Intelligent Agents

Knowledge and Seeing

Özgür L. Özçep

Universität zu Lübeck

Institut für Informationssysteme

Todays lecture based on

- The AAMAS 2019 Tutorial **„EPISTEMIC REASONING IN MULTI-AGENT SYSTEMS“, Part 2: Knowledge and Seeing**
  http://people.irisa.fr/Francois.Schwarzentruber/2019AAMAStutorial/

MOTIVATION

The Main Scenario

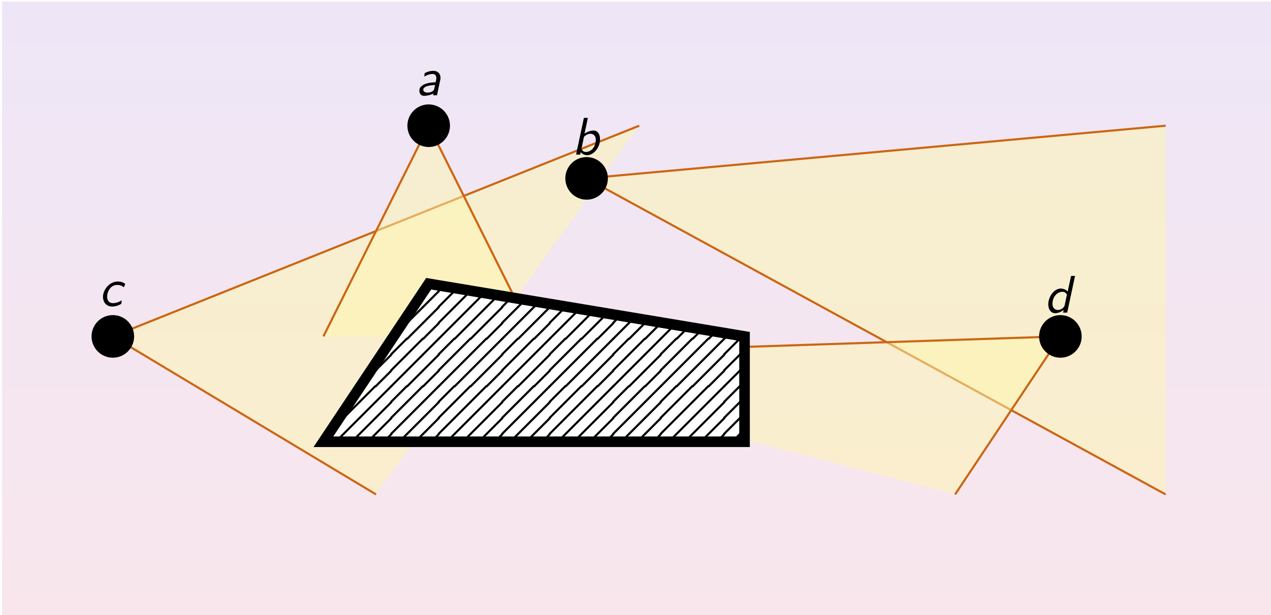

- Agents equipped with vision devices, positioned in the plane / space, e.g. robots that cooperate
- Aim: Represent and compute visual-epistemic reasoning of agents

Spatial reasoning

- Kripke models/epistemic models: abstract notion of possible world and of accessibility

- But agents usally act in space (and time)
  - Should be accounted for
  - The approach discussed here does this within the semantics of specific of form $a\ sees\ b$
  - Leads to ability to express (qualitative9 spatial notions

- Spatial Reasoning and spatial logics (temporal logics, se next lecture next to ime) is a huge topic (see (Aiello et al, 2007))

MODELING

Modeling

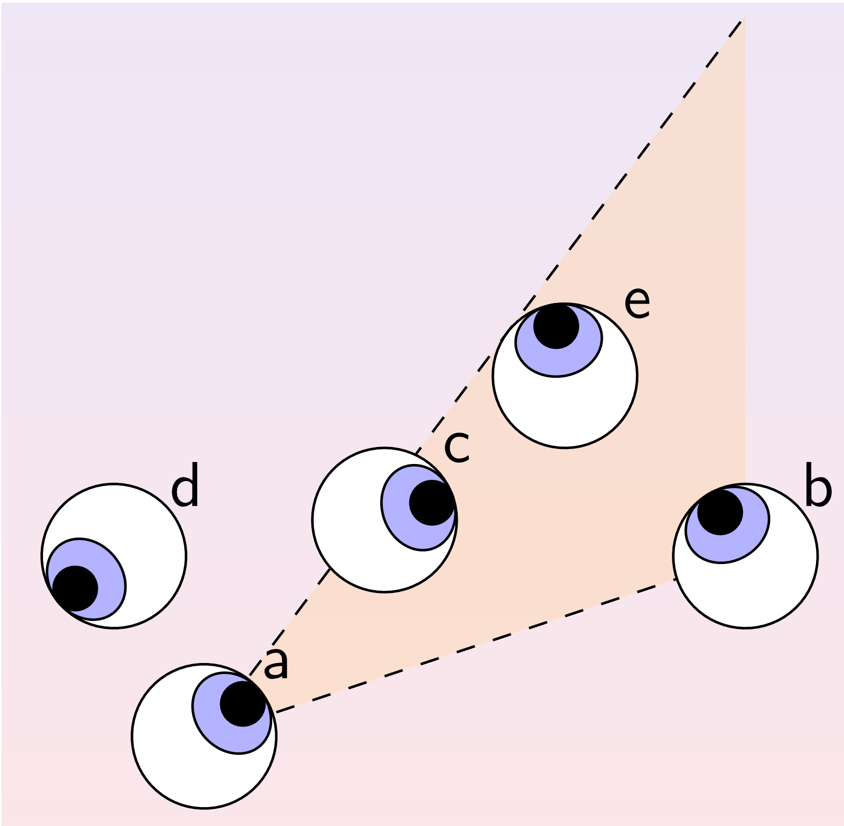

Each agent has a sector (cone) of vision

Assumptions (common knowledge)

- Agents are transparent points in the plane
- All objects of interest are agents
- Agents see infinite sectors
- Angles of vision are the same
- No obstacles (yet)

Possible Worlds

U is set of unit vetors of $\mathbb{R}^2$

Definition

A geometrical possible world is a tuple $w = (pos, dir)$ where:

- $pos: Agt \to \mathbb{R}^2$
- $dir: Agt \to U$

Remember: Agt = set of agents

- $dir(a)$ is the bisector of the sector of vision with angle α
- $C_{p,u,\alpha}$: closed sector with vertex at the point p, angle α and bisector in direction u
- The region seen by a is $C_{pos(a),dir(a),\alpha}$

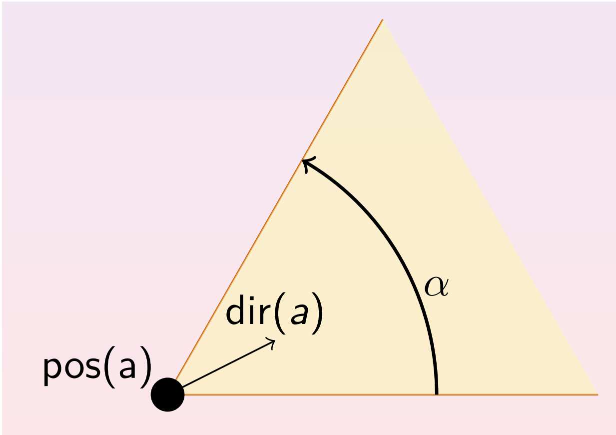

An agent sees another

Definition

a sees b in $w = (pos, dir)$ iff $pos(b) \in C_{pos(a),dir(a),\alpha}$

- a sees a
- a sees b
- a does not see c

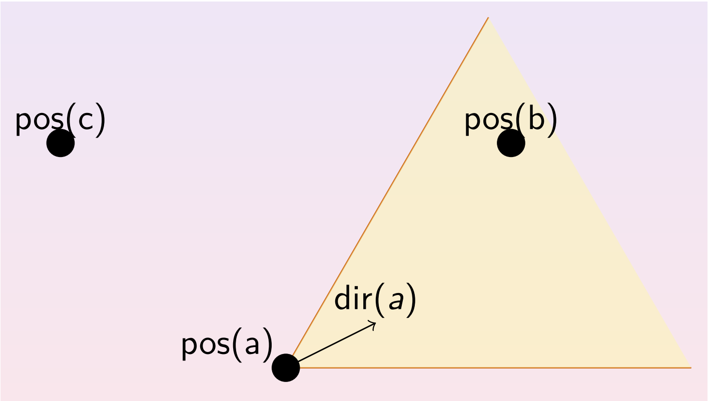

Epistemic model $\mathcal{M}_{flatland}$

Definition

$\mathcal{M}_{flatland} = (W, (\sim_a)_{a \in AGT}, V)$ with

- W is the set of geometrical possible worlds
- $w \sim_a u$ iff agent a sees the same agents in both w and u and these agents have the same position and directions in both w and u

Accessibility relation $\sim_a$ is an equivalence relation. (logic: S5)

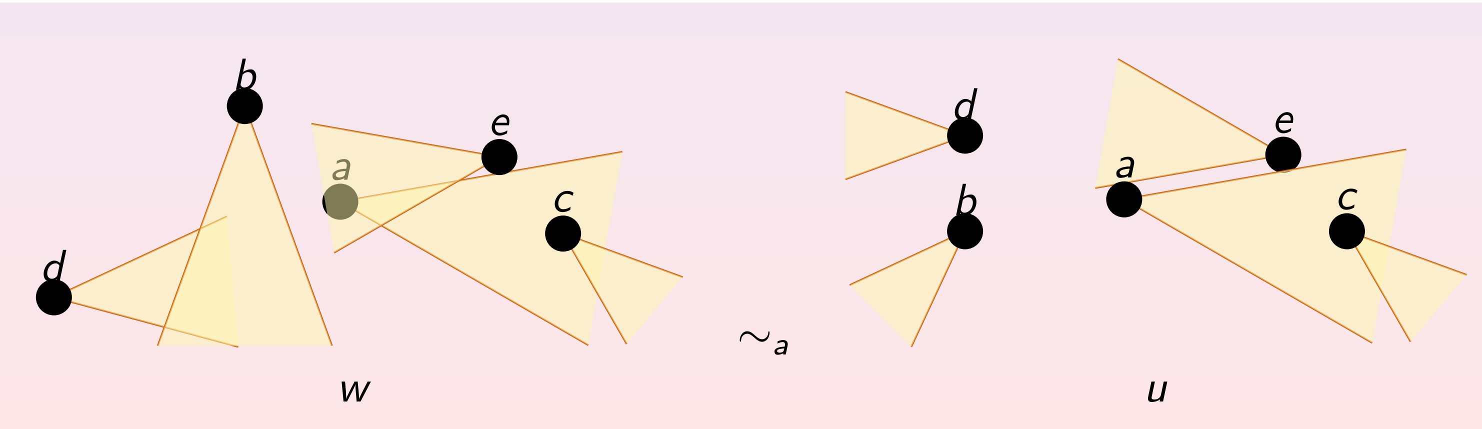

Axiomatization: Disjunctive surprises!

- $\vDash (K_a a\ sees\ b) \lor K_a a\ \sout{sees}\ b$
  (Note that this is not an instance of a tautology)
- $\vDash K_a(\ b\ sees\ c\ \lor d\ sees\ e) \leftrightarrow$
  $$K_a(b\ sees\ c) \lor K_a(d\ sees\ e)$$

Example

- $K_a K_B CK_{c,d,e}(f\ sees\ g)$
  (Note that we use now CK instead of C to denote common knowledge operator)

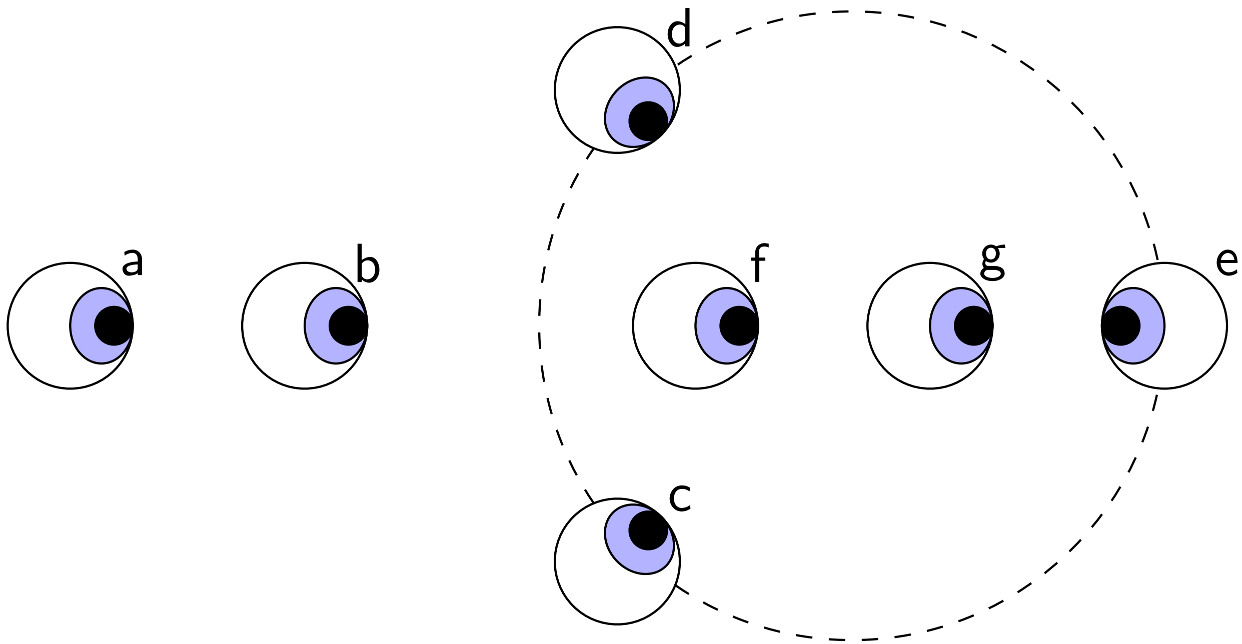

In 1D, only qualitative positions matter

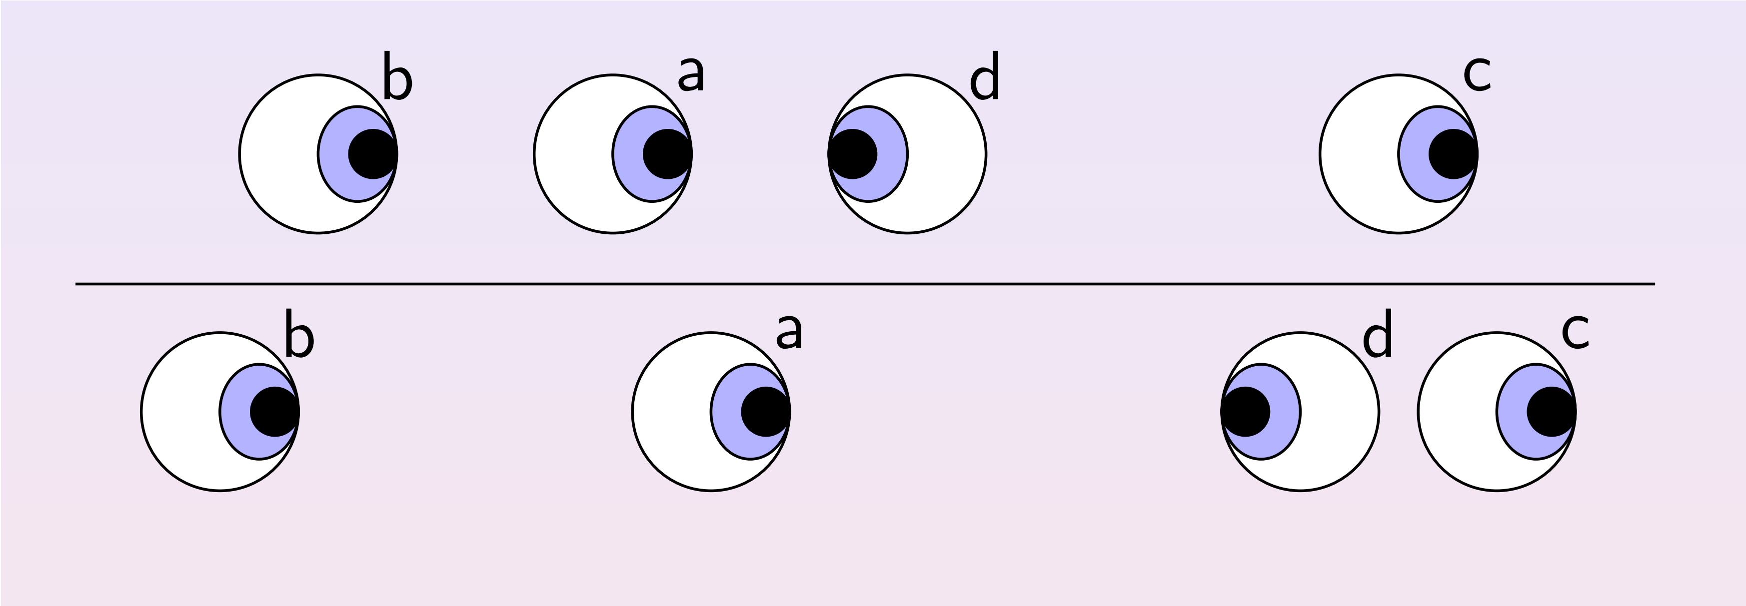

Observation

Qualitative positions are expressible in the language

- $sameDir(a,b) \coloneqq (a\ sees\ b \leftrightarrow b\ \cancel{sees}\ a)$
- $a\ is\ between\ b,c \coloneqq (a\ sees\ b \leftrightarrow a\ \cancel{sees}\ c)$

Abstraction of the Kripke model in 1D

Definition

$abs(w) = \{b \text{ sees } c \mid \mathcal{M}_{robots,1D}, w \vDash b \text{ sees } c\}$ with

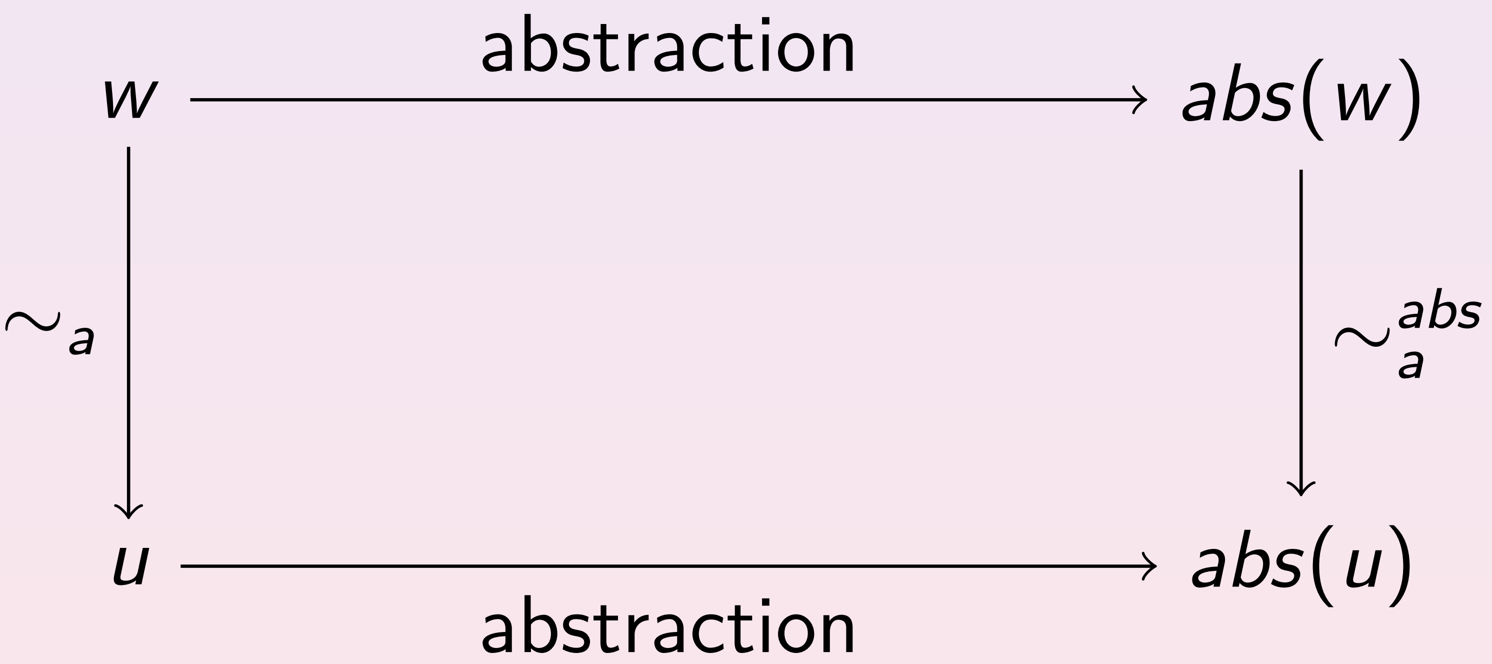

Axiomatization in 1D

- Propositional tautologies;
- $(sameDir(a,b) \leftrightarrow sameDir(b,c)) \rightarrow sameDir(a,c)$;
- $\neg\ (a\ isBetween\ b,c) \lor \neg(b\ isBetween\ a,c)$;
- $(K_a a\ sees\ b) \lor (K_a a\ \cancel{sees}\ b)$
- $a\ sees\ b\ \rightarrow ((K_a b\ sees\ c) \lor (K_a b\ \cancel{sees}\ c))$
- $\chi\ \rightarrow \widehat{K}_a \psi$
  where χ, ψ are complete descriptions with $\chi \sim_a^{abs}\ \psi$
- $K_a \phi \rightarrow \phi$

A complete description is a conjunction that

- contains $a\ sees\ b$ or $a\ \cancel{sees}\ b$ for all agents a, b
- is satisfiable

In 2D, qualitative representation is open issue

Example

$K_b(a \; sees \; b \; \wedge a \; sees \; d \rightarrow a \; sees \; c)$

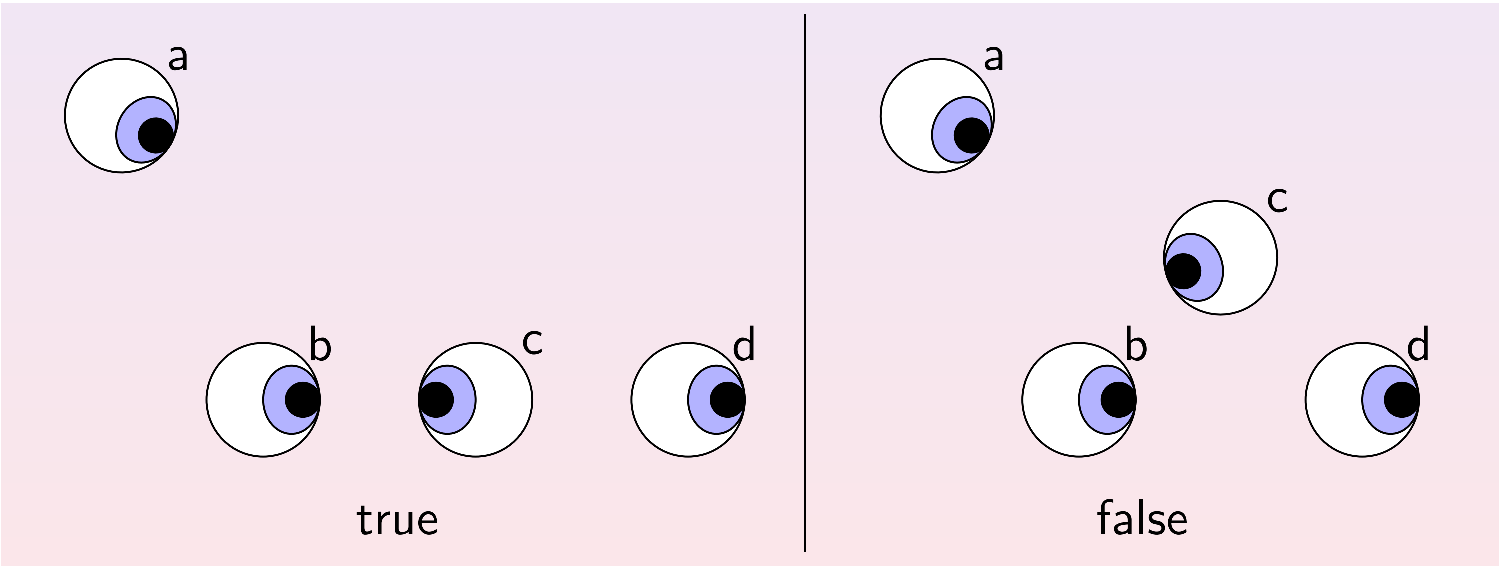

(Assuming here that cone of vision is 1-D: $\alpha = 0$)

Abstraction of the Kripke model in 2D

Definition

$abs(w) = \{b\ sees\ c \mid \mathcal{M}_{robots,2D}, w \vDash b\ sees\ c\}$ with

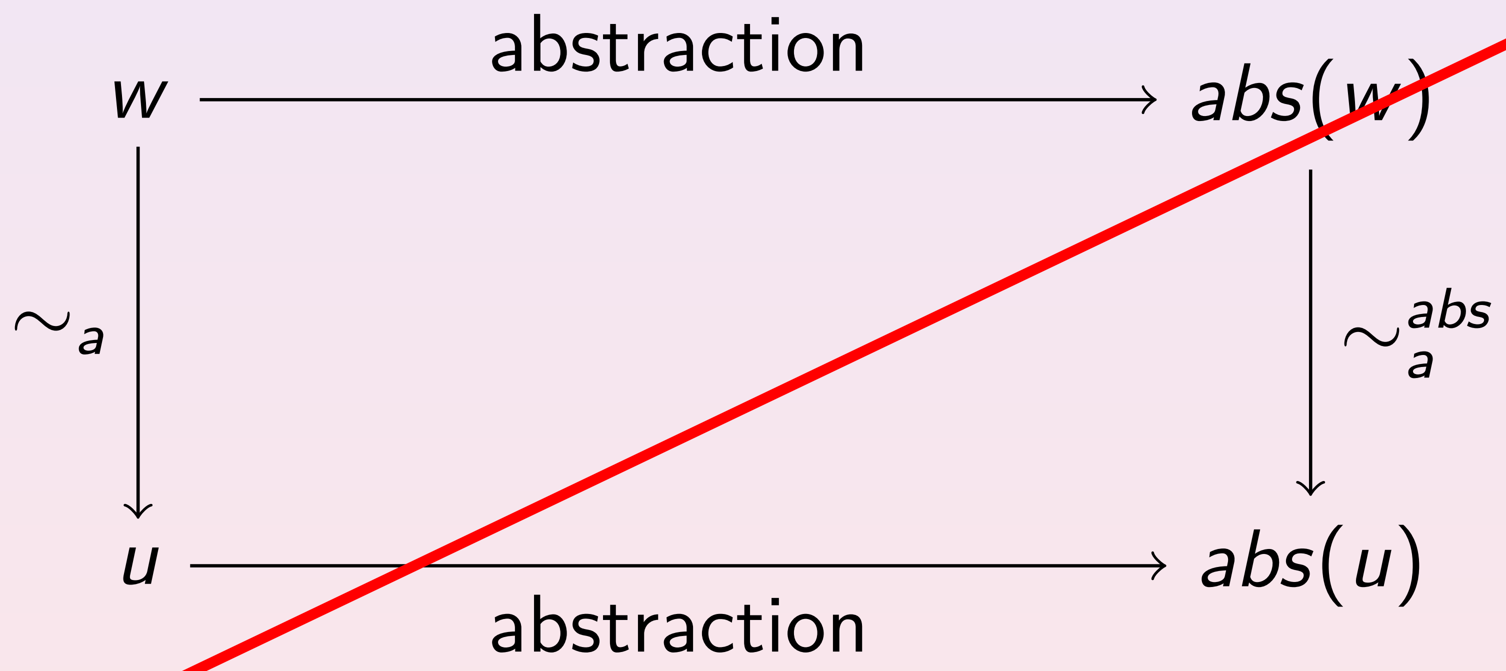

Model checking

- Input
  - A description of a world w (not the whole model)
  - A formula ϕ
- Output: yes iff $w \vDash \phi$

Complexity

Lineland	flatland
PSPACE-complete	PSPACE-hard, and in EXPSPACE (the latter shown by reduction to ℝ-FOL-theory)

ℝ-FOL-theory = elementary algebra : First-oder logic (FOL) of the reals
Language:

- FOL with equality and
- Constants 0, 1
- Functions symbols $+, \times$
- Relation symbols $<$

- Can define, e.g., reals as solutions of polynomials

- Validity of elemantary algebra is known to be in EXSPACE

Complexity

Definition

Standard translation from modal logic to first-order logic

- Atomic propositions p are rewritten to unary predicates P
- $K_a p$ rewritten to $\forall u(R(w,u) \rightarrow P(u))$

(see e.g. Blackburn et al. Modal logic, 2001)

Observation (Adapted translation to $\mathbb{R}$-FOL-theory

$K_a\ (b\ sees\ c)$ rewritten into

$$\forall pos'_a \forall pos'_b \ldots \forall dir'_a \forall dir'_b \ldots$$
$$\{ \bigwedge_{b \in AGT} [(pos_b \in C_{pos(a),dir(a),\alpha}) \rightarrow (pos'_b = pos_b \wedge dir'_b = dir_b)] \wedge$$
$$[(pos_b \notin C_{pos(a),dir(a),\alpha}) \rightarrow (pos'_b \notin C_{pos(a),dir(a),\alpha})]\}$$
$$\rightarrow (pos'_c \notin C_{pos(b),dir(b),\alpha})$$

VARIANT WITH CAMERAS

Agents are cameras

- Cameras
  - Can turn
  - Can NOT move

- Common knowledge
  - Of the positions of agents
  - Of the abilities of perception

Semantics: restricted set of worlds

Definition

Given a fixed $pos' : AGENTS \rightarrow \mathbb{R}^2$, worlds are $w = (pos, dir)$ s.t. $pos = pos'$

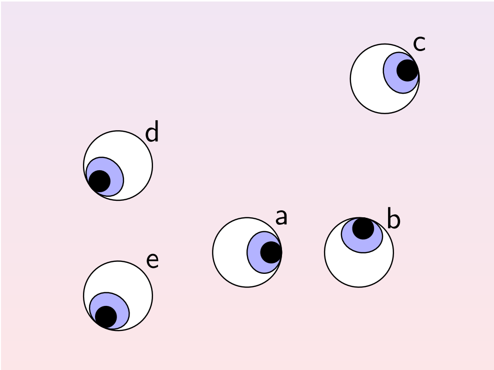

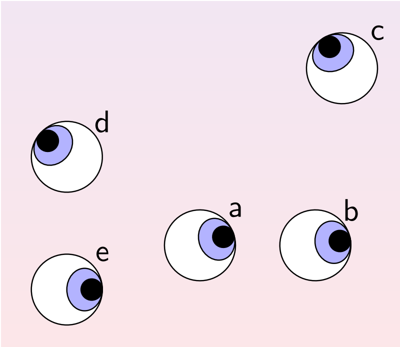

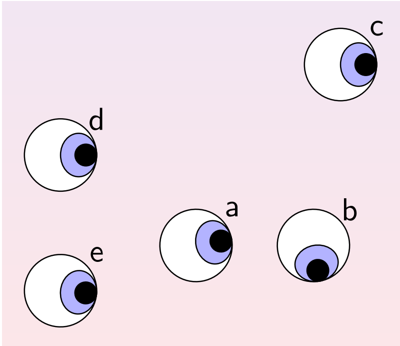

Semantics: $\mathcal{M}_{cameras}$

Definition

$\mathcal{M}_{cameras}$ is $\mathcal{M}_{flatland}$ where we publicly announced the current positions of the agents

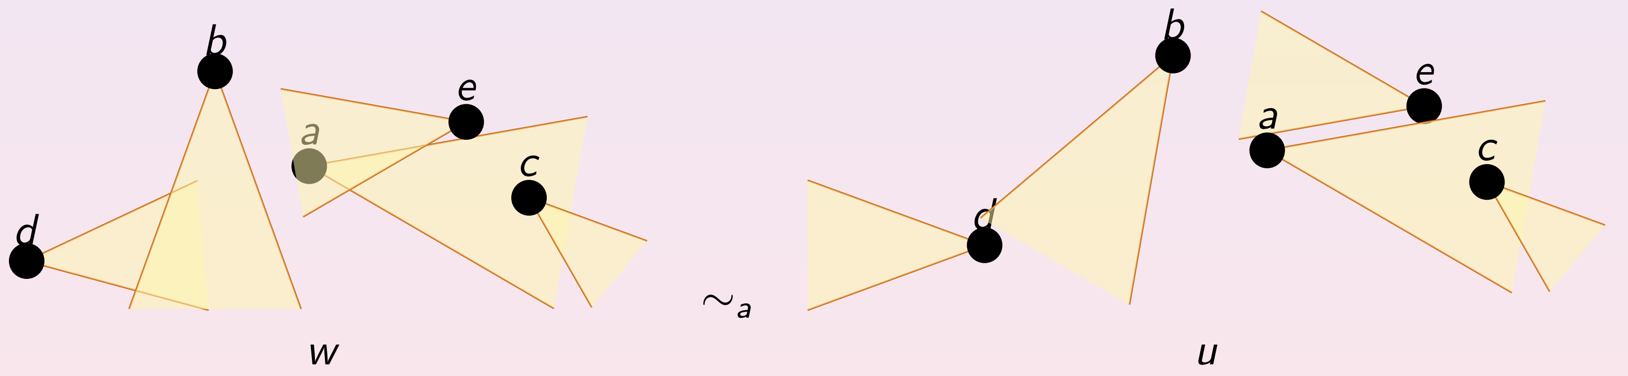

Abstraction of the Kripke model in 2D works

Definition

$abs(w) = \{b \text{ sees } c \mid \mathcal{M}_{cameras,}, w \vDash b \text{ sees } c\}$ with

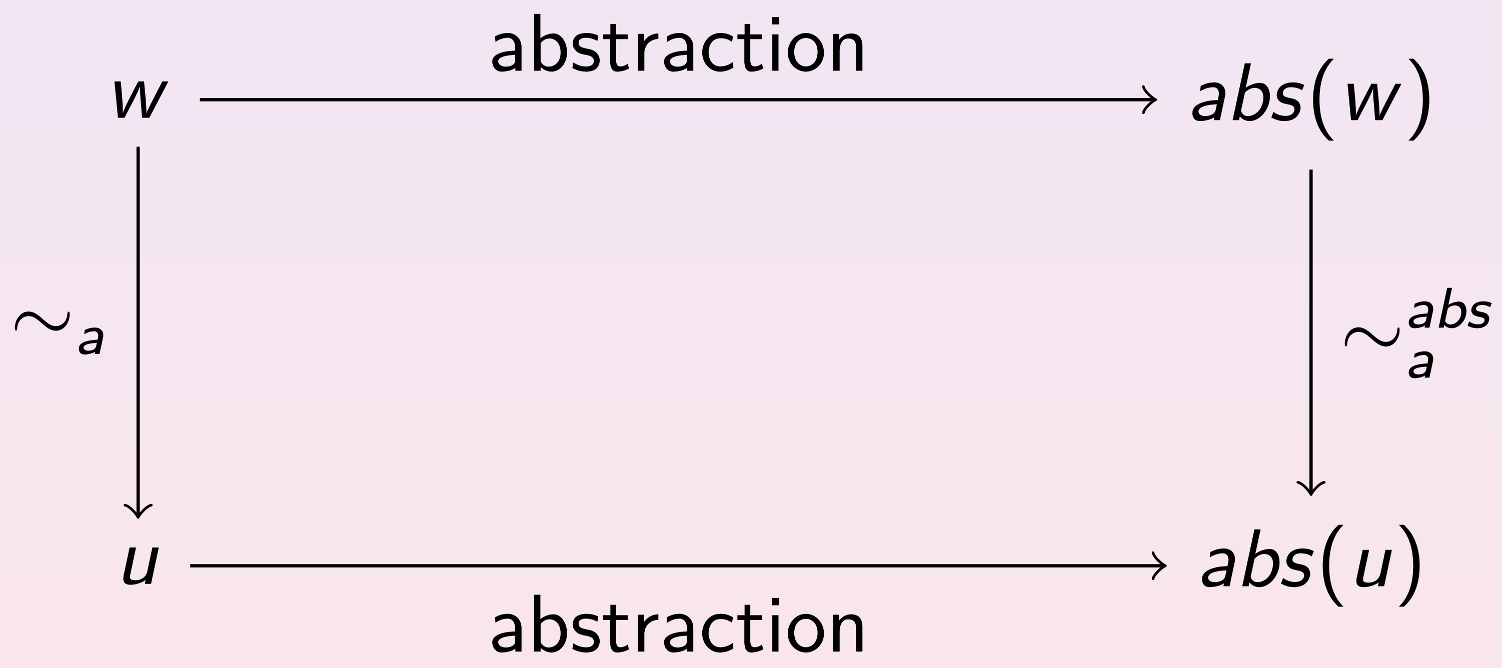

Spectrum of vision

Example (Family of vision sets of agent a)

$$S_a = \{\{b\}, \emptyset, \{c\}, \{d\}, \{d, f\}, \{d, f, e\}, \{f, e\}, \{e\}\}$$

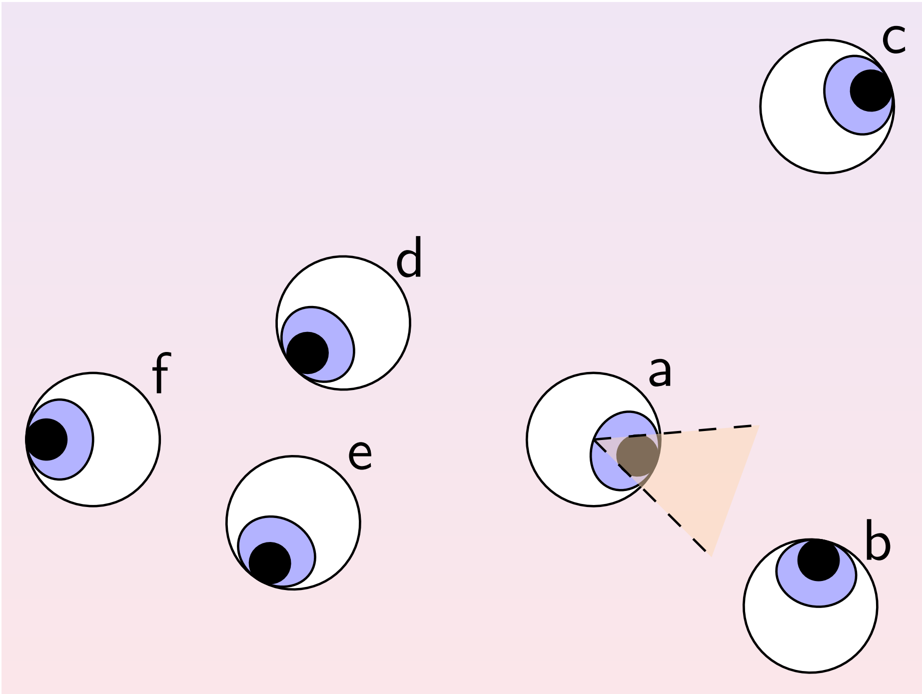

A possible world (dir) can be described as $\{S_a\}_{a \in AGT}$

(configurations on next slides by a moving counterclockwise)

Spectrum of vision

Example (Family of vision sets of agent a)

$$S_a = \{\{b\}, \emptyset, \{c\}, \{d\}, \{d,f\}, \{d,f,e\}, \{f,e\}, \{e\}\}$$

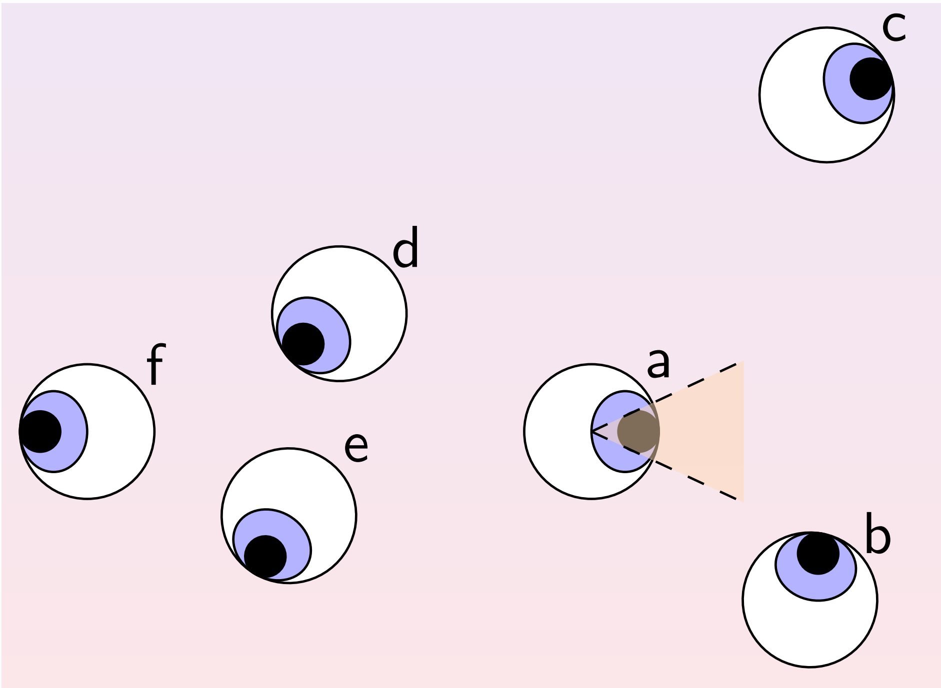

Spectrum of vision

Example (Family of vision sets of agent a)

$$S_a = \{\{b\}, \emptyset, \{c\}, \{d\}, \{d, f\}, \{d, f, e\}, \{f, e\}, \{e\}\}$$

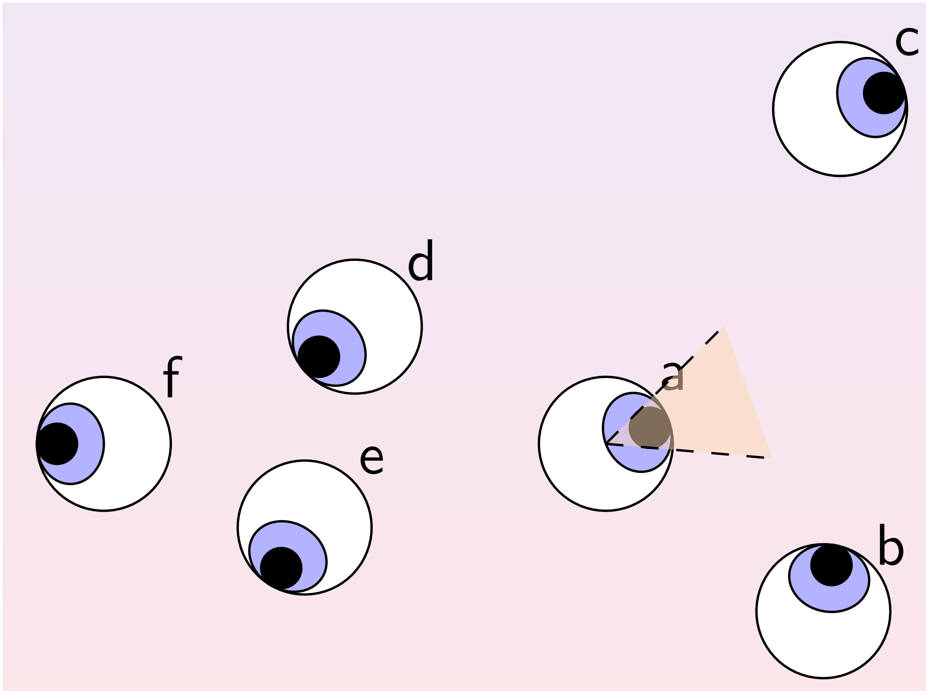

Spectrum of vision

Example (Family of vision sets of agent a)

$$S_a = \{\{b\}, \emptyset, \{c\}, \{d\}, \{d, f\}, \{d, f, e\}, \{f, e\}, \{e\}\}$$

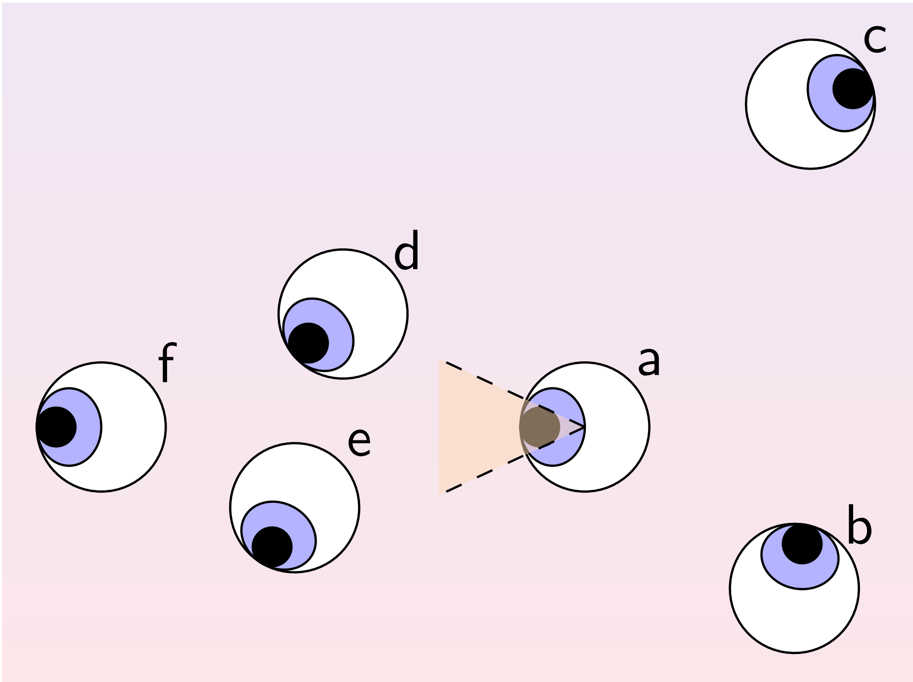

NB:

- each S_a is computed in $O(k \log k)$ steps, where $k = \#(Agt)$.

PDL (Propositional Dynamic Logic)

Definition (PDL Syntax)

$$\phi ::= a\ sees\ b \mid \neg\phi \mid \phi \vee \psi \mid [\pi]\phi$$

- Intended semantics for $[\pi]\phi$: after all executions of program π, it holds that ϕ

PDL Language

Definition (Syntax of programs)

$$\pi_{\ldots} ::= a^{\frown} \mid \phi? \mid \pi;\pi' \mid \pi \cup \pi' \mid \pi^*$$

- Intended semantics for
  - $a^{\frown}$: a turns;
  - $\phi?$: the program succeeds when ϕ is true
  - $\pi;\pi'$: π followed by π'
  - $\pi \cup \pi'$: non-deterministically execute π or π'
  - π^*: repeat π a finite, but non-deterministically, number of times

Translating epistemic operators in programs

- K_a is simulated by

$$\underbrace{\left[(a\ sees\ b_1?) \cup (a\ \cancel{sees}\ b_1?\,;\ b_1^{\frown})\right);\ \ldots;\ (a\ sees\ b_n?) \cup (a\ \cancel{sees}\ b_n?\,;\ b_n^{\frown}))]}_{\pi_a}$$

(b_i = all agents except for a)

- Each component program $\left[(a\ sees\ b_i?) \cup (a\ \cancel{sees}\ b_i?\,;\ b_i^{\frown})\right)$ says: can turn view of b_i iff a does not see b_i
- Thus the program may change arbitrarily all agents, other than a, that a cannot see
- And this is exactly the semantics of K_a

Model checking

Observation

Model checking of PDL for cameras is PSPACE-complete

(Gasquet et al. 2014)

Summary: Visual-epistemic reasoning of agents

- Epistemic language involving atomic propositions 'a $sees$ b'.
- Semantics in geometric and Kripke models.
- 1D case and 2D case with cameras (spectrum of vision):
  - Finite abstraction in the 1D case and in the 2D case with cameras (spectrum of vision).
  - Optimal PSPACE model checking.
- Open problem for the full 2D case: finite abstraction?

Future work

- Obstacles (occlusion)
- Moving agents/cameras in the plane: mathematically more complex; finite abstractions may not work
- Agents/cameras in the 3D space

Uhhh, a lecture with a hoepfully useful

APPENDIX

References

- M. Aiello, I. Pratt-Hartmann, and J. Benthem, editors. Handbook of Spatial Logics. Springer, 2007.
- P. Balbiani, O. Gasquet, and F. Schwarzentruber. Agents that look at one another. Logic Journal fo the IGPL, 21(3):438–467, 12 2012.
- P. Blackburn, M. de Rijke, and Y. Venema. Modal Logic, volume 53 of Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2. edition, 2002.
- O. Gasquet, V. Goranko, and F. Schwarzentruber. Big brother logic: logical modeling and reasoning about agents equipped with surveillance cameras in the plane. In Proceedings of AAMAS ’14, 2014.

Color Convention in this course

- Formulae, when occurring inline
- Newly introduced terminology and definitions
- Important results (observations, theorems) as well as emphasizing some aspects
- Examples are given with standard orange with possibly light orange frame
- Comments and notes
- Algorithms