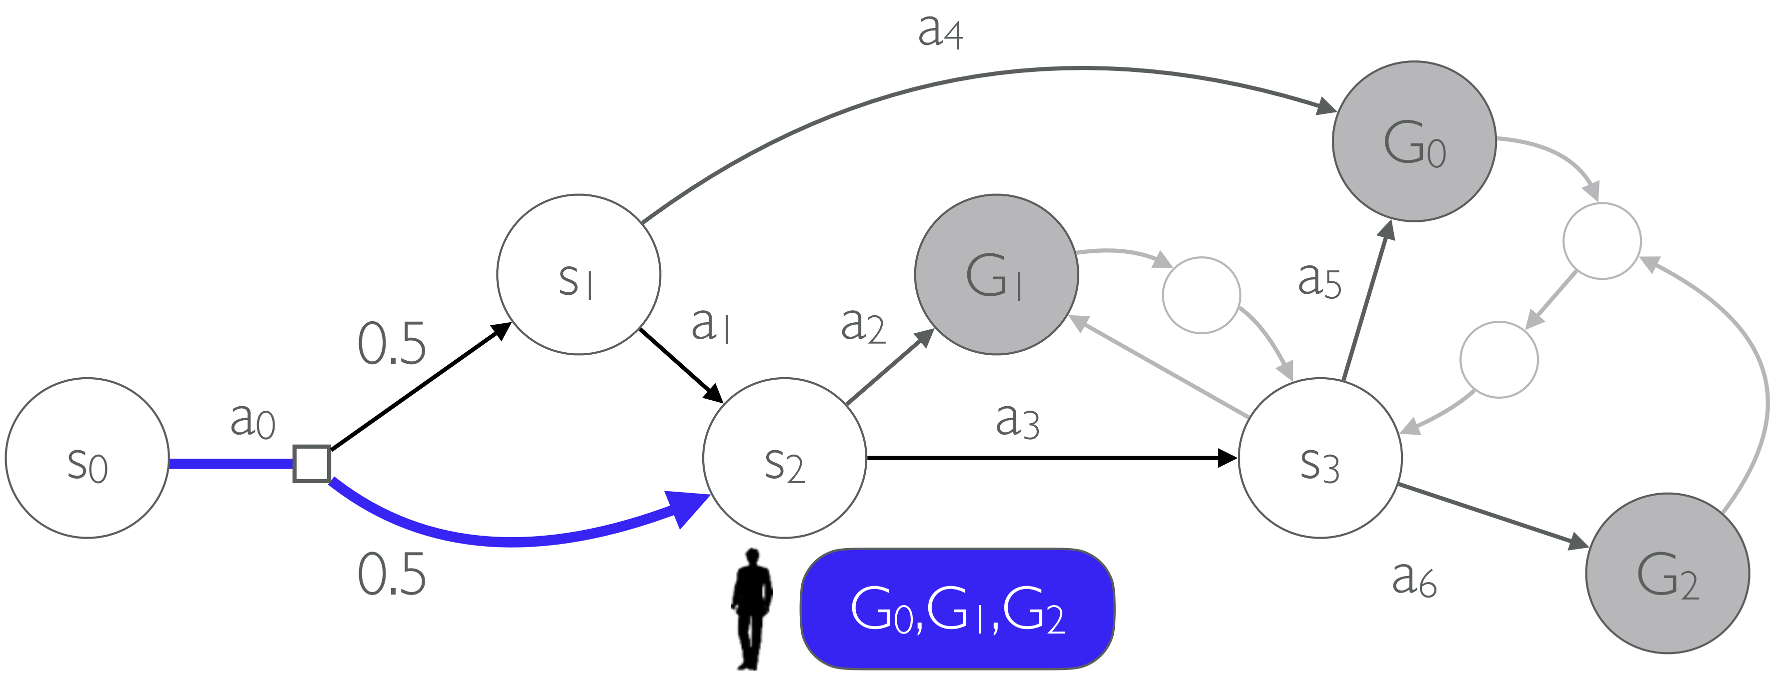


*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*

# Stochastic GRD

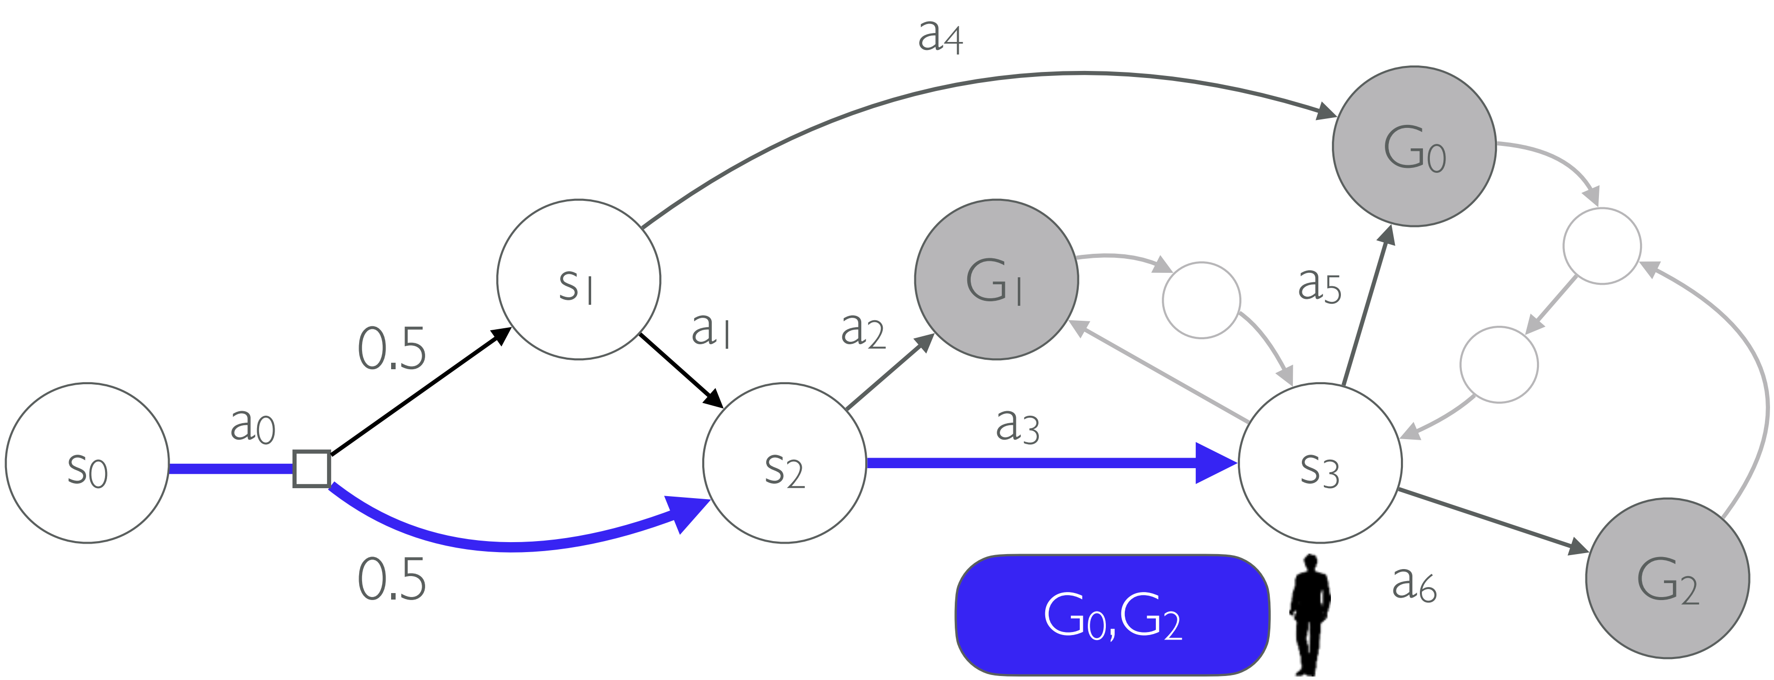


Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016

# Stochastic GRD

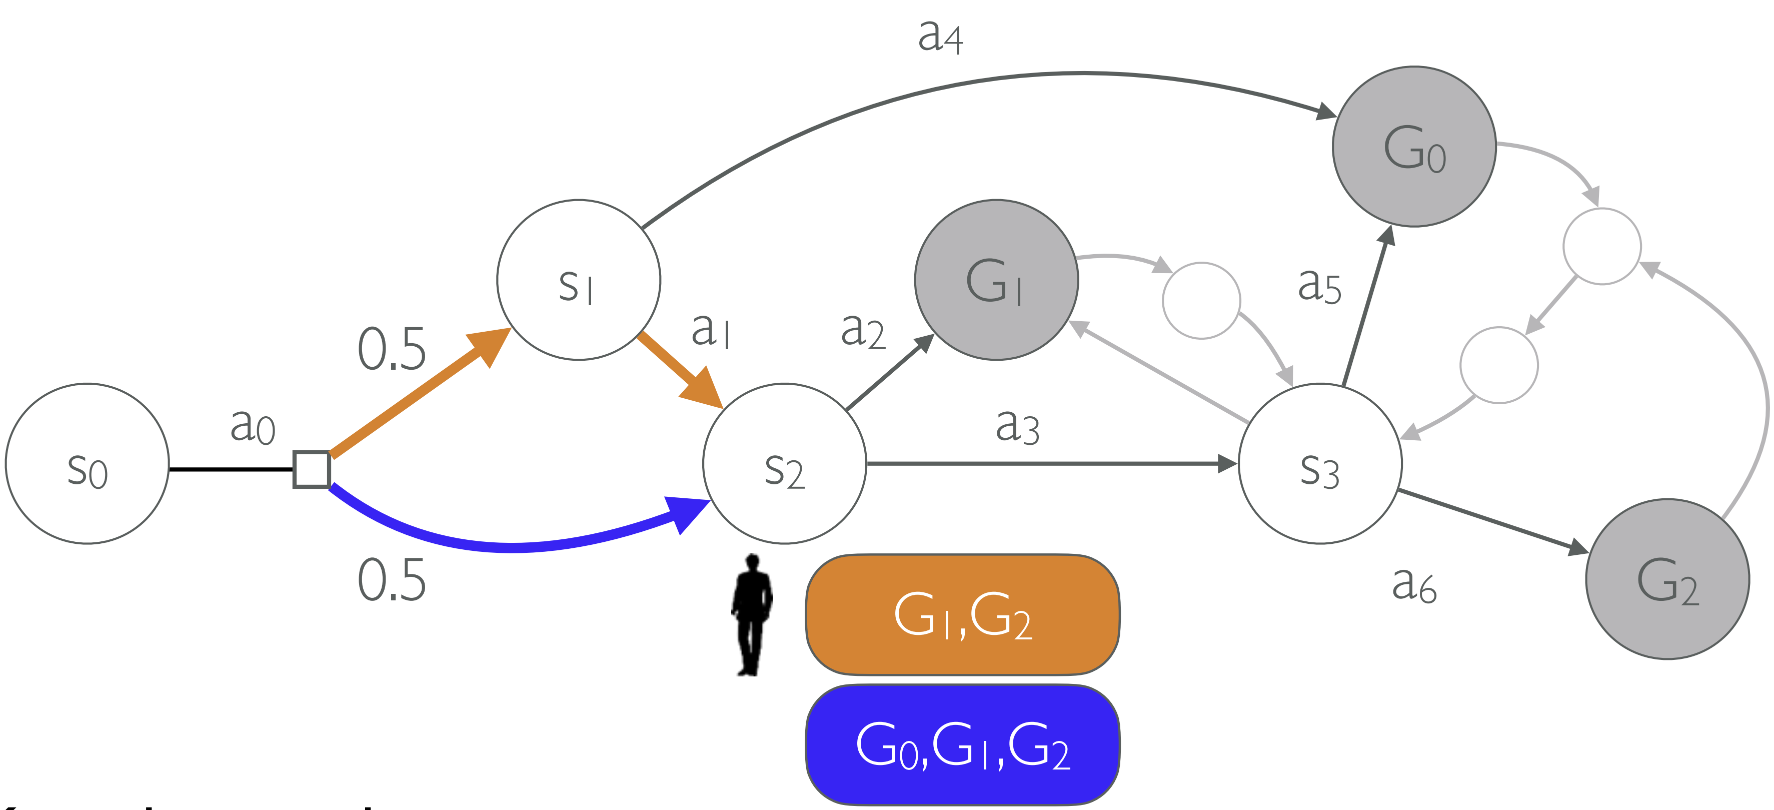


- Key observations:
  - Set of possible goals depends on the observed path to the state
  - *wcd* computation is no longer Markovian in the original state space

*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*



YODAlab

# Stochastic GRD

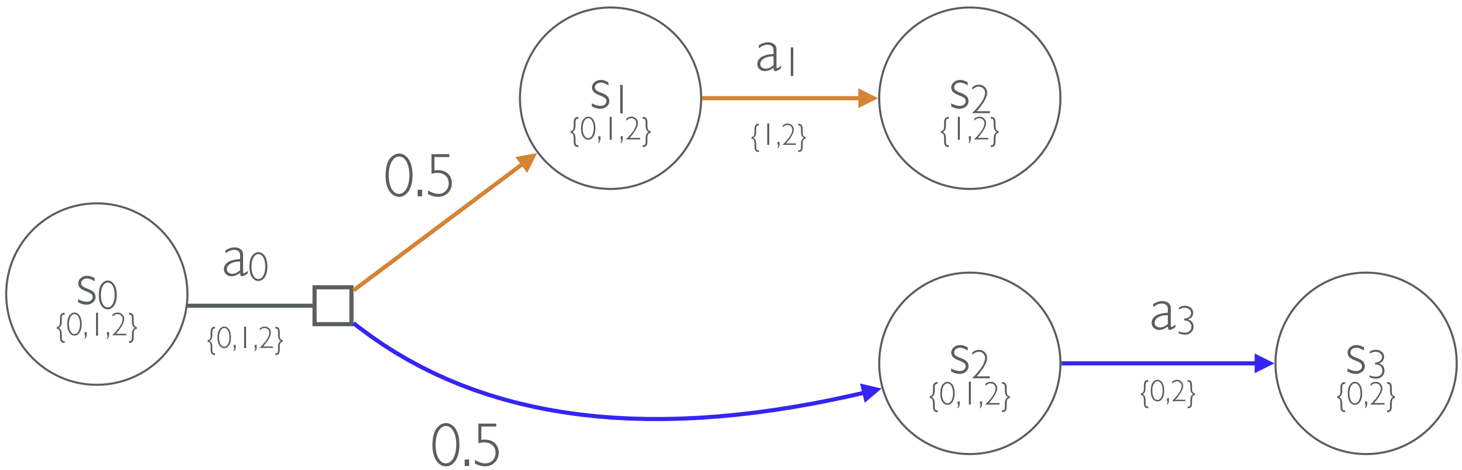


- Approach: Model the problem using augmented MDPs
  - *wcd* computation is now Markovian in the augmented state space
  - Use standard MDP algorithms (e.g., VI) to compute *wcd*
  - Agent can take max of two actions without revealing its goal (*wcd* = 2)
    - Paths: $s_0, a_0, s_1, a_1, s_2$ — or — $s_0, a_0, s_2, a_3, s_3$

*Wayllace, Hou, Yeoh, and Son: Goal Recognition Design with Stochastic Agent Action Outcomes. IJCAI 2016*

# Content

- Goal Recognition
- Goal Recognition Design
- Stochastic Goal Recognition Design
- Partially-Observable Stochastic Goal Recognition Design
- Ongoing Work: Data-driven Goal Recognition Design

# Partially-Observable S-GRD

- **Goal Recognition Design (GRD):**
  - Agent acts optimally
  - Environment is fully observable
  - Agent's action outcomes are **deterministic**

- **Stochastic GRD (S-GRD):**
  - Agent's action outcomes are **stochastic**

- **Partially-Observable S-GRD (S-GRD):**
  - Agent's action outcomes are **stochastic**
  - Environment is **partially-observable**
    - agent actions are not observable; states are partially observable
    - more realistic in some applications (robotics, navigation, etc.)

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

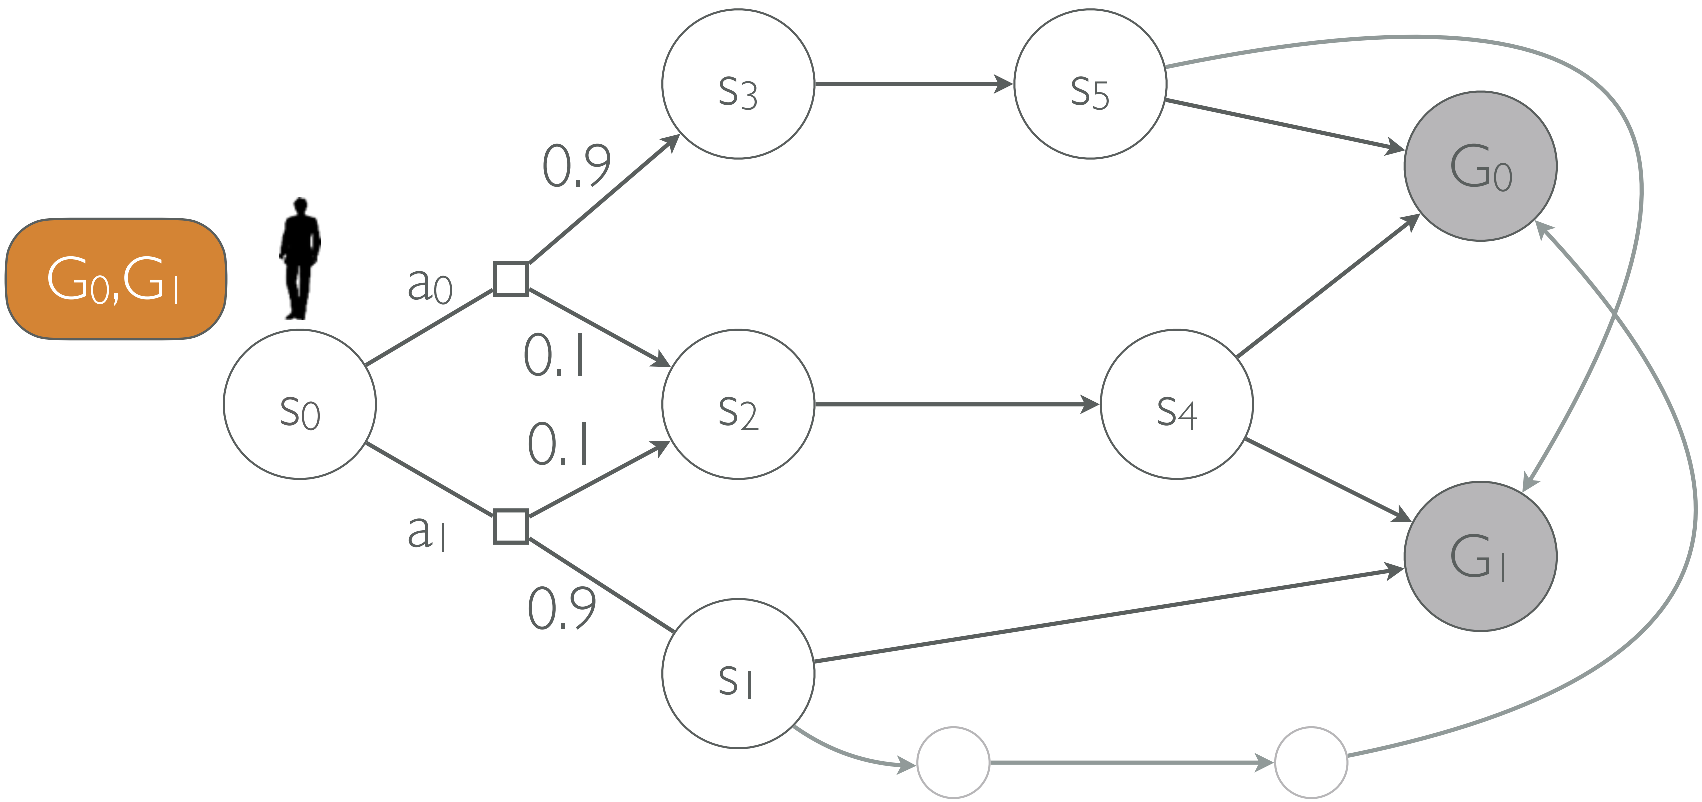


Setting: Unobservable actions, fully-observable states

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

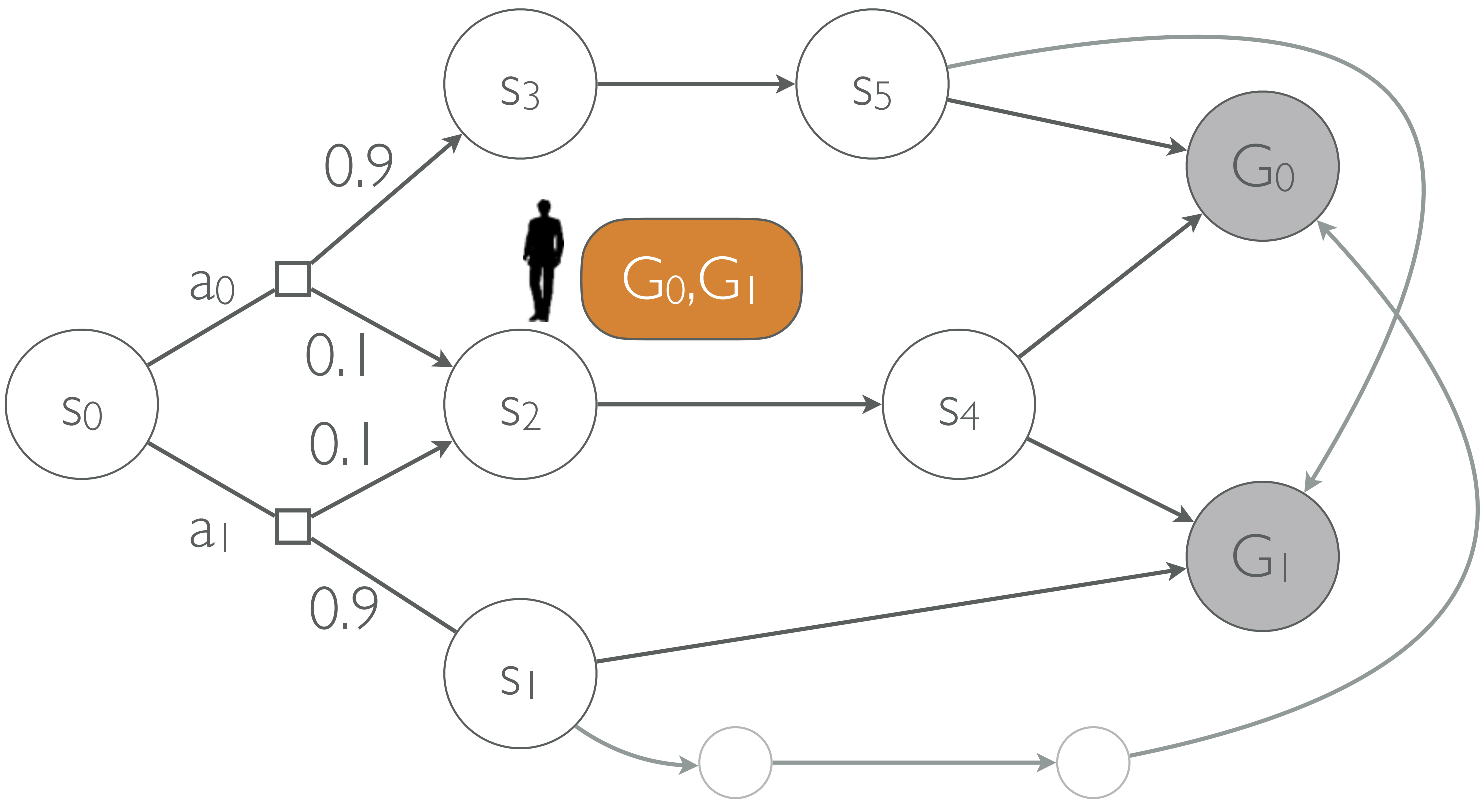


Setting: Unobservable actions, fully-observable states

Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020

# Partially-Observable S-GRD

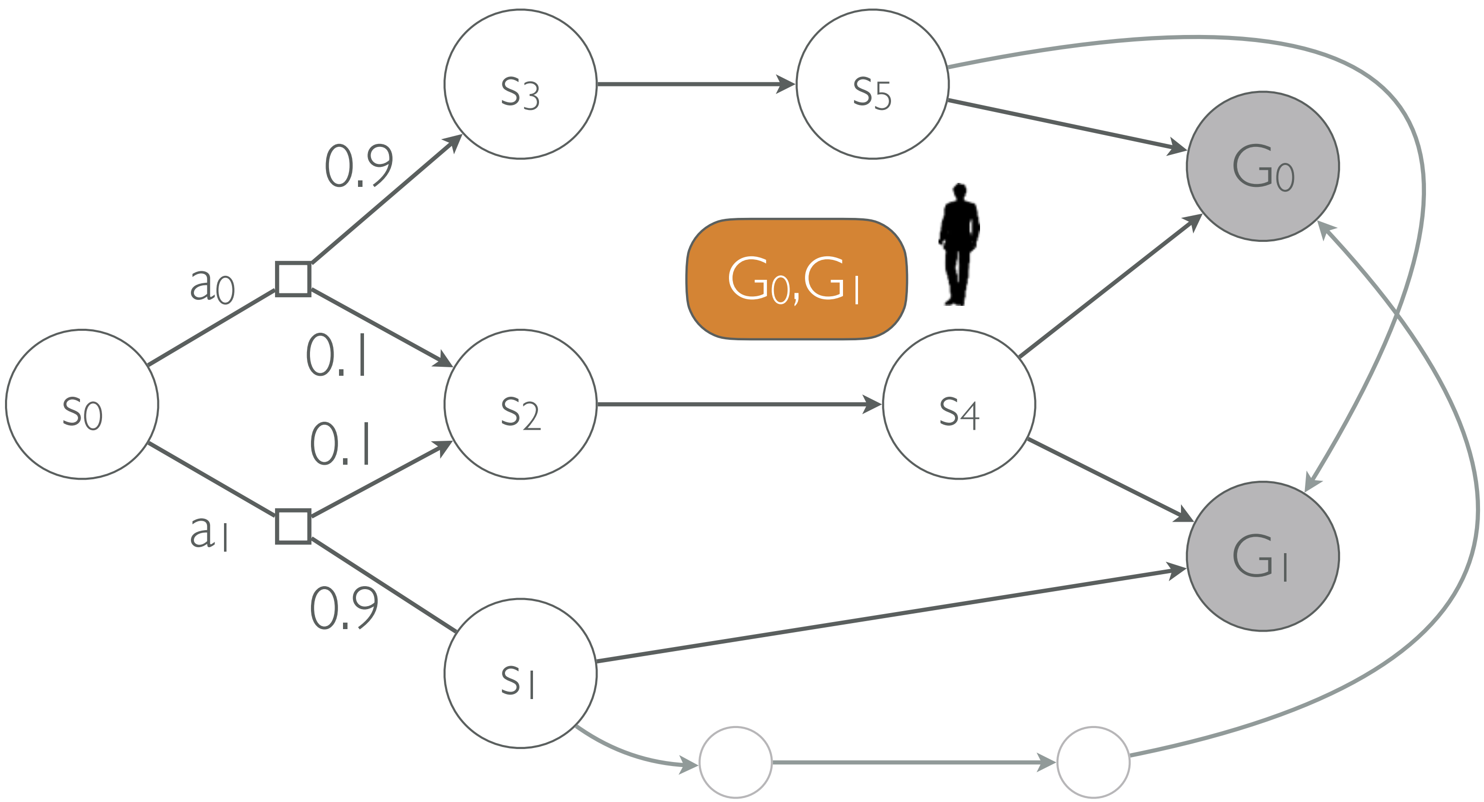


Setting: Unobservable actions, fully-observable states

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

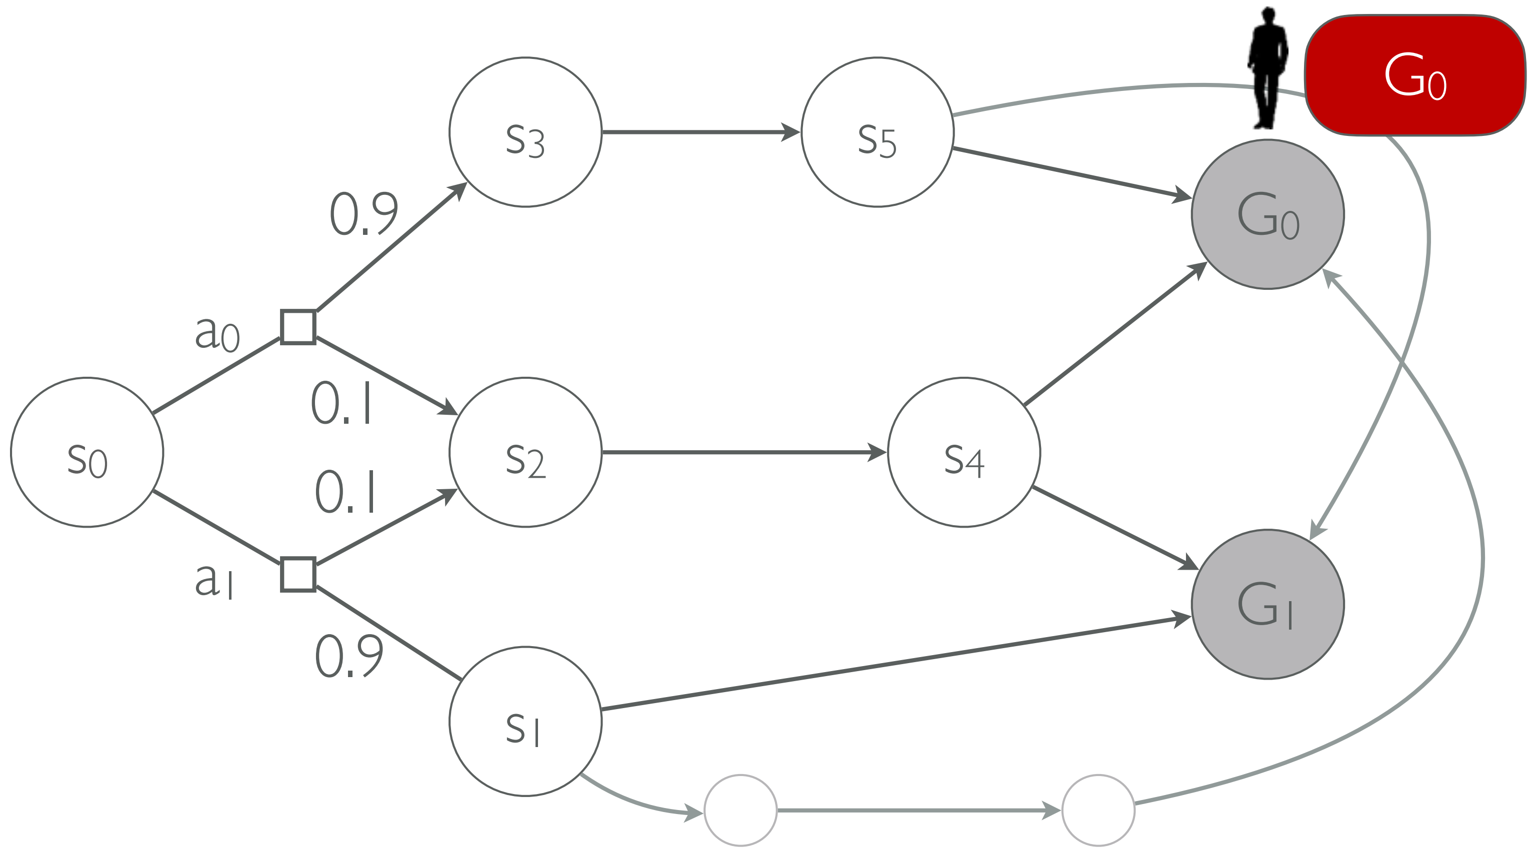


Setting: Unobservable actions, fully-observable states

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

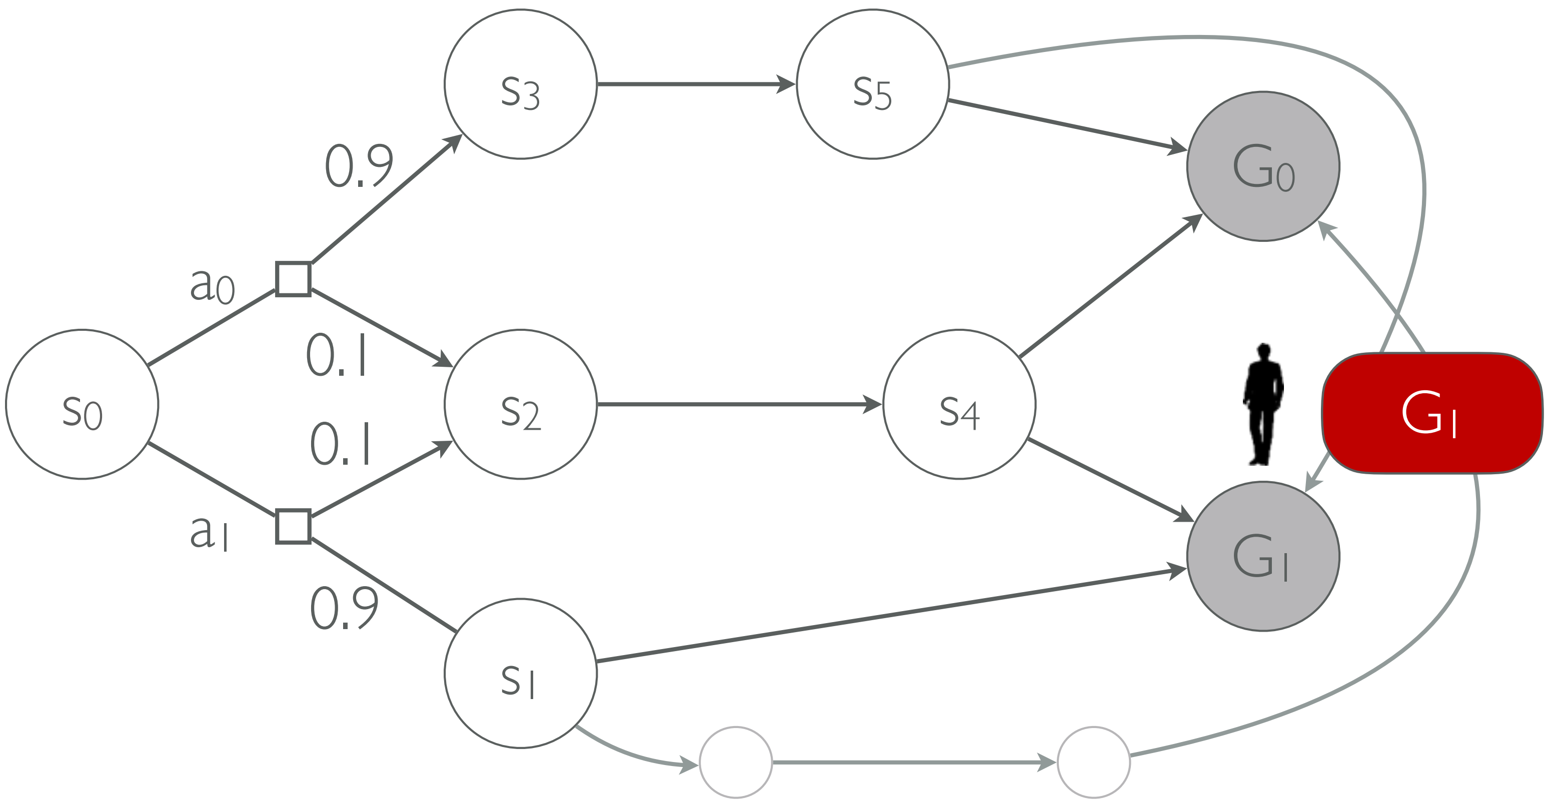


Setting: Unobservable actions, fully-observable states

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

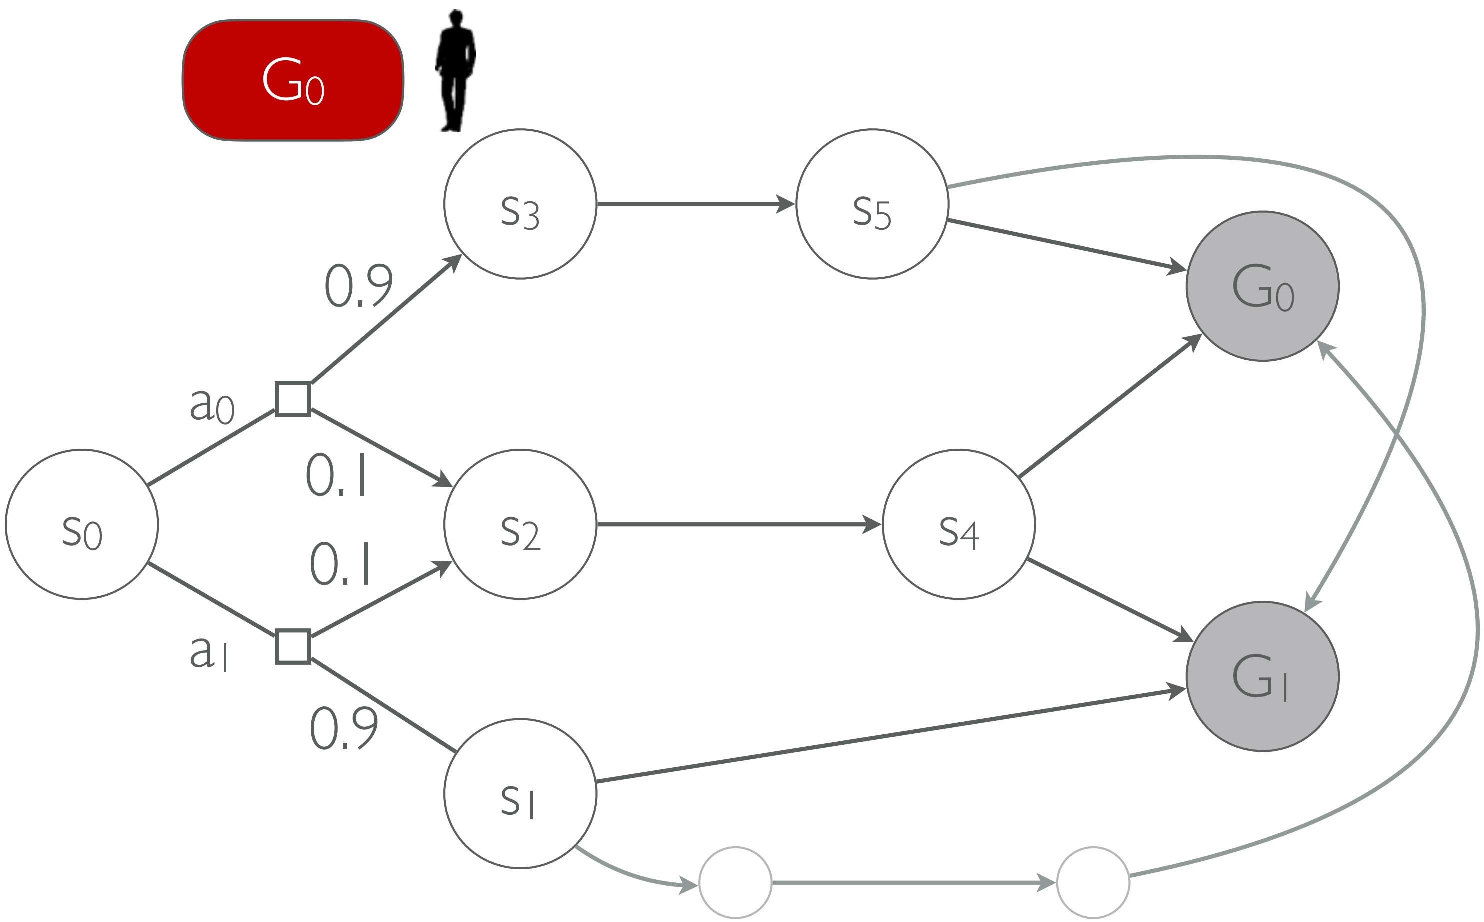


Setting: Unobservable actions, fully-observable states

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

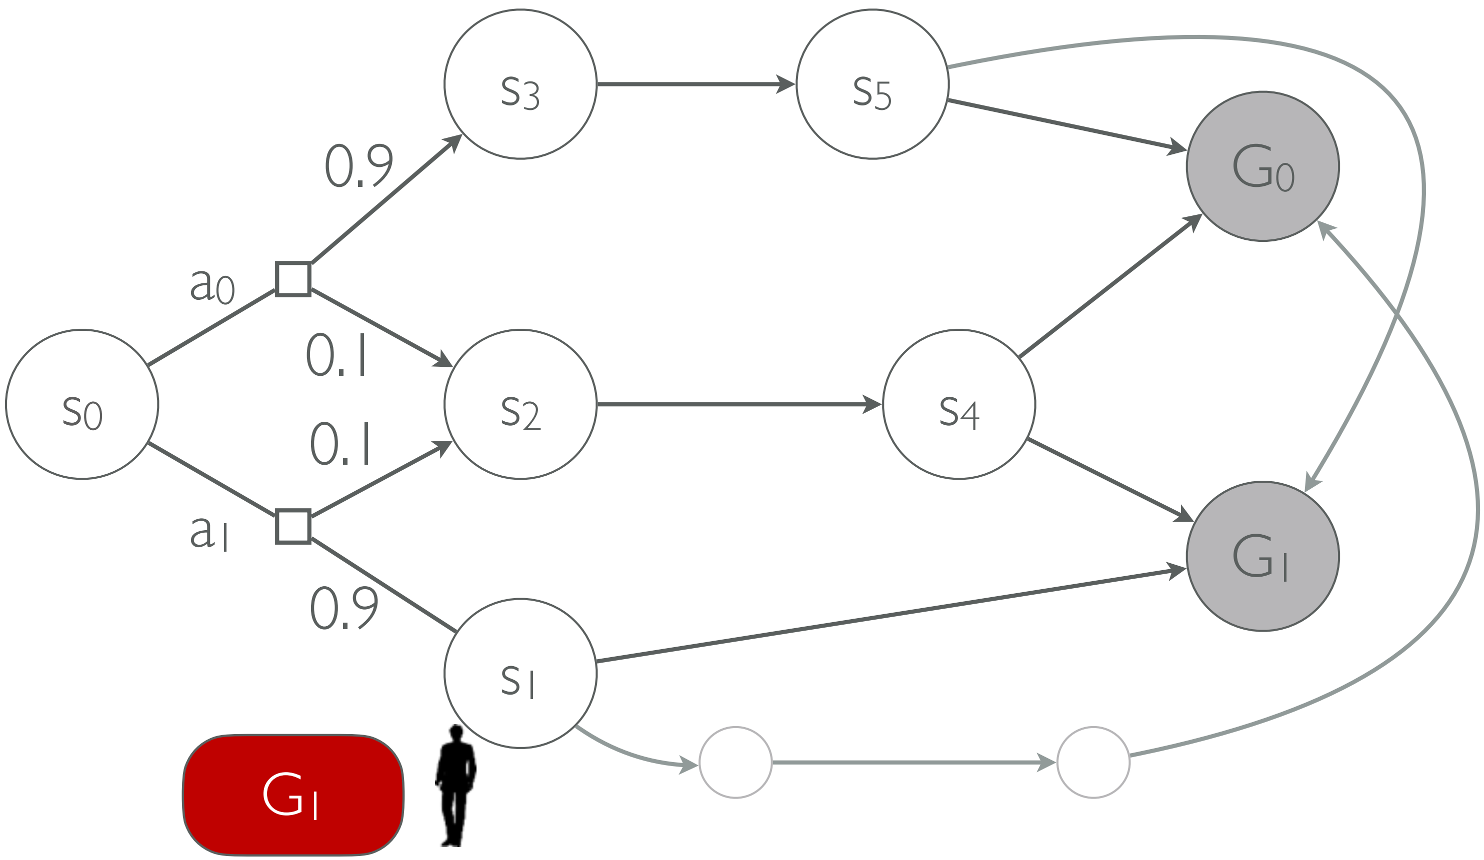


Setting: Unobservable actions, fully-observable states

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

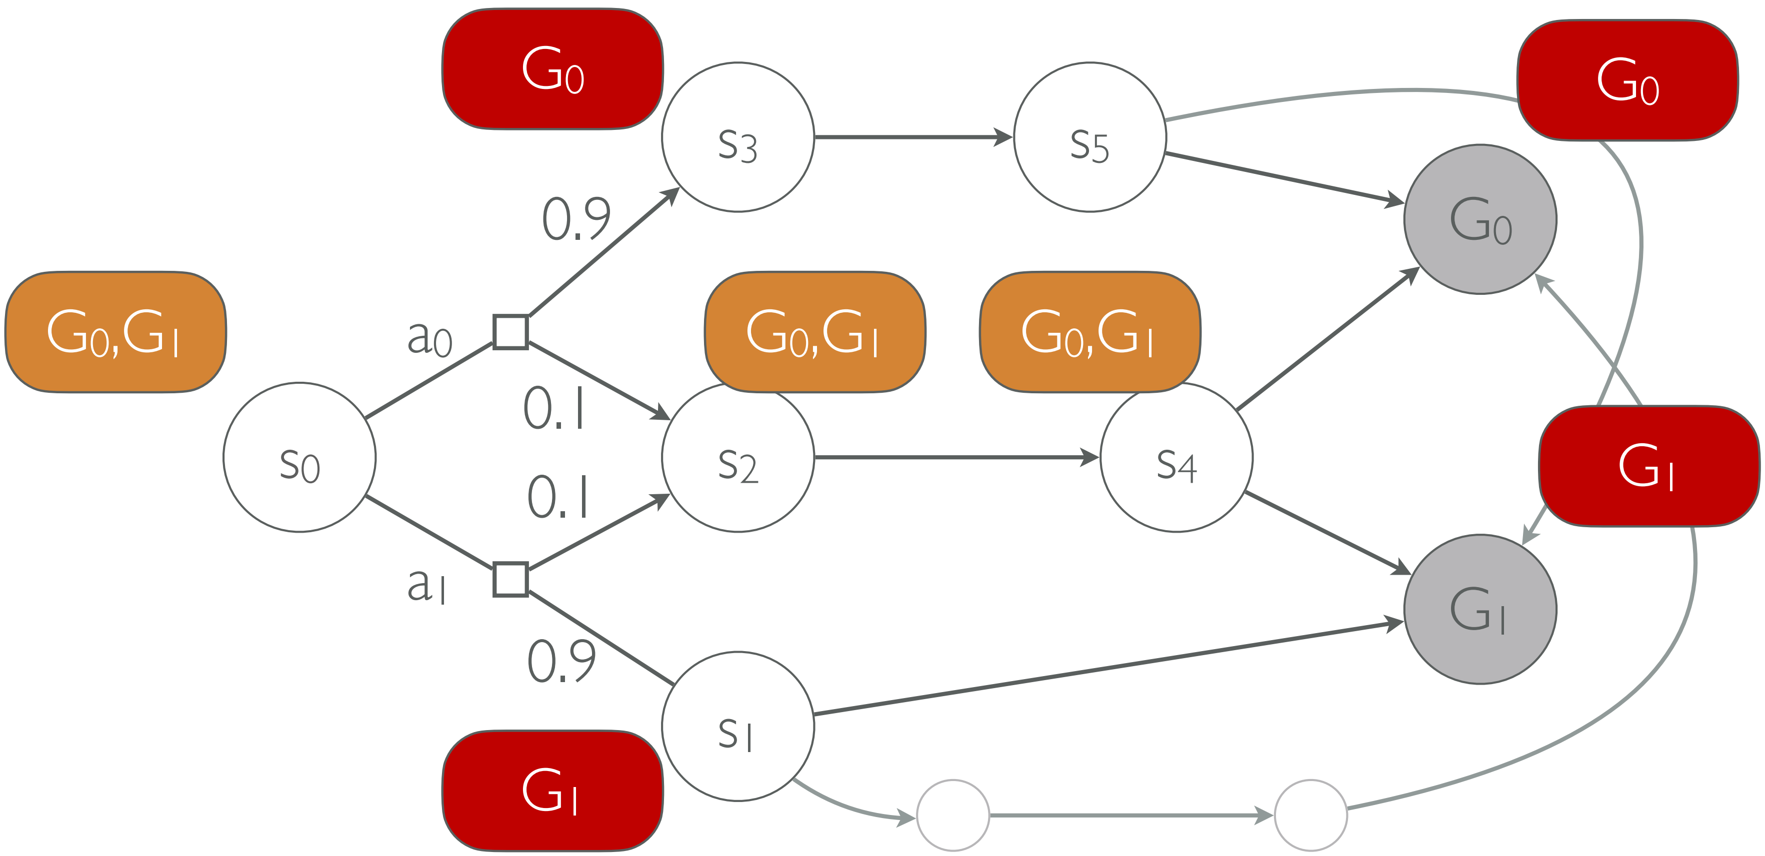


Setting: Unobservable actions, fully-observable states

$wcd = \max\,(0.9*0 + 0.1*2 \text{ for } a_0,\ \ 0.9*0 + 0.1*2 \text{ for } a_1) = 0.2$

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

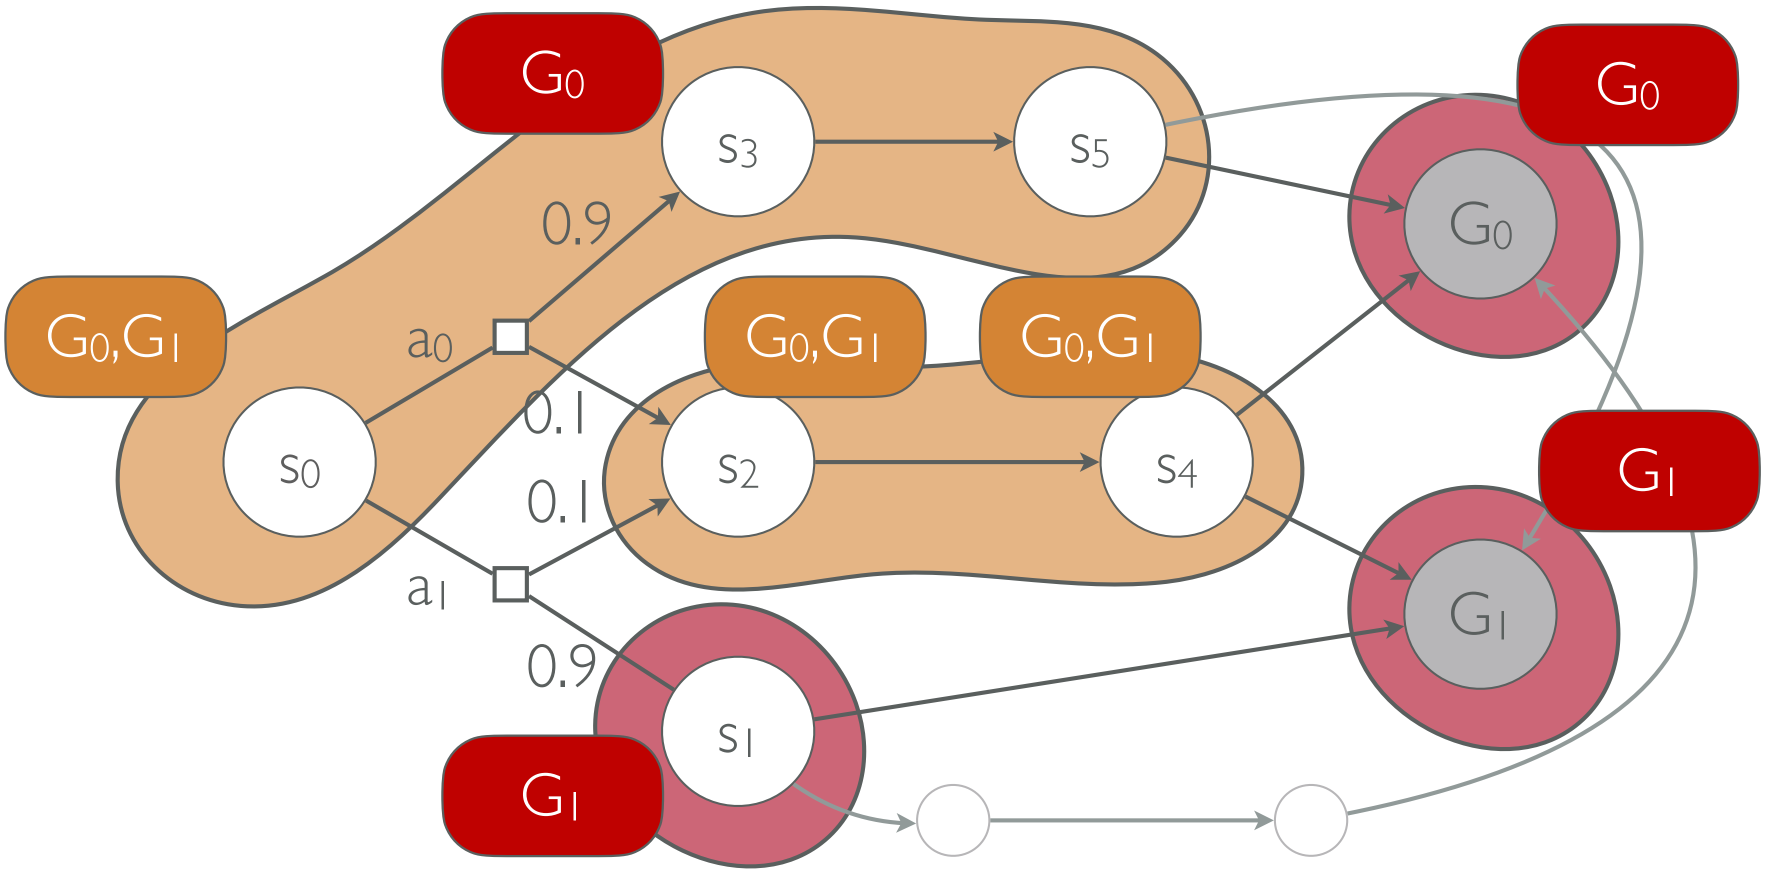


Setting: Unobservable actions, **partially-observable** states

$wcd$ = max (0.9*0 + 0.1*2 for $a_0$, 0.9*0 + 0.1*2 for $a_1$) = 0.2

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*

# Partially-Observable S-GRD

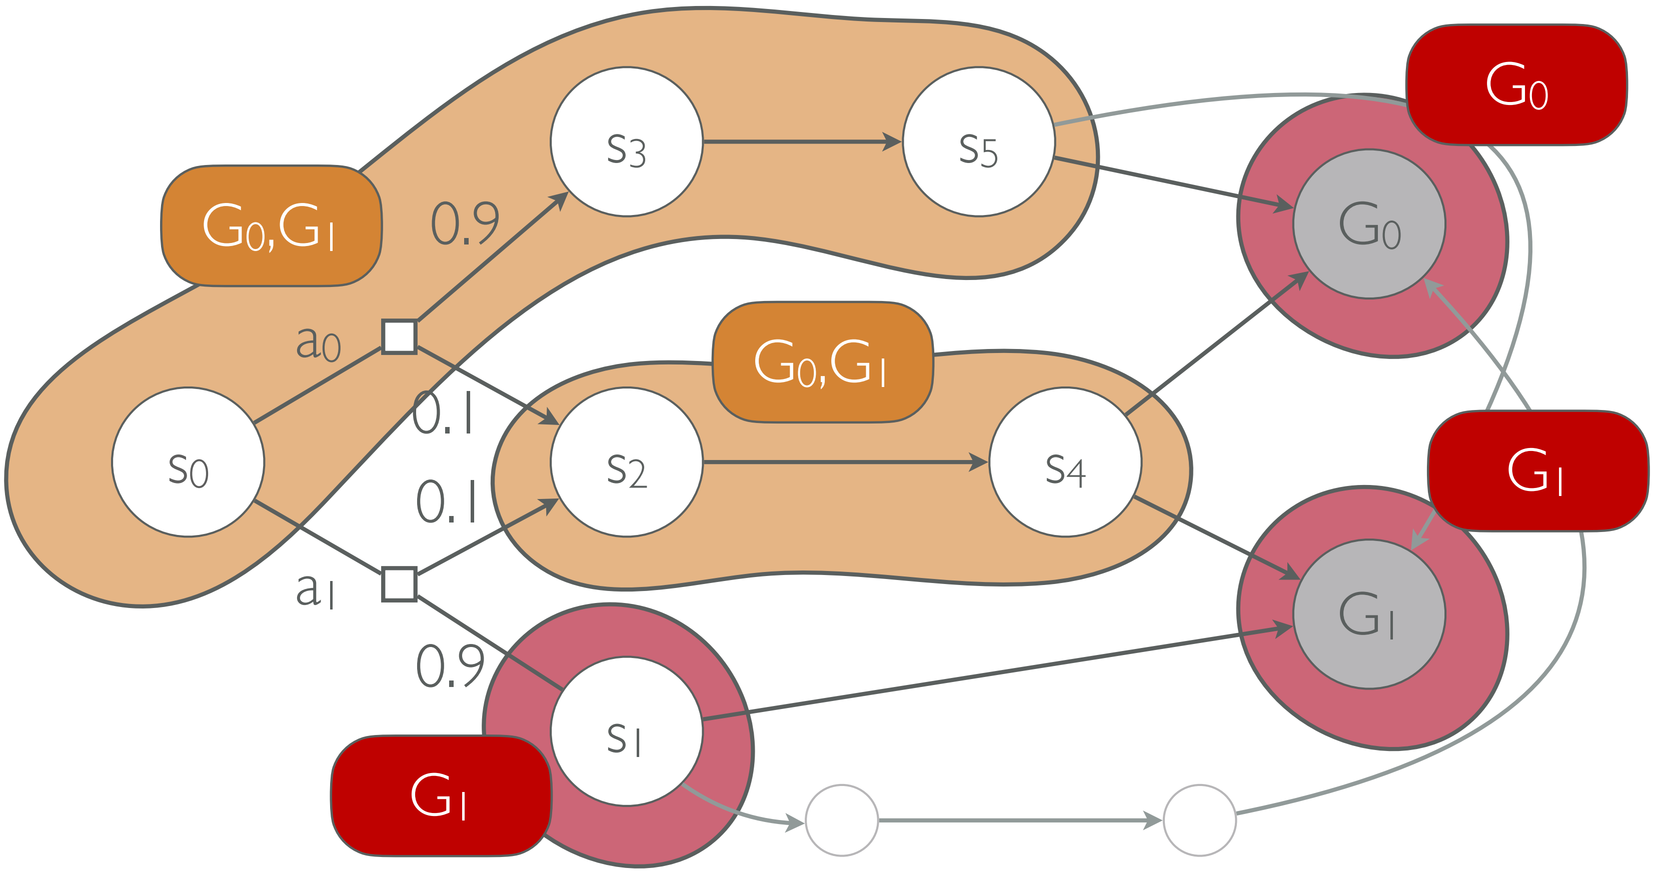


Setting: Unobservable actions, **partially-observable** states

*wcd* = max (0.9***2** + 0.1*2 for $a_0$, 0.9*0 + 0.1*2 for $a_1$) = **2**


*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*




YODAlab

# Partially-Observable S-GRD

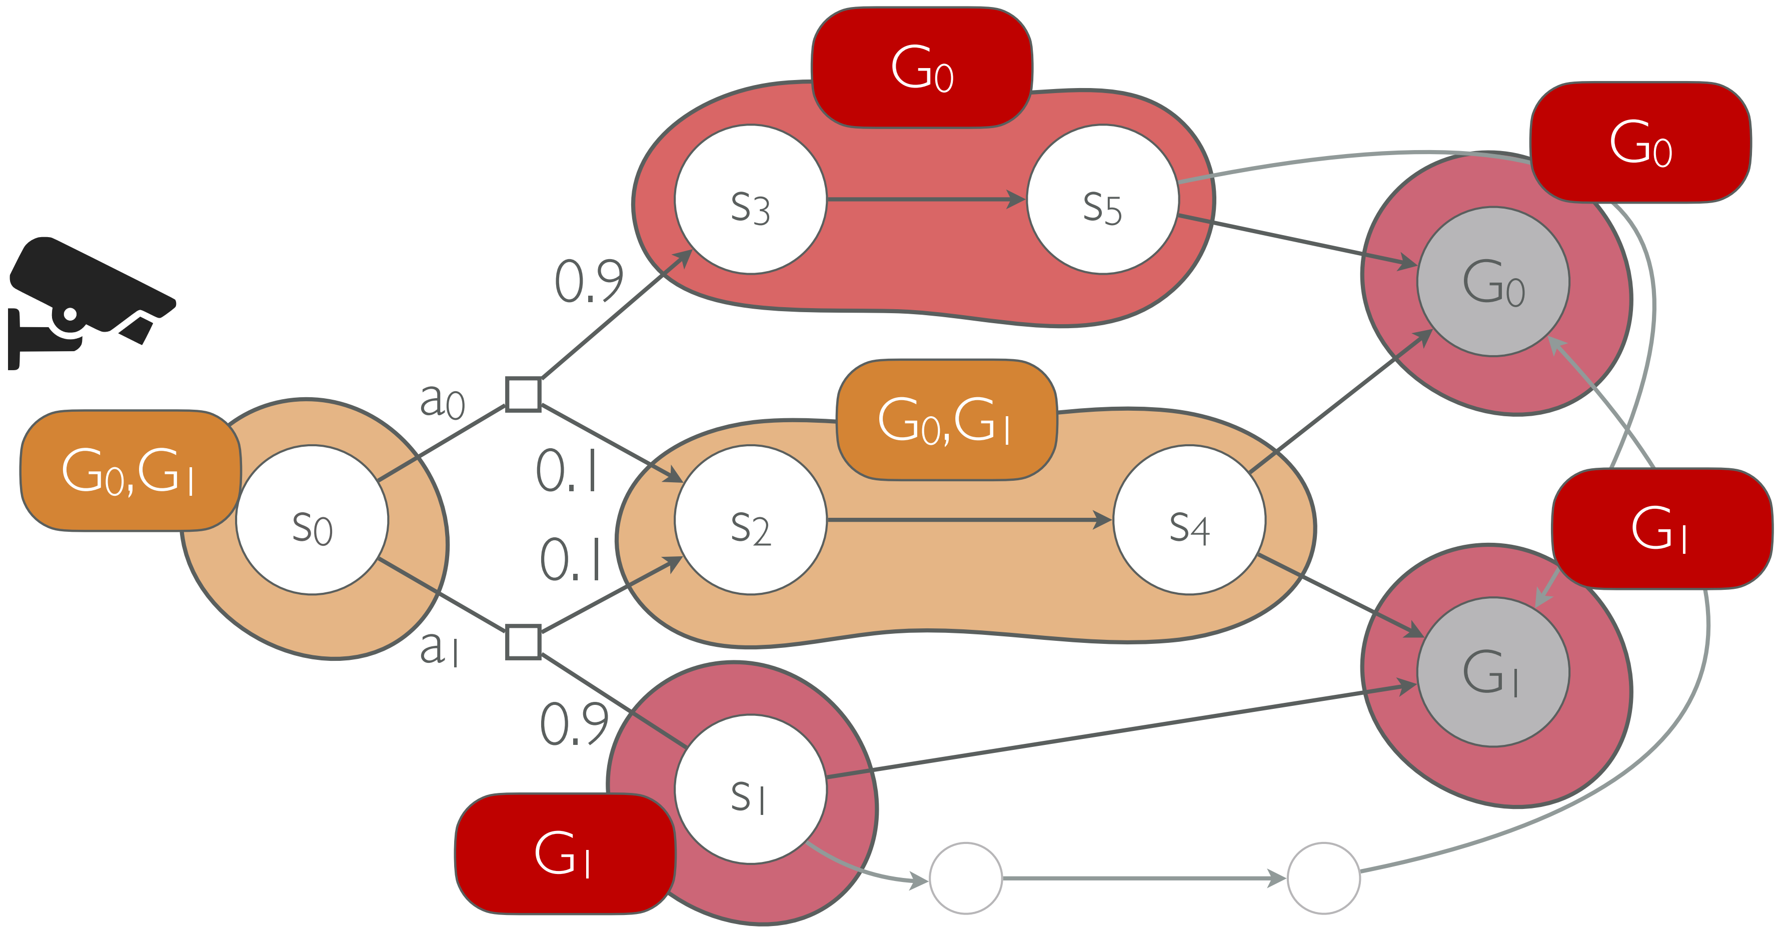


Setting: Unobservable actions, **partially-observable** states

*wcd* = max (0.9***0** + 0.1*2 for $a_0$, 0.9*0 + 0.1*2 for $a_1$) = **0.2**

*Wayllace, Keren, Gal, Karpas, Yeoh, and Zilberstein: Accounting for Observer's Partial Observability in Stochastic Goal Recognition Design. ECAI 2020*



YODAlab

# Content

- Goal Recognition
- Goal Recognition Design
- Stochastic Goal Recognition Design
- Partially-Observable Stochastic Goal Recognition Design
- Ongoing Work: Data-driven Goal Recognition Design

# Data-Driven GRD

| Literature | Suboptimal Agent | Stochastic Actions | Partially Observable Environment | Action Removal | Sensor Refinement | Action Conditioning |
|---|---|---|---|---|---|---|
| Keren *et al*. (ICAPS 2014) | | | | ✓ | | |
| Keren *et al*. (AAAI 2015) | ✓ | | | ✓ | | |
| Son *et al*. (AAAI 2016) | | | | ✓ | | |
| Keren *et al*. (AAAI 2016) | ✓ | | ✓ | ✓ | ✓ | |
| Keren *et al*. (IJCAI 2016) | ✓ | | ✓ | ✓ | | |
| Wayllace *et al*. (IJCAI 2016) | | ✓ | | ✓ | | |
| Ang *et al*. (IJCAI 2017) | | | | ✓ | | |
| Wayllace *et al*. (IJCAI 2017) | | ✓ | | ✓ | | |
| Keren *et al*. (ICAPS 2018) | ✓ | | ✓ | ✓ | ✓ | ✓ |
| Keren *et al*. (JAIR 2018) | ✓ | | ✓ | ✓ | ✓ | ✓ |
| Wayllace *et al*. (AAAI 2019) | ✓ | ✓ | | ✓ | | |
| Wayllace *et al*. (ECAI 2020) | | ✓ | ✓ | ✓ | ✓ | |

# Data-Driven GRD

**Common Assumption:**
Uses a worst-case measure (across all possible agent behaviors) for the difficulty of the goal recognition problem

The worst-case measure often does not reflect the expected agent behavior, especially if agent is human

Also computationally expensive and does not scale well

| | Suboptimal | Stochastic | Partially | Action | Sensor | Action |
|---|---|---|---|---|---|---|
| Kere | | | | | | |
| Ker | | | | | | |
| Sc | | | | | | |
| Ker | | | | | | |
| Ker | | | | | | |
| Wayl | | | | | | |
| Ar | | | | | | |
| Wayllace *et al*. (IJCAI 2017) | | ✓ | | ✓ | | |
| Keren *et al*. (ICAPS 2018) | ✓ | | ✓ | ✓ | ✓ | ✓ |
| Keren *et al*. (JAIR 2018) | ✓ | | ✓ | ✓ | ✓ | ✓ |
| Wayllace *et al*. (AAAI 2019) | ✓ | ✓ | | ✓ | | |
| Wayllace *et al*. (ECAI 2020) | | ✓ | ✓ | ✓ | ✓ | |



YODAlab

# Data-Driven GRD

**Common Assumption:**
Uses a worst-case measure (across all possible agent behaviors) for the difficulty of the goal recognition problem

The worst-case measure often does not reflect the expected agent behavior, especially if agent is human

Also computationally expensive and does not scale well

**Data-Driven GRD:**
Use a data-driven approach to learn a predictor for the expected difficulty of the problem for a variety of agent behaviors, *including human behavior*

# Data-Driven GRD

- **Predictive Module:**
  - Curate a training dataset: Tuplets of environment, behavior, and *wcd*
  - Behaviors: Optimal, bounded suboptimal, human
    - Collected human behavioral data for navigating to a goal in a grid
    - Trained a multilayer perceptron model to predict the next action
  - CNN-based model that takes as input an environment and outputs a predicted *wad*

*Kasumba, Yu, Ho, Keren, and Yeoh: Data-Driven Goal Recognition Design for General Behavioral Agents. arXiv 2024*

# Data-Driven GRD

- **Predictive Module:**
  - Curate a training dataset: Tuplets of environment, behavior, and *wcd*
  - Behaviors: Optimal, bounded suboptimal, human
  - CNN-based model that takes as input an environment and outputs a predicted *wcd*

- **Design Module:**
  - Transforms the GRD problem into an unconstrained optimization problem using Lagrangian relaxation:

$$L = wcd(w', h) + \lambda(c(w, w') - B)$$

    - $wcd(w', h)$: *wcd* of environment $w'$ with behavioral model $h$
    - $c(w, w')$: cost of changing current environment $w$ to environment $w'$
    - $B$: cost budget

*Kasumba, Yu, Ho, Keren, and Yeoh: Data-Driven Goal Recognition Design for General Behavioral Agents. arXiv 2024*

# Data-Driven GRD

- **Predictive Module:**
  - Curate a training dataset: Tuplets of environment, behavior, and *wcd*
  - Behaviors: Optimal, bounded suboptimal, human
  - CNN-based model that takes as input an environment and outputs a predicted *wcd*

- **Design Module:**
  - Transforms the GRD problem into an unconstrained optimization problem using Lagrangian relaxation:

$$L = wcd(w', h) + \lambda(c(w, w') - B)$$

  - Perform gradient descent on the relaxed Lagrangian; at each step:
    - obtain a vector of possible changes and their magnitude
    - select element with the highest gradient value and make the corresponding change

*Kasumba, Yu, Ho, Keren, and Yeoh: Data-Driven Goal Recognition Design for General Behavioral Agents. arXiv 2024*

# Data-Driven GRD

- **User Study: Accuracy of human goal inference:**
  - *Does modifying the environment to reduce predicted wcd result in environments that are easier for humans to infer goals?*

  - Generated 30 initial environments
  - Modified them using:
    - Greedy: Using predicted *wcd* from our predictive module
    - Proposed (opt-bhvr): Using our design module, but assuming optimal agent behavior
    - Proposed (data-driven): Using our predictive and design modules
  - Asked users to guess the goal of the observed agent

*Kasumba, Yu, Ho, Keren, and Yeoh: Data-Driven Goal Recognition Design for General Behavioral Agents. arXiv 2024*

# Data-Driven GRD

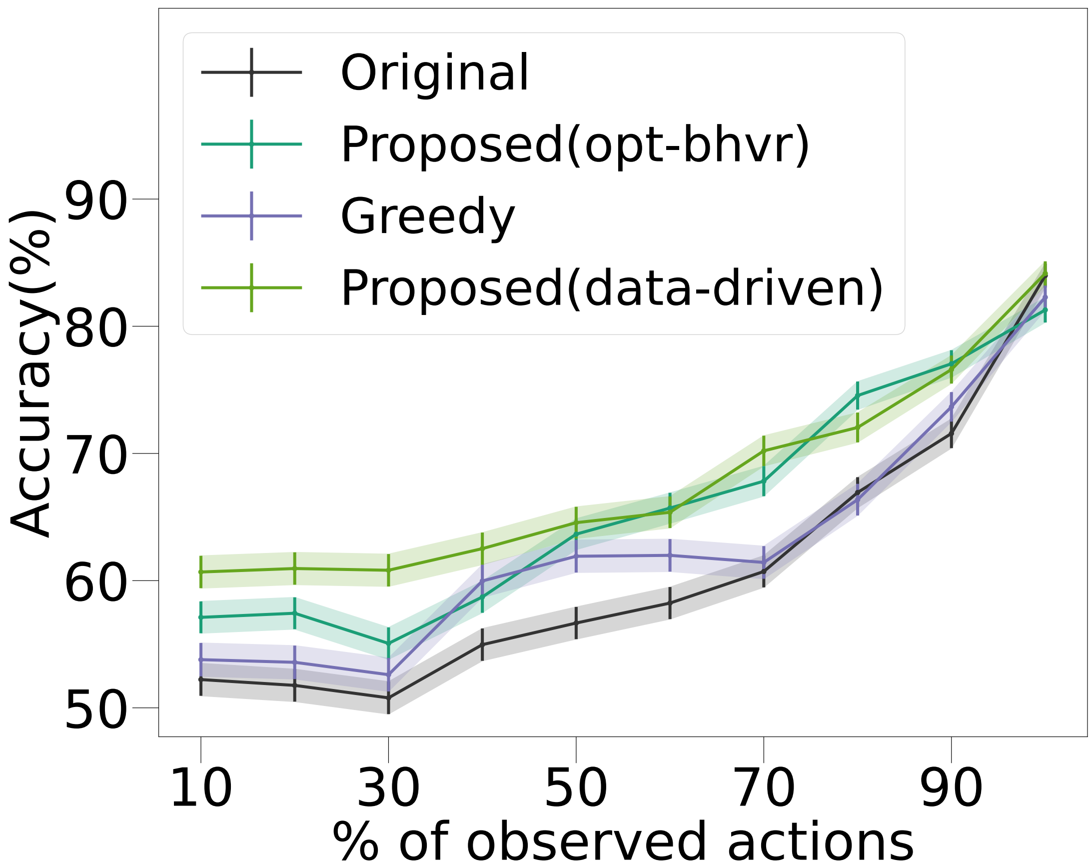


- Data-driven approach allows users to more accurately guess the goal of the observed agent

*Kasumba, Yu, Ho, Keren, and Yeoh: Data-Driven Goal Recognition Design for General Behavioral Agents. arXiv 2024*

# Conclusions

- **Goal Recognition:**
  - Seek to identify the goal G of an agent based on its observations O
- **Goal Recognition Design (GRD):**
  - Seek to ***modify/design*** the underlying environment to improve goal recognition
  - Orthogonal to goal recognition; advances made will complement advances in goal recognition
- **Partially-Observable Stochastic GRD:**
  - Generalizes GRD to partially-observable environments and stochastic action outcomes
- **Data-Driven GRD:**
  - Uses ML to account for human behaviors in GRD

# Acknowledgments

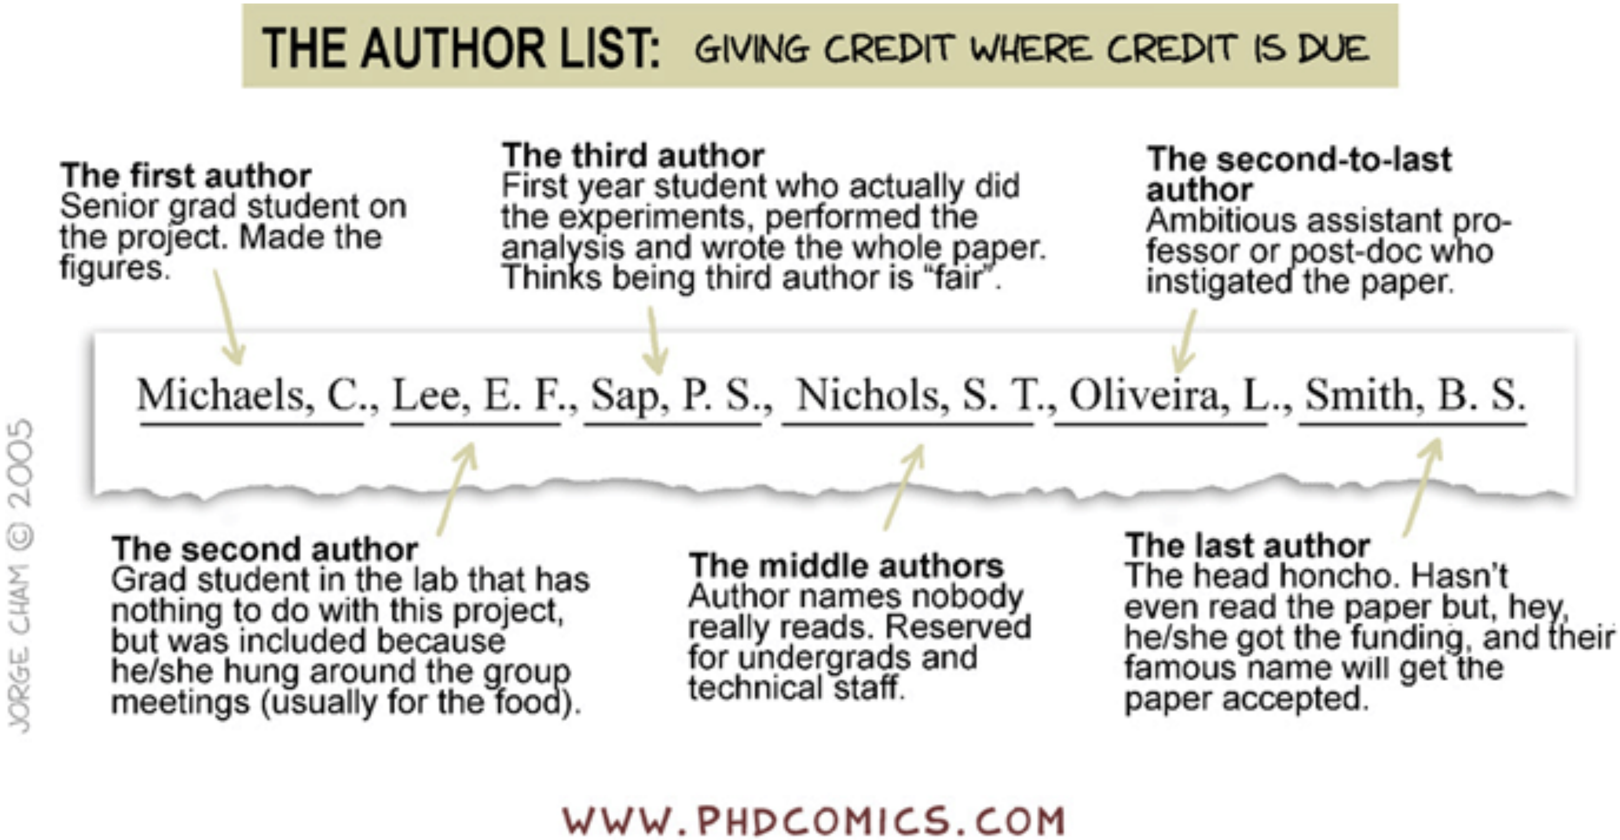




YODAlab

# Acknowledgments

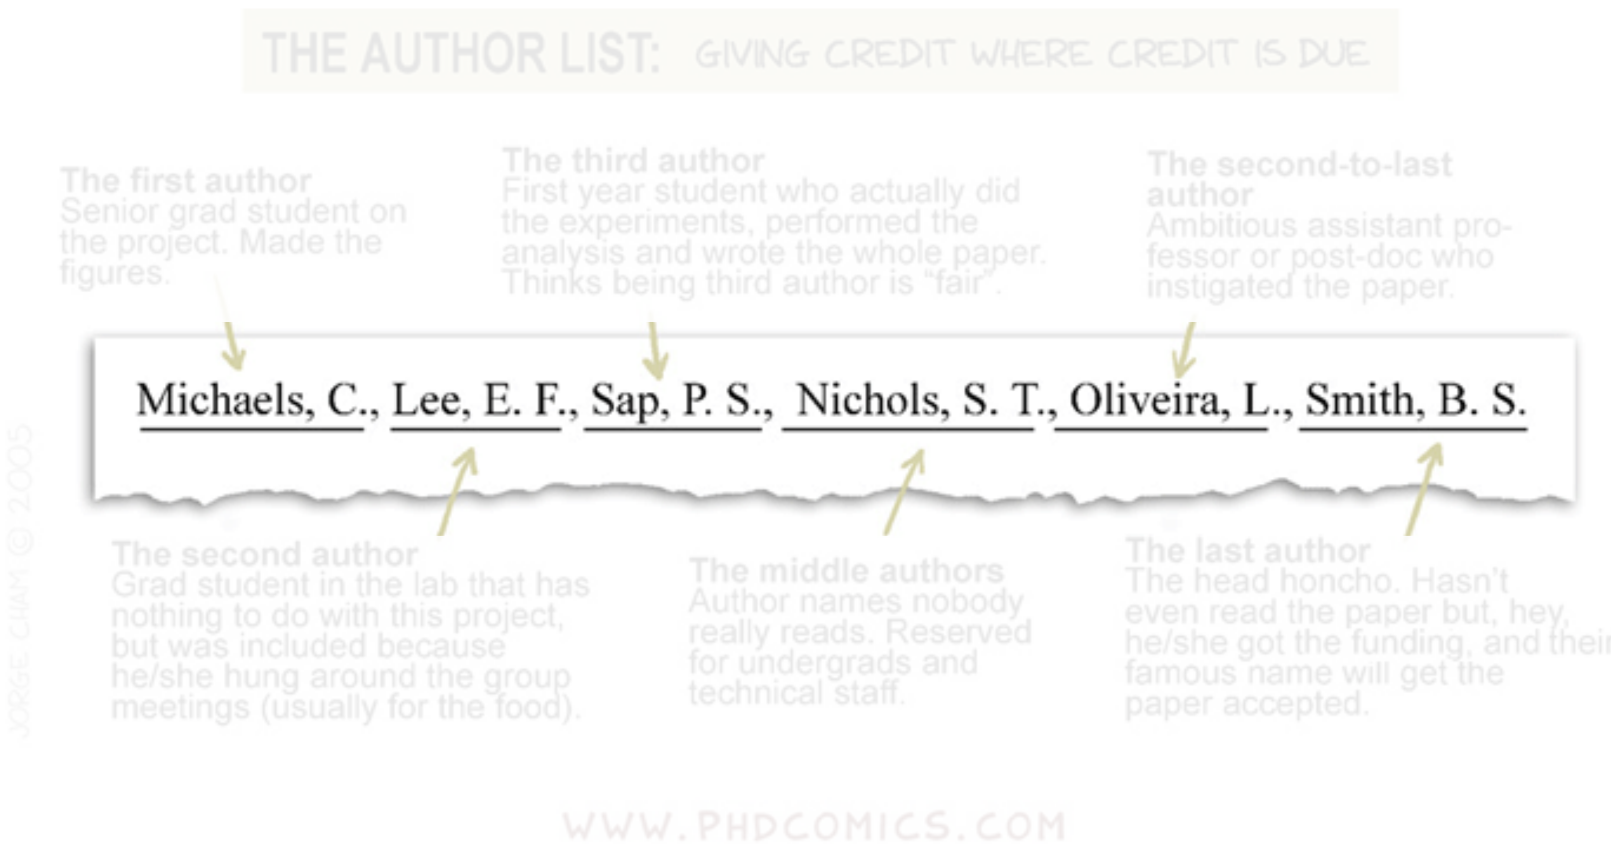




YODAlab

# Acknowledgments

Student Collaborators

Faculty Collaborators

# SPOT Survey Time [15min]

(I'll leave the room for 15 minutes)

Use this link to fill in the SPOT survey:

https://p20.courseval.net/etw/ets/et.asp?nxappid=UA2&nxmid=start