Experimental and causal view on information integration in autonomous agents

Philipp Geiger (MPI for Intelligent Systems),
Katja Hofmann, Bernhard Schölkopf

August 30, 2016

Slides available on `pgeiger.org`

Introduction

Internet of things, self-driving cars, etc.: many trends that increase the number of connected intelligent agents (sensors and actuators)

Introduction

Internet of things, self-driving cars, etc.: many trends that increase the number of connected intelligent agents (sensors and actuators)

Problem statement:

> How can an agent autonomously integrate as much relevant data (or higher level information) as possible from others to inform causal model/ actions?

Introduction

Internet of things, self-driving cars, etc.: many trends that increase the number of connected intelligent agents (sensors and actuators)

Problem statement:

> How can an agent autonomously integrate as much relevant data (or higher level information) as possible from others to inform causal model/ actions?

Examples:

- Road experience transfer between different self-driving cars
- Path descriptions based on landmarks or maps

Previous work

Various approaches to various versions of this problem:

- Reinforcement learning (RL)
- Learning from demonstrations (LfD)
- Transfer learning for agents (TLA)
- Multi-agent systems (MAS)
- Knowledge representation

Previous work

Various approaches to various versions of this problem:

- Reinforcement learning (RL)
- Learning from demonstrations (LfD)
- Transfer learning for agents (TLA)
- Multi-agent systems (MAS)
- Knowledge representation

(Inaccurate? Missing something?)

Our two perspectives on this problem:

1. Simulated experiments – to obtain better understanding
2. Causal models – e.g. for transfer across different agent hardware

Our two perspectives on this problem:

1. Simulated experiments – to obtain better understanding
2. Causal models – e.g. for transfer across different agent hardware

Structure for both:

- introduce toy instance of the problem
- illustrate approach

Problem instance: navigation from video in 'Malmo'

Background: AI experimentation platform 'Malmo': library for programming agents for 'Minecraft' (computer game) [Bignell2016]

Problem instance: navigation from video in ‘Malmo’

Background: AI experimentation platform ‘Malmo’: library for programming agents for ‘Minecraft’ (computer game) [Bignell2016]

Task: unknown landscape; navigate to visually recognizable goal

Problem instance: navigation from video in ‘Malmo’

Background: AI experimentation platform ‘Malmo’: library for programming agents for ‘Minecraft’ (computer game) [Bignell2016]

Task: unknown landscape; navigate to visually recognizable goal

Available heterogeneous information:

- agent’s own sensors (position q, image y) and action (move left/right/forward/backward) at each time t
- “local controller” (past experience on “local physical laws”)
- video $y^*_{0:L}$ of a different (“source”) agent that gets to the goal

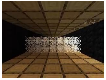

NB: no actions given! – allows e.g. for differing action spaces

A simple integrating agent algorithm

(Given: local controller *ctl*, source agent's video $y^*_{1:L}$)

A simple integrating agent algorithm

(Given: local controller *ctl*, source agent's video $y^*_{1:L}$)

For $i = 1, \ldots, L$

1. Use *ctl* and interaction with environment
   to search locally around position q_{i-1}
   for position q_i with image y most similar to y^*_i

   (formally: $q_i := \arg\min_q \|Gauss * (\overline{y^*_i} - \overline{\mathbb{E}(Y|Q=q)})\|_2$)
2. Use *ctl* to go to q_i

A simple integrating agent algorithm

(Given: local controller *ctl*, source agent's video $y^*_{1:L}$)

For $i = 1, \ldots, L$

1. Use *ctl* and interaction with environment
   to search locally around position q_{i-1}
   for position q_i with image y most similar to y^*_i
   (formally: $q_i := \arg\min_q \| Gauss * (\overline{y^*_i} - \overline{\mathbb{E}(Y|Q=q)})\|_2$)
2. Use *ctl* to go to q_i

Proof-of-concept implementation - evaluation on next slide

- *ctl* := proportional controller based on previous experience
- uses teleportation in search for q_i

Evaluation on "Malmo"

Evaluation on "Malmo"

Source agents trajecory (blue) and integration method (red):

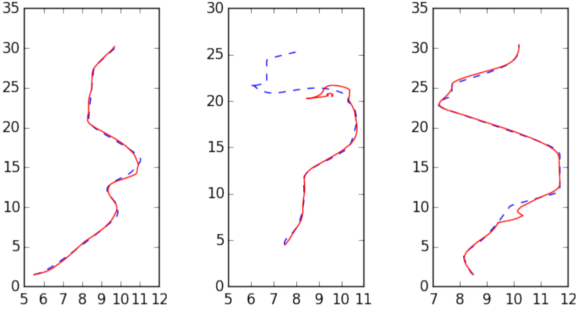

success fail success

Problem instance: experience transfer between cars

Setup: two (or more) self-driving cars with different hardware

Problem instance: experience transfer between cars

Setup: two (or more) self-driving cars with different hardware

Task – w.l.o.g. for car 1: safely follow some trajectory (e.g. road)

Problem instance: experience transfer between cars

Setup: two (or more) self-driving cars with different hardware

Task – w.l.o.g. for car 1: safely follow some trajectory (e.g. road)

Available heterogeneous information:

- hardware specifications of all cars (e.g. table with HP, ...)
- past experiences (actions/observations) of all cars
- influence structure between relevant variables ("causal DAG", see next slide)

Background: causal models & transportability

Def.: diagram (DAG) plus factorizing distribution over set of random variables [Pearl2000]

Background: causal models & transportability

Def.: diagram (DAG) plus factorizing distribution over set of random variables [Pearl2000]

Reason about (identifiability of) **outcomes of manipulations** of the underlying system

Background: causal models & transportability

Def.: diagram (DAG) plus factorizing distribution over set of random variables [Pearl2000]

Reason about (identifiability of) **outcomes of manipulations** of the underlying system

Main example: “X causes Y” := “intervening on X changes Y”

Background: causal models & transportability

Def.: diagram (DAG) plus factorizing distribution over set of random variables [Pearl2000]

Reason about (identifiability of) **outcomes of manipulations** of the underlying system

Main example: "X causes Y" := "intervening on X changes Y"

But useful for reasoning about related systems in general - example:

$$X \longrightarrow Z \longleftarrow Y$$

$\Rightarrow$ $P(z, y|x) = P(z|x)P(y)$

$\Rightarrow$ system $P(z, y|x_1)$ contains information $P(y)$ about modified system $P(z, y|x_2)$ [Pearl2011]

Causal reasoning for toy scenario

$$
\begin{tikzcd}
 & u(t) \arrow[d] & & y(t) \arrow[d] \\
HP \arrow[r] & F(t) \arrow[dr] & & G(t) \arrow[dl] \\
 & & \ddot{y}(t) &
\end{tikzcd}
$$

- y: position
- F: force from engine
- G: other forces (friction etc.)
- HP: horse powers
- u: control signal

Causal reasoning for toy scenario

$$
\begin{tikzcd}
 & u(t) \arrow[d] & & y(t) \arrow[d] \\
HP \arrow[r] & F(t) \arrow[dr] & & G(t) \arrow[dl] \\
 & & \ddot{y}(t) &
\end{tikzcd}
$$

y: position

F: force from engine

G: other forces (friction etc.)

HP: horse powers

u: control signal

1. Assume two cars only differ in $HP = hp_1, hp_2$

Causal reasoning for toy scenario

$$
\begin{tikzcd}
 & u(t) \arrow[d] & y(t) \arrow[d] \\
HP \arrow[r] & F(t) \arrow[dr] & G(t) \arrow[d] \\
 & & \ddot{y}(t)
\end{tikzcd}
$$

y: position

F: force from engine

G: other forces (friction etc.)

HP: horse powers

u: control signal

1. Assume two cars only differ in $HP = hp_1, hp_2$
2. causal DAG $\Rightarrow$ car 2's experience about mechanism $p(G|y)$ transferable to car 1.

Causal reasoning for toy scenario

$$\begin{tikzcd}
 & u(t) \arrow[d] & & y(t) \arrow[d] \\
HP \arrow[r] & F(t) \arrow[dr] & & G(t) \arrow[dl] \\
 & & \ddot{y}(t) &
\end{tikzcd}$$

y: position
F: force from engine
G: other forces (friction etc.)
HP: horse powers
u: control signal

1. Assume two cars only differ in $HP = hp_1, hp_2$
2. causal DAG $\Rightarrow$ car 2's experience about mechanism $p(G|y)$ transferable to car 1.
3. Additivity of y & knowing $p(F|u, hp_1)$ $\Rightarrow$ identify dynamics of car 1

Causal reasoning for toy scenario

$$
\begin{tikzcd}
& u(t) \arrow[d] & & y(t) \arrow[d] \\
HP \arrow[r] & F(t) \arrow[dr] & & G(t) \arrow[dl] \\
& & \ddot{y}(t) &
\end{tikzcd}
$$

y: position

F: force from engine

G: other forces (friction etc.)

HP: horse powers

u: control signal

1. Assume two cars only differ in $HP = hp_1, hp_2$
2. causal DAG $\Rightarrow$ car 2's experience about mechanism $p(G|y)$ transferable to car 1.
3. Additivity of y & knowing $p(F|u, hp_1)$ $\Rightarrow$ identify dynamics of car 1

E.g.: Car 1 avoids slipping on oil spill at position not visited before

Conclusions

Experimental view

Conclusions

Experimental view

- Simple "integrating agent" partially succeeded in toy naviagtion task on "Malmo"

Conclusions

Experimental view

- Simple “integrating agent” partially succeeded in toy naviagtion task on “Malmo”
- Important: take several measuremens then averaging; problem: repetitive structures

Conclusions

Experimental view

- Simple "integrating agent" partially succeeded in toy naviagtion task on "Malmo"
- Important: take several measuremens then averaging; problem: repetitive structures
- NB: Other AI platforms exist, such as "OpenAI Gym"

Conclusions

Experimental view

- Simple "integrating agent" partially succeeded in toy naviagtion task on "Malmo"
- Important: take several measuremens then averaging; problem: repetitive structures
- NB: Other AI platforms exist, such as "OpenAI Gym"

Causal view

Conclusions

Experimental view

- Simple “integrating agent” partially succeeded in toy naviagtion task on “Malmo”
- Important: take several measuremens then averaging; problem: repetitive structures
- NB: Other AI platforms exist, such as “OpenAI Gym”

Causal view

- encode mechanics and reason about transferability

Conclusions

Experimental view

- Simple "integrating agent" partially succeeded in toy naviagtion task on "Malmo"
- Important: take several measuremens then averaging; problem: repetitive structures
- NB: Other AI platforms exist, such as "OpenAI Gym"

Causal view

- encode mechanics and reason about transferability
- Unclear: can this be done by classical say Bayes nets?

Conclusions

Experimental view

- Simple "integrating agent" partially succeeded in toy naviagtion task on "Malmo"
- Important: take several measuremens then averaging; problem: repetitive structures
- NB: Other AI platforms exist, such as "OpenAI Gym"

Causal view

- encode mechanics and reason about transferability
- Unclear: can this be done by classical say Bayes nets?

Future directions

- Use **machine learning** to infer "integration mapping"
- "Universal representation" $\rightsquigarrow$ n instead of n^2 mappings

References

- [Geiger2016] Philipp Geiger, Katja Hofmann, Bernhard Schoelkopf: Experimental and causal view on autonomous information integration in agents.
- [Bignell2016] David Bignell, Katja Hofmann, Tim Hutton, and Matthew Johnson, 'The Malmo platform for artificial intelligence experimentation', in IJCAI, (2016).
- [Pearl2011] Judea Pearl and Elias Bareinboim: Transportability of causal and statistical relations: A formal approach. AAAI 2011.
- [Pearl2000] Judea Pearl: Causality. Cambridge University Press, 2000.
- `https://philippgeiger.org`
- `https://ei.is.tuebingen.mpg.de/research_groups/causal-inference-group`