Universität Bamberg

Faithful Geometric Models for Integrating Learning and Reasoning

Diedrich Wolter
diedrich.wolter@uni-bamberg.de
University of Bamberg • Germany

Universität Bamberg

Faithful Geometric Models for Integrating Learning and Reasoning

Acknowledgement:
Özgür L. Özçep & Mena Leemhuis
University of Lübeck • Germany

Diedrich Wolter
diedrich.wolter@uni-bamberg.de
University of Bamberg • Germany

Overview

- **Knowledge Graph Embeddings**
  - melting pot of symbolic AI and ML
  - fully expressive models are challenging
- **Geometric-logic commitments**
- **Cone-Based Geometric Models**
  - faithful models that capture uncertainty

L E A R N I N G
+ R E A S O N I N G
E M B E D D I N G

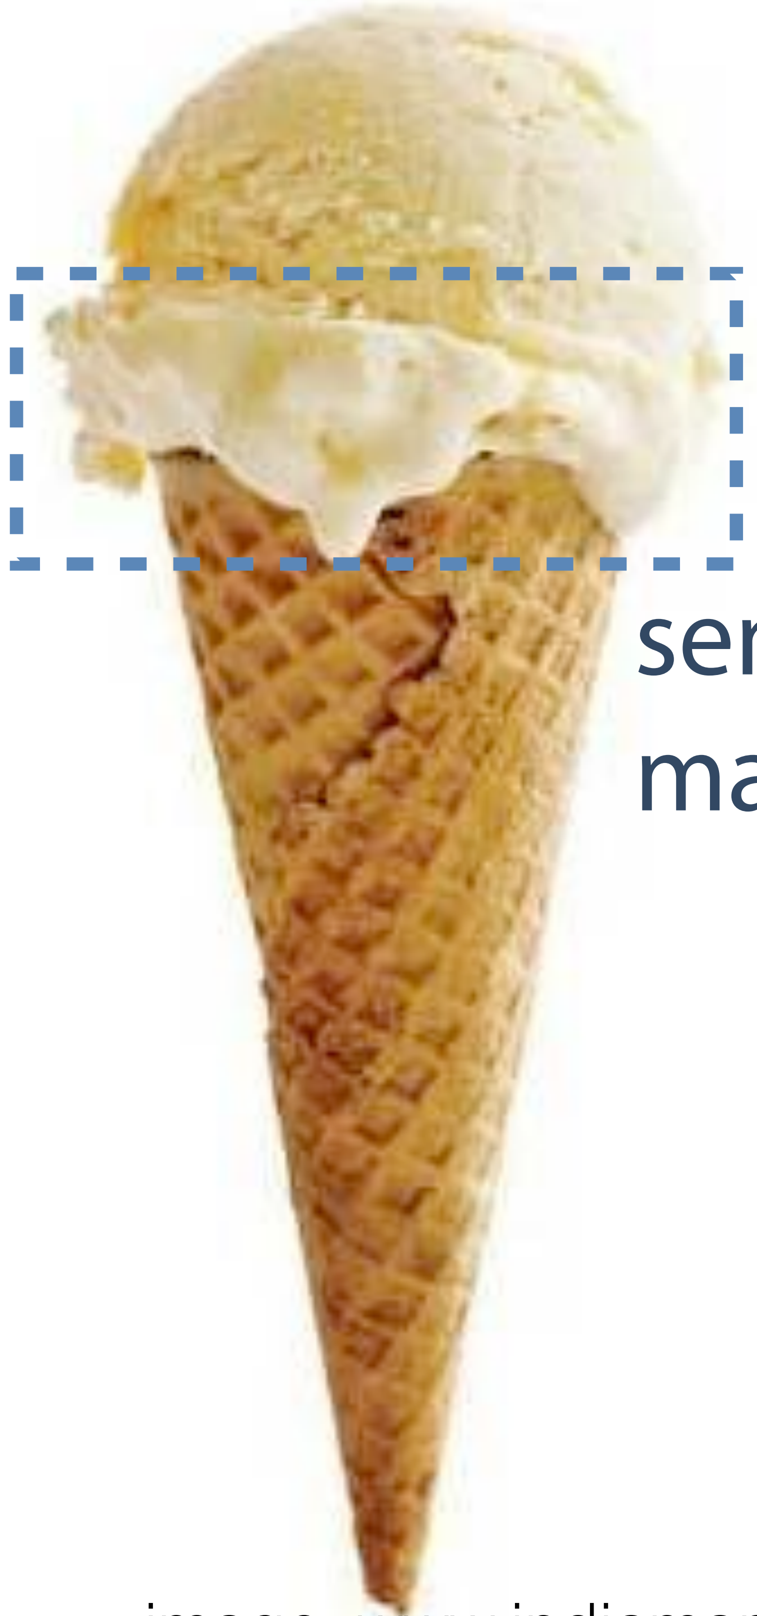

semantic alignment matters

image: www.indiamart.com

Learning + Reasoning

- motivating hypothesis: neither learning nor reasoning alone sufficient to master challenging tasks
  - reasoning lacks data to operate on – knowledge engineering bottleneck
  - learning lacks unbiased calculus

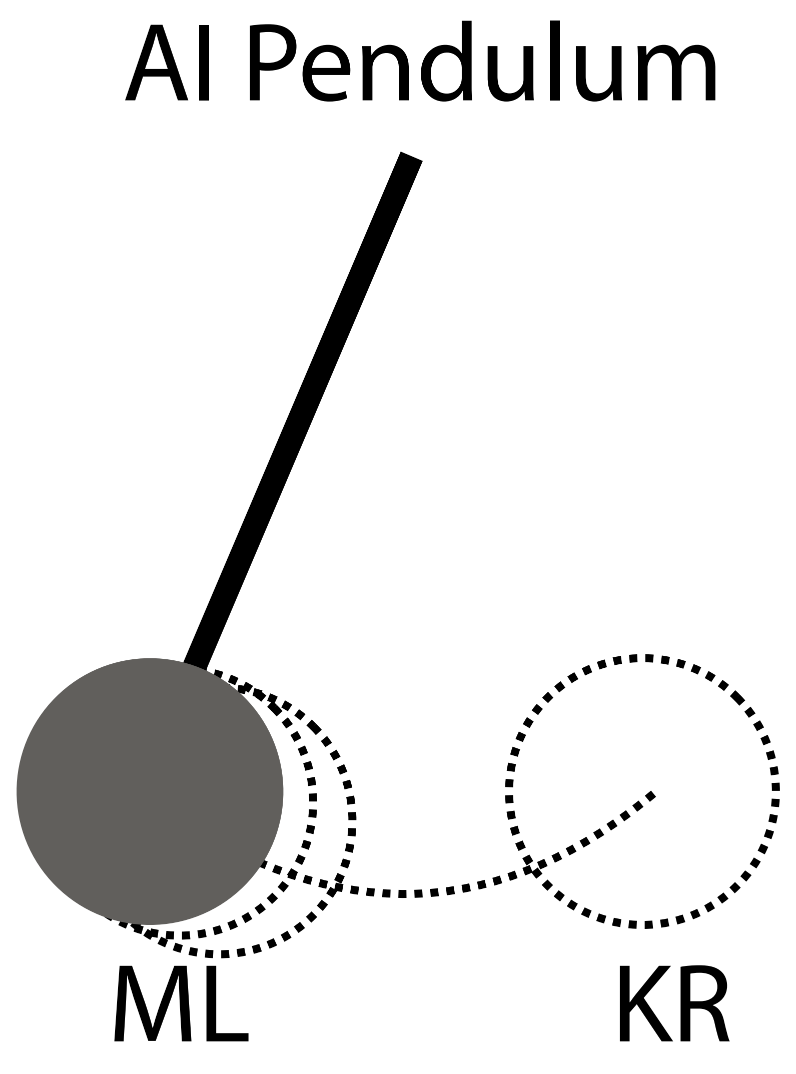

- active field of **neuro-symbolic AI (hybrid AI)**
  - **Logic Tensor Networks** (Serafini & d'Avila Garcez, 2016; Badreddine et al., 2022)
  - **Logical Neural Networks** (Riegel et al. 2020, Sen et al. 2022)
  - **logic-based knowledge graph embeddings** (e.g., Gutiérres-Basulto & Schockaert 2018, Kulmanov et al. 2019)

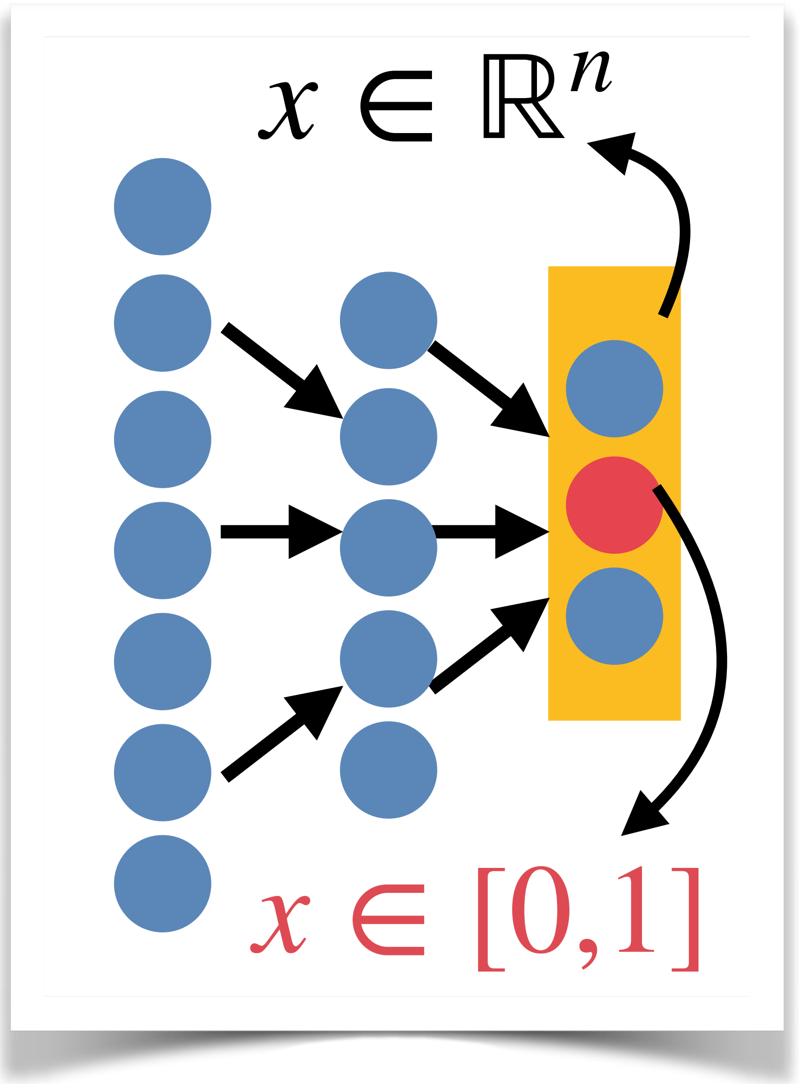

Knowledge Graph (KG)

- graph-like representation of knowledge
  - vertices represent entities
  - edges represent binary relations
  - Boolean validity
- large-scale databases, so-called **triplestore**: (subj rel obj)
  - semantic queries, e.g., using Wikipedia's Wikidata
  - application example: question answering
- KGs are often incomplete
  - **link prediction** as widely considered task:
    predict validity of unseen triple (x r y) from seen triples
  - **example:** (robin eats worm), (seagull eats worm)
    ⇝ (blackbird eats worm)

KG for question answering extracted from Wikipedia:
"Who invented Coca-Cola?"

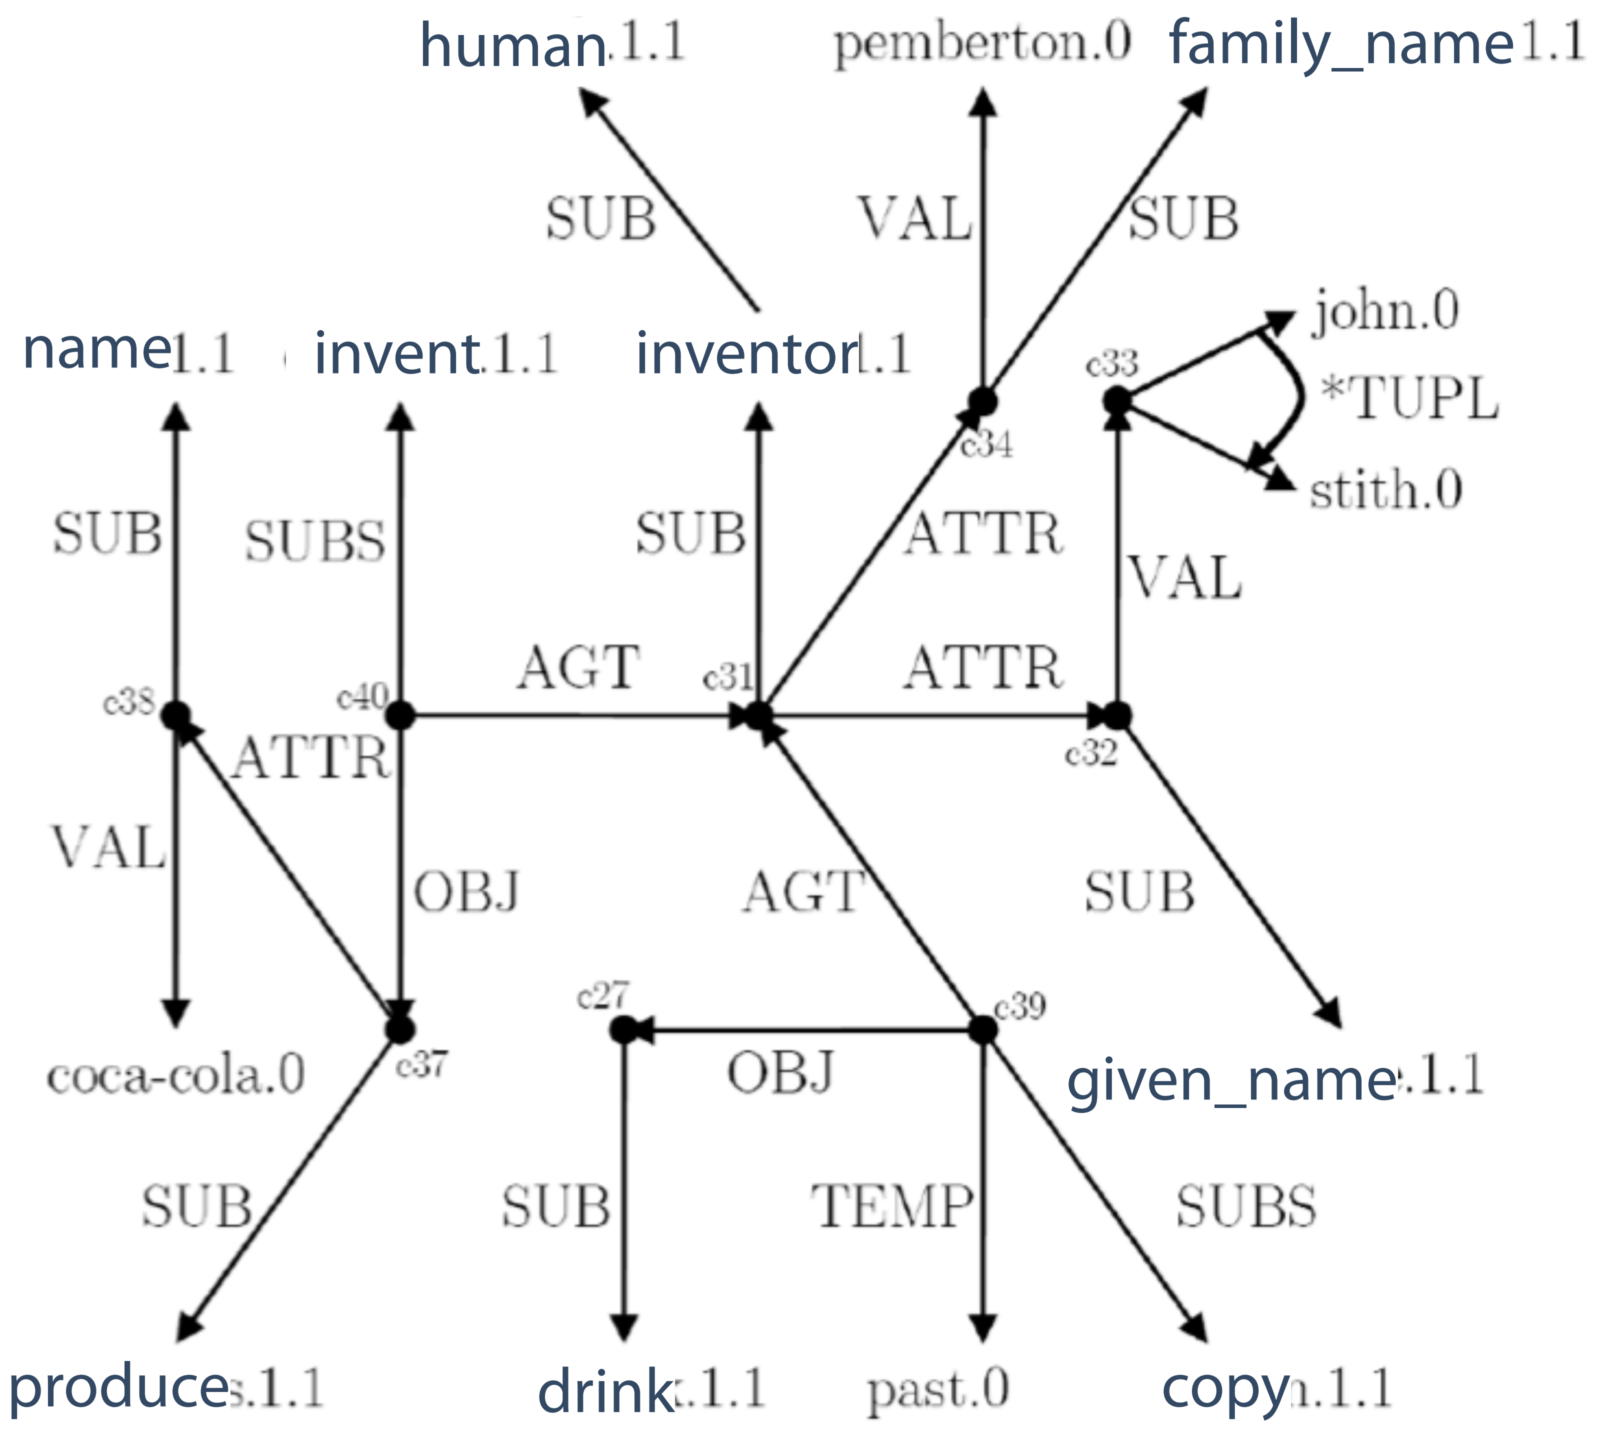

(Furbach et al., 2010)

Knowledge Graph Embeddings (KGEs)

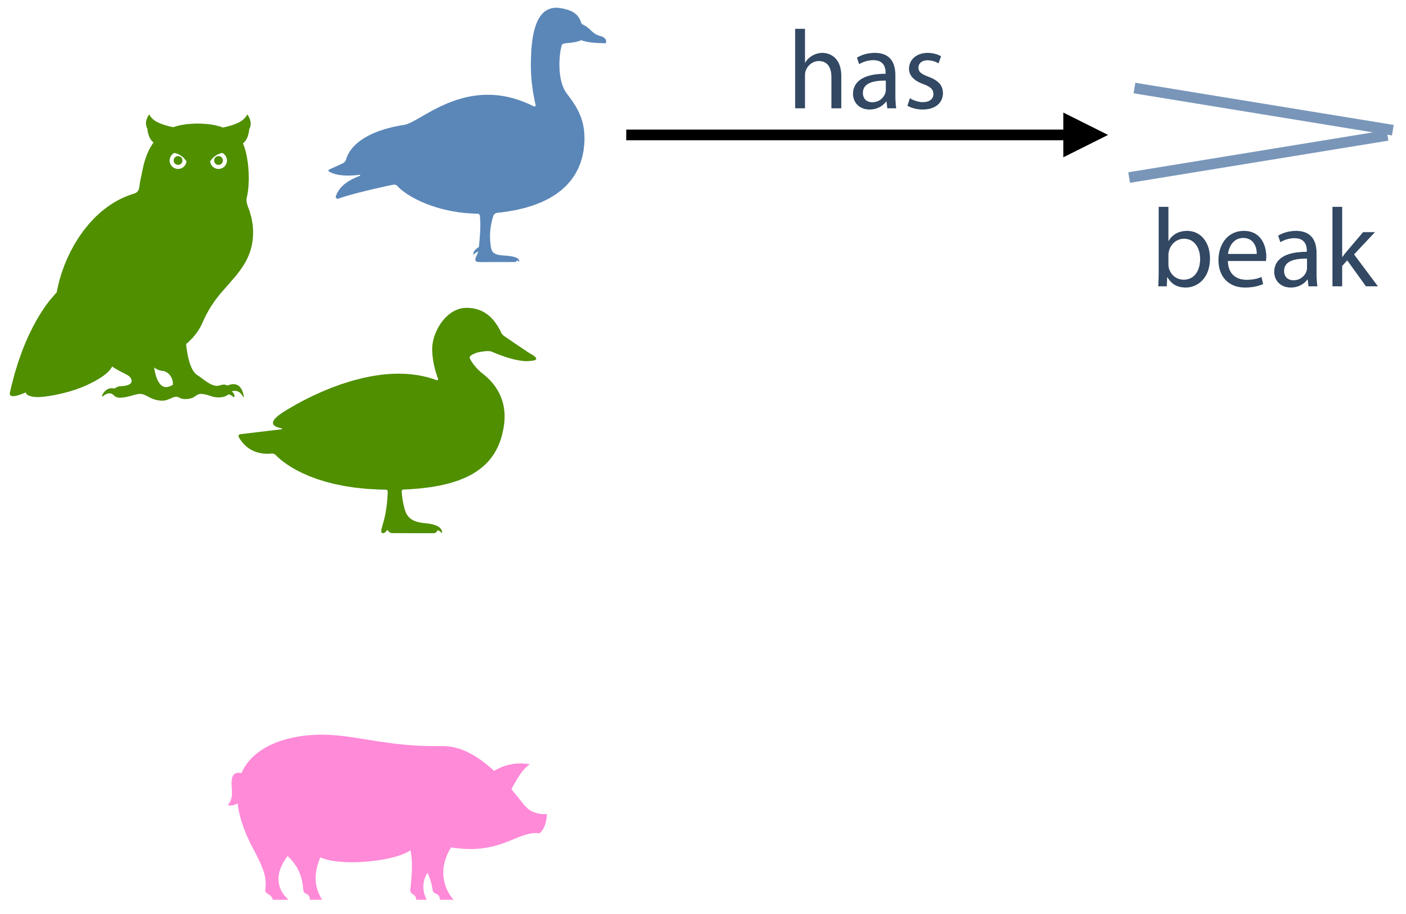

- **idea**: link prediction achievable by exploiting similarities in geometrical space
  - learn geometric representation reflecting data
- KGE comprises
  - **vector representation** x^I of object x in $\mathbb{R}^n$
    - $(\cdot)^I$ interprets an object geometrically
  - **scoring function** $s_r : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ for relations r
    - if $s_r(x^I, y^I)$ small, $r(x, y)$ is assumed to hold
    - scoring function induces geometry of relations
- variety of approaches: how interpretation is learned, how scoring functions are defined

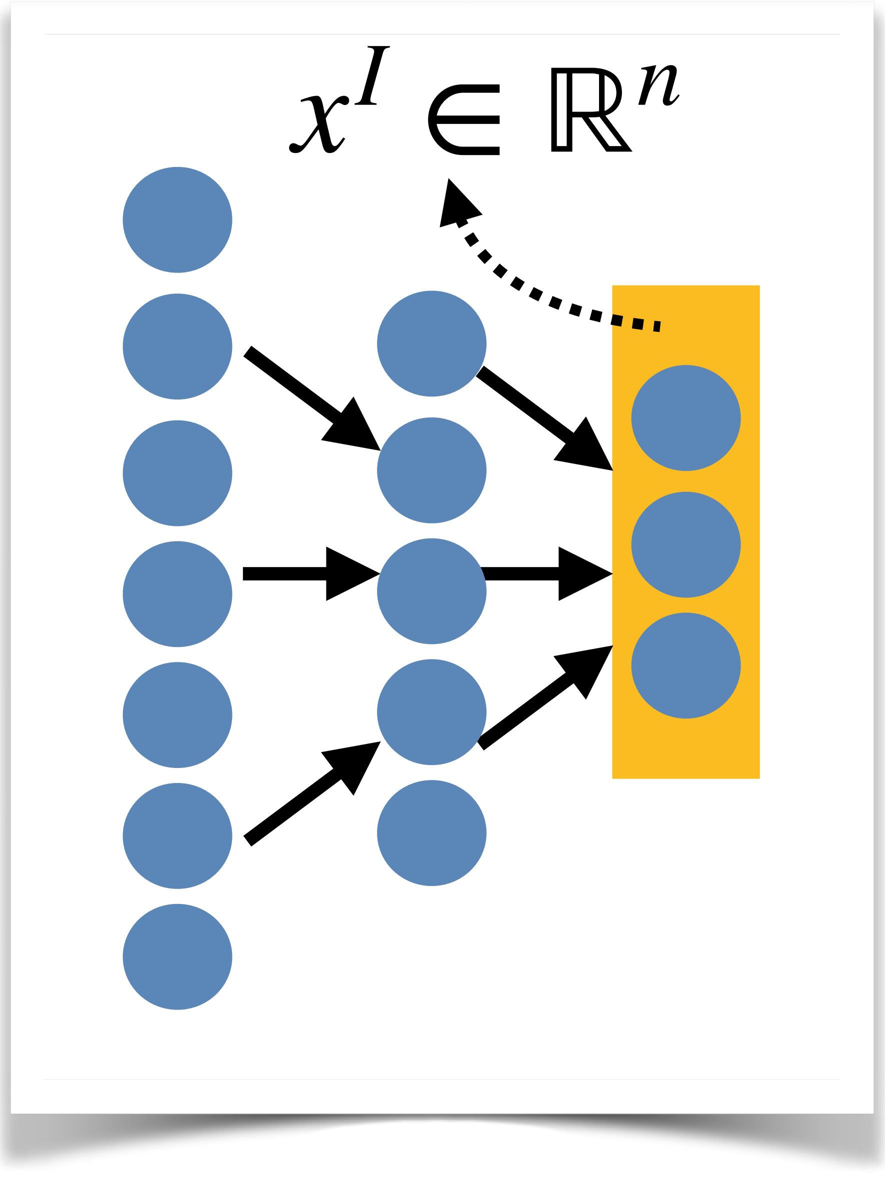

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole ’18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible

▶ **Can we obtain an exact logic characterisation?**
  - …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

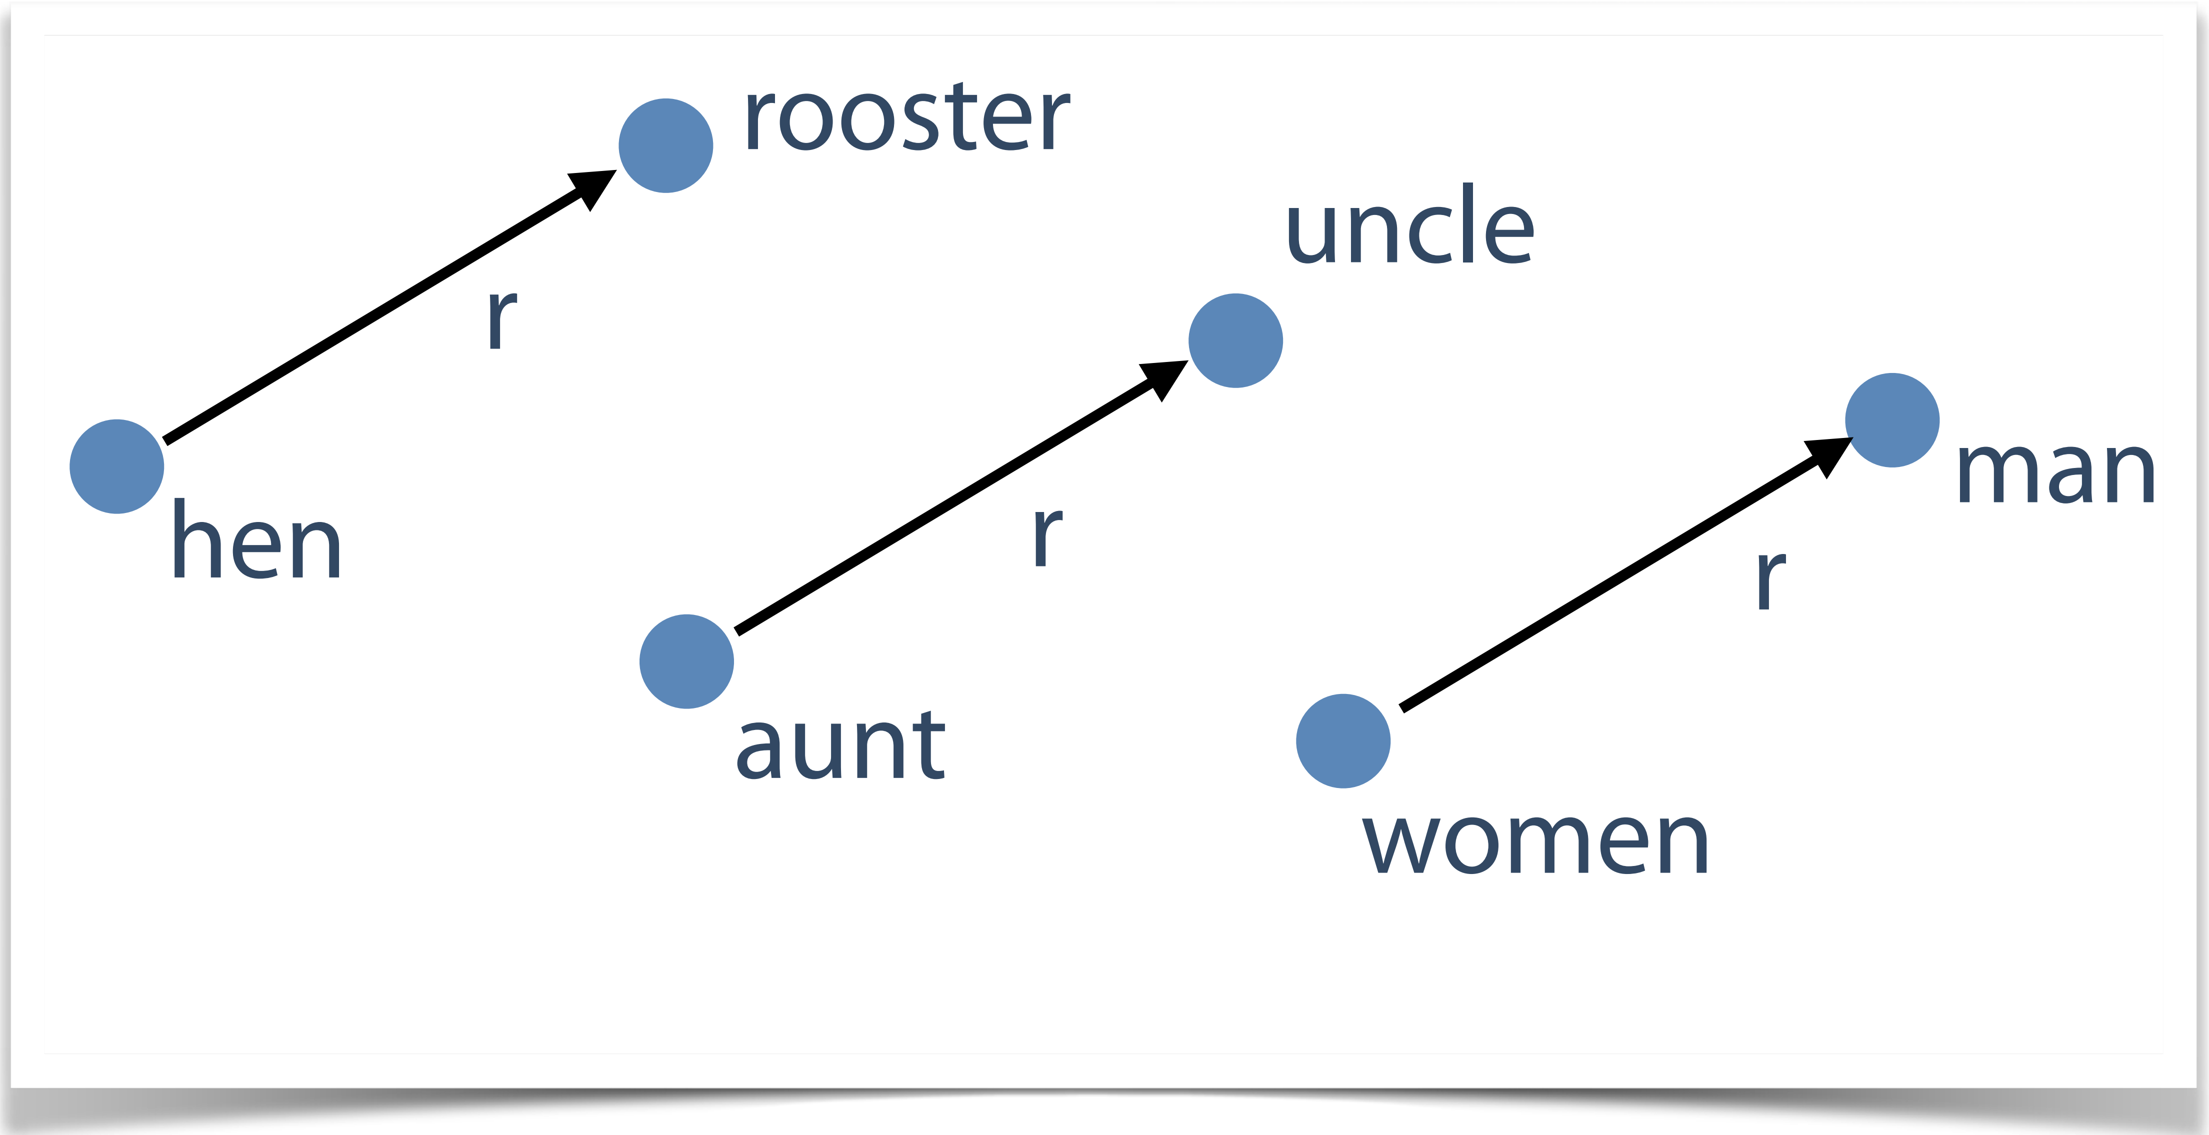

$$s_r(x, y) := ||x + r - y||_2$$

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole '18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible

▶ **Can we obtain an exact logic characterisation?**

- …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

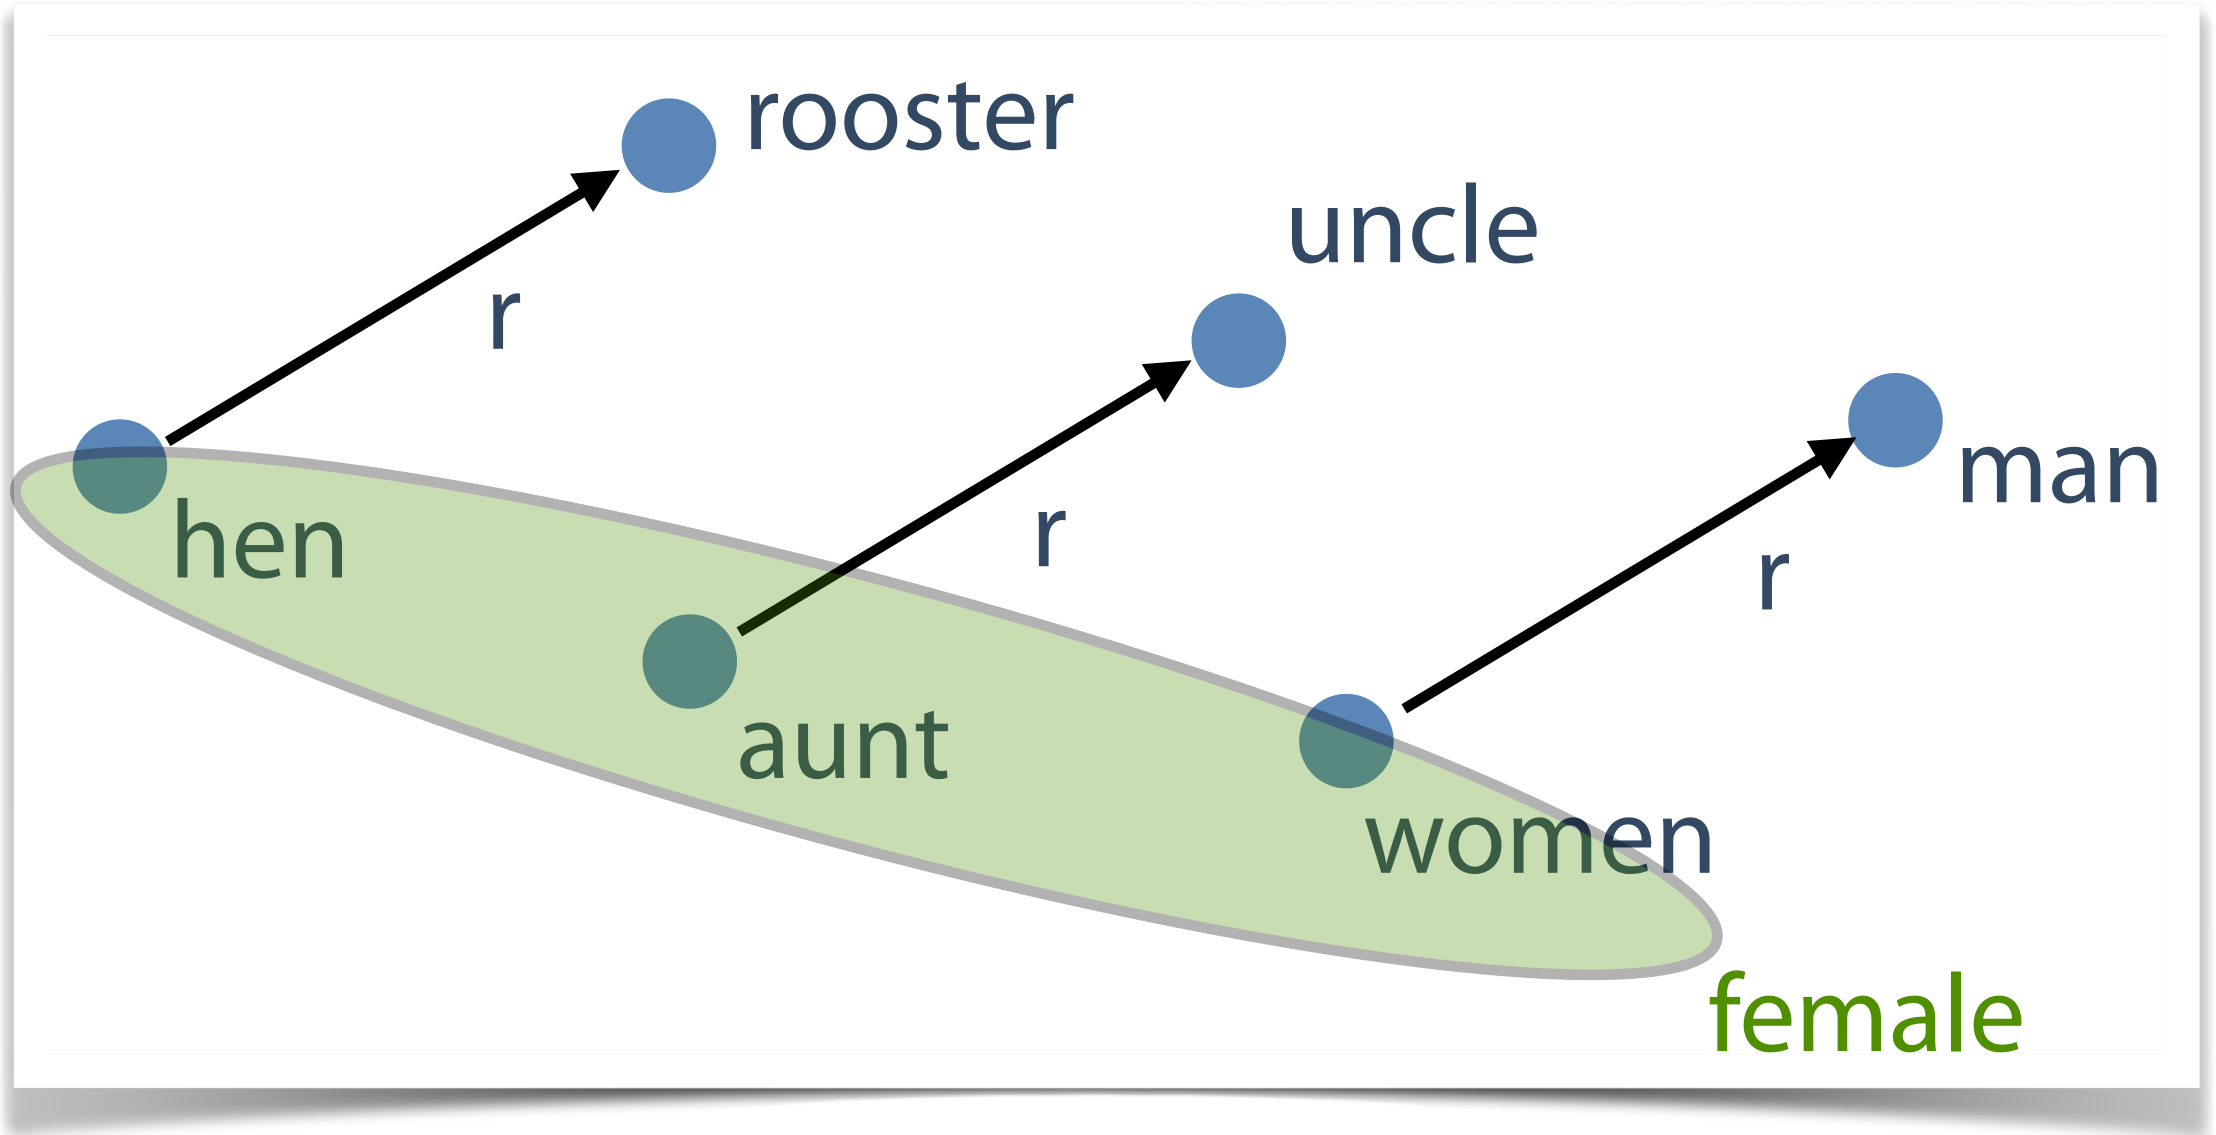

$$s_r(x, y) := ||x + r - y||_2$$

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole '18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible
- ▶ **Can we obtain an exact logic characterisation?**
  - …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

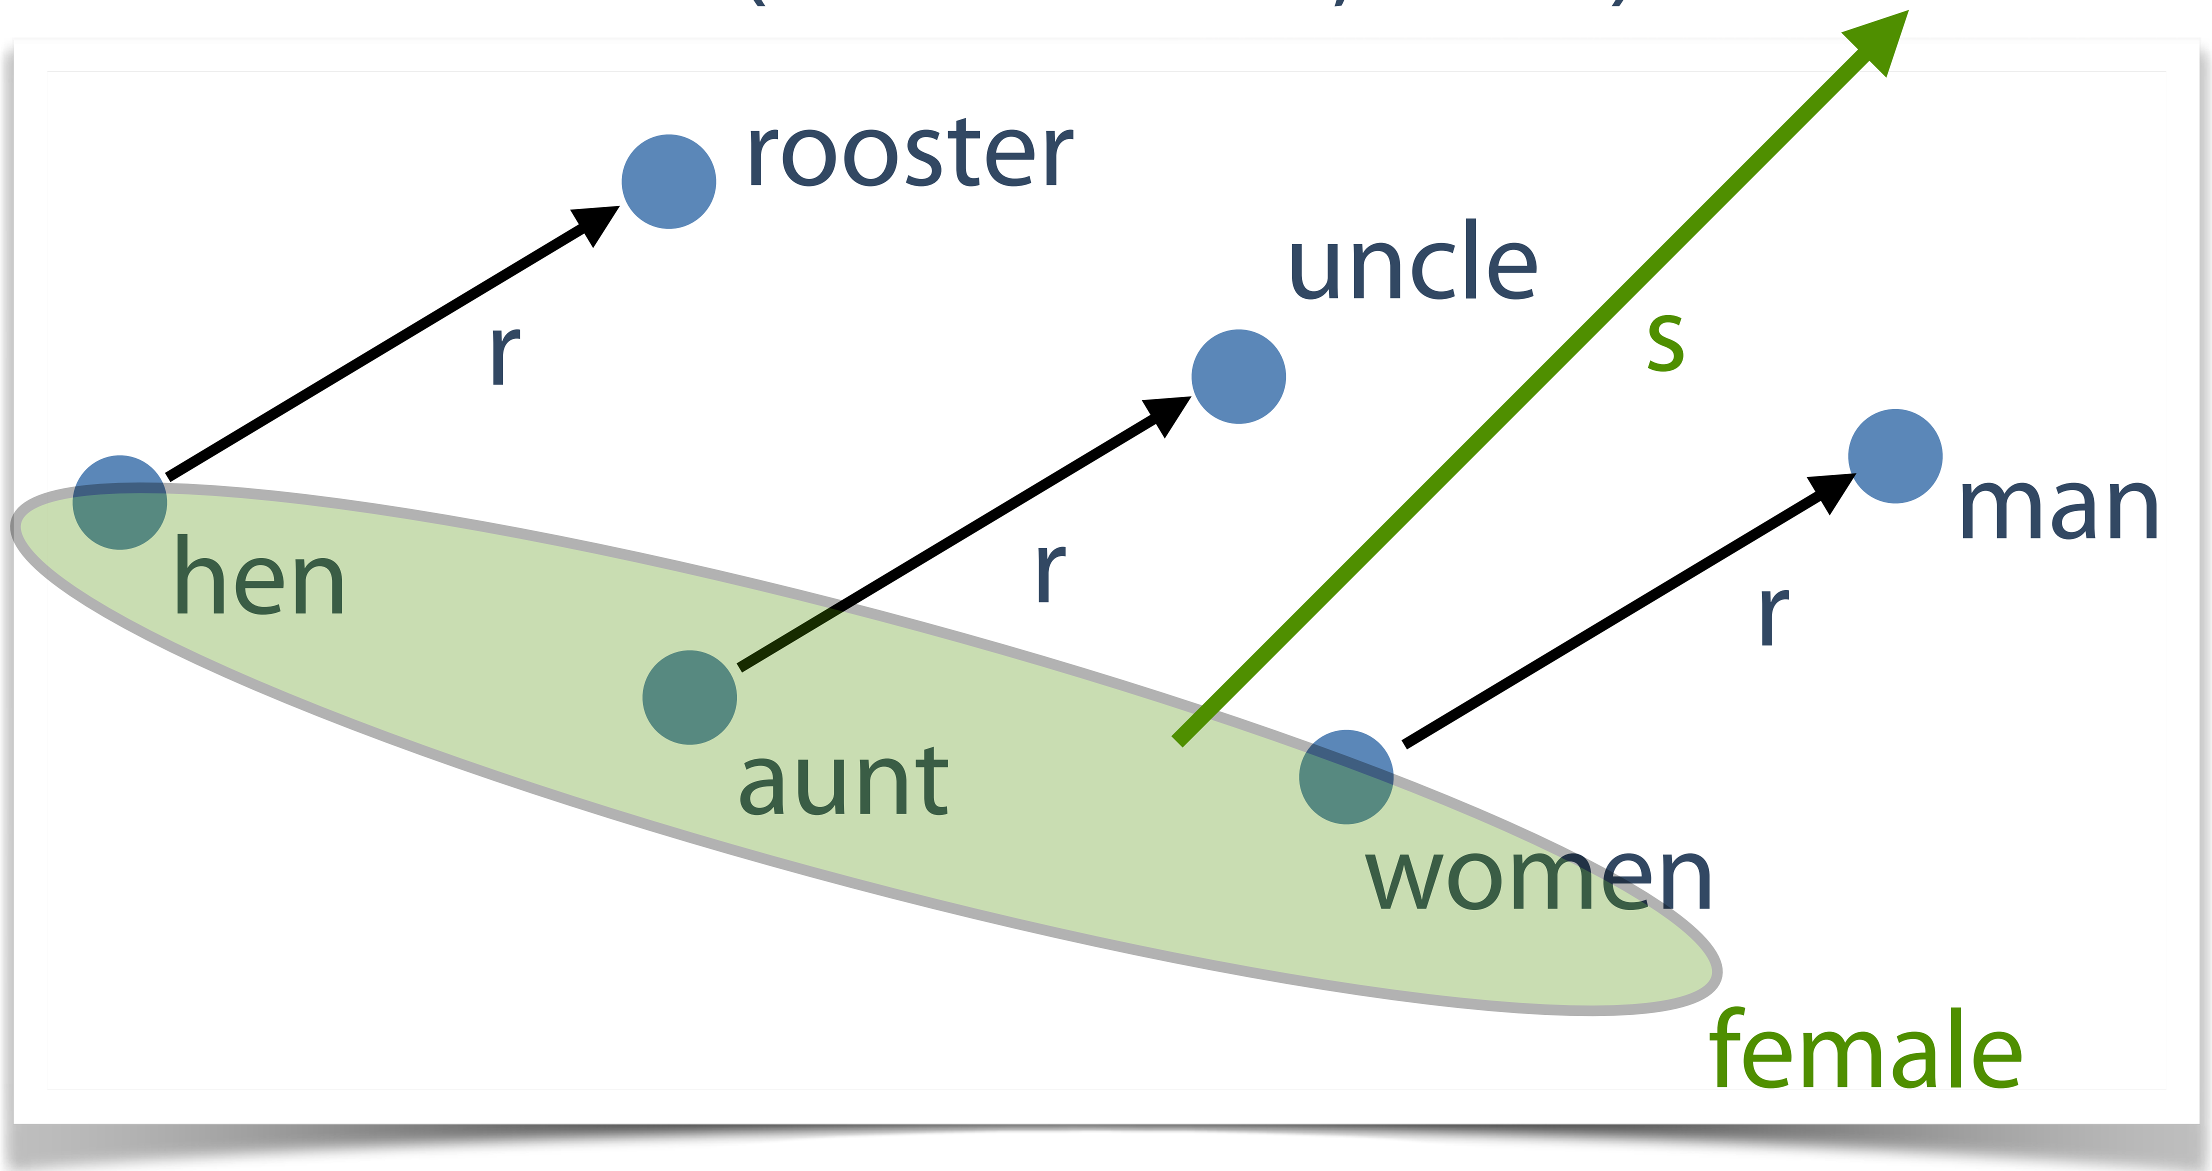

$$s_r(x, y) := ||x + r - y||_2$$

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole '18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible
- ▶ **Can we obtain an exact logic characterisation?**
  - …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

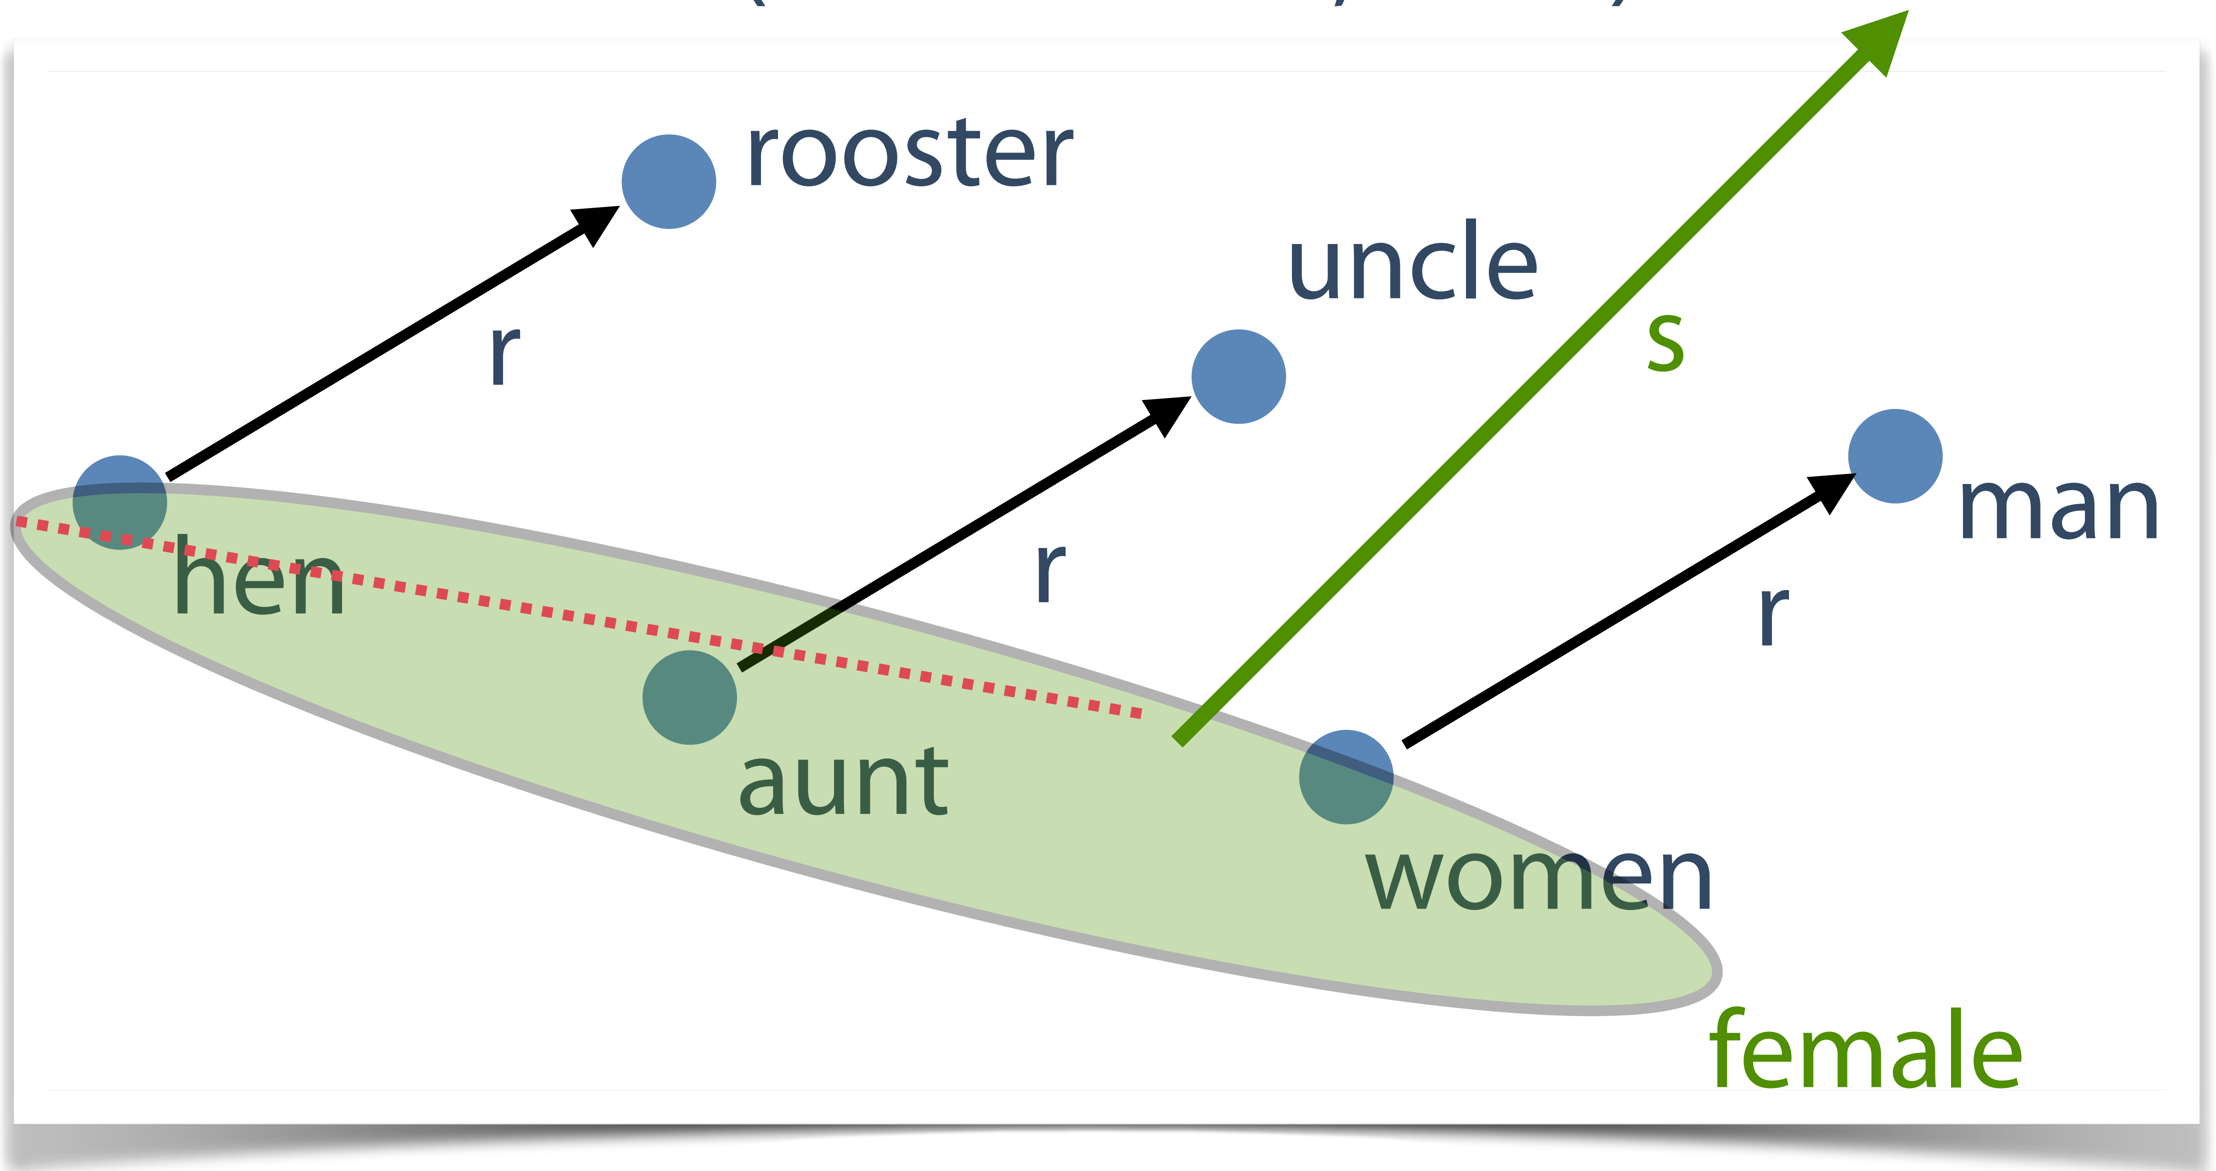

$$s_r(x, y) := ||x + r - y||_2$$

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole '18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible

▶ **Can we obtain an exact logic characterisation?**
  - …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

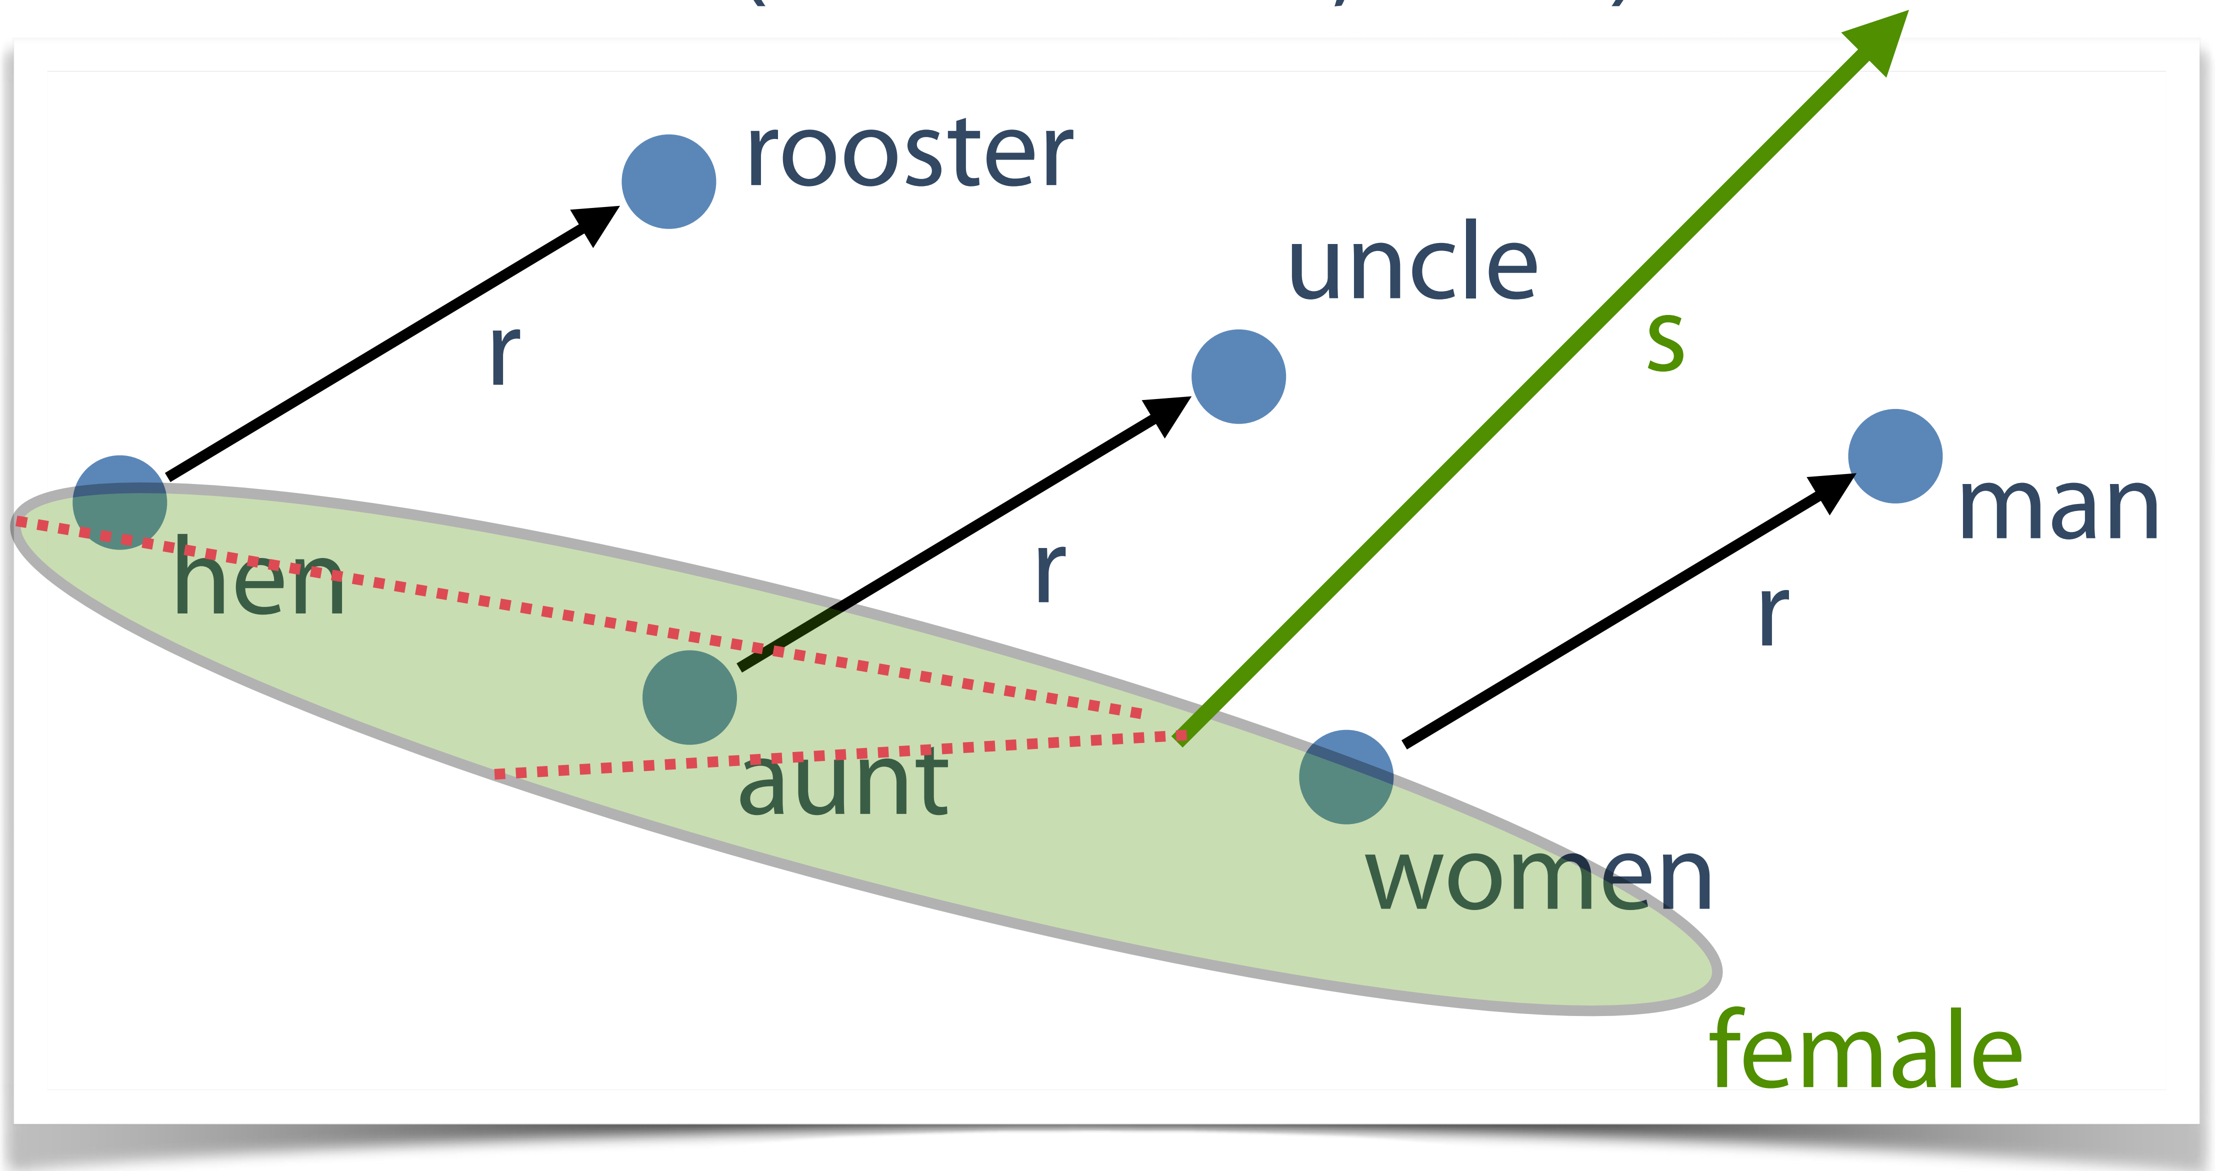

$$s_r(x, y) := ||x + r - y||_2$$

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole '18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible
- ▶ **Can we obtain an exact logic characterisation?**
  - …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

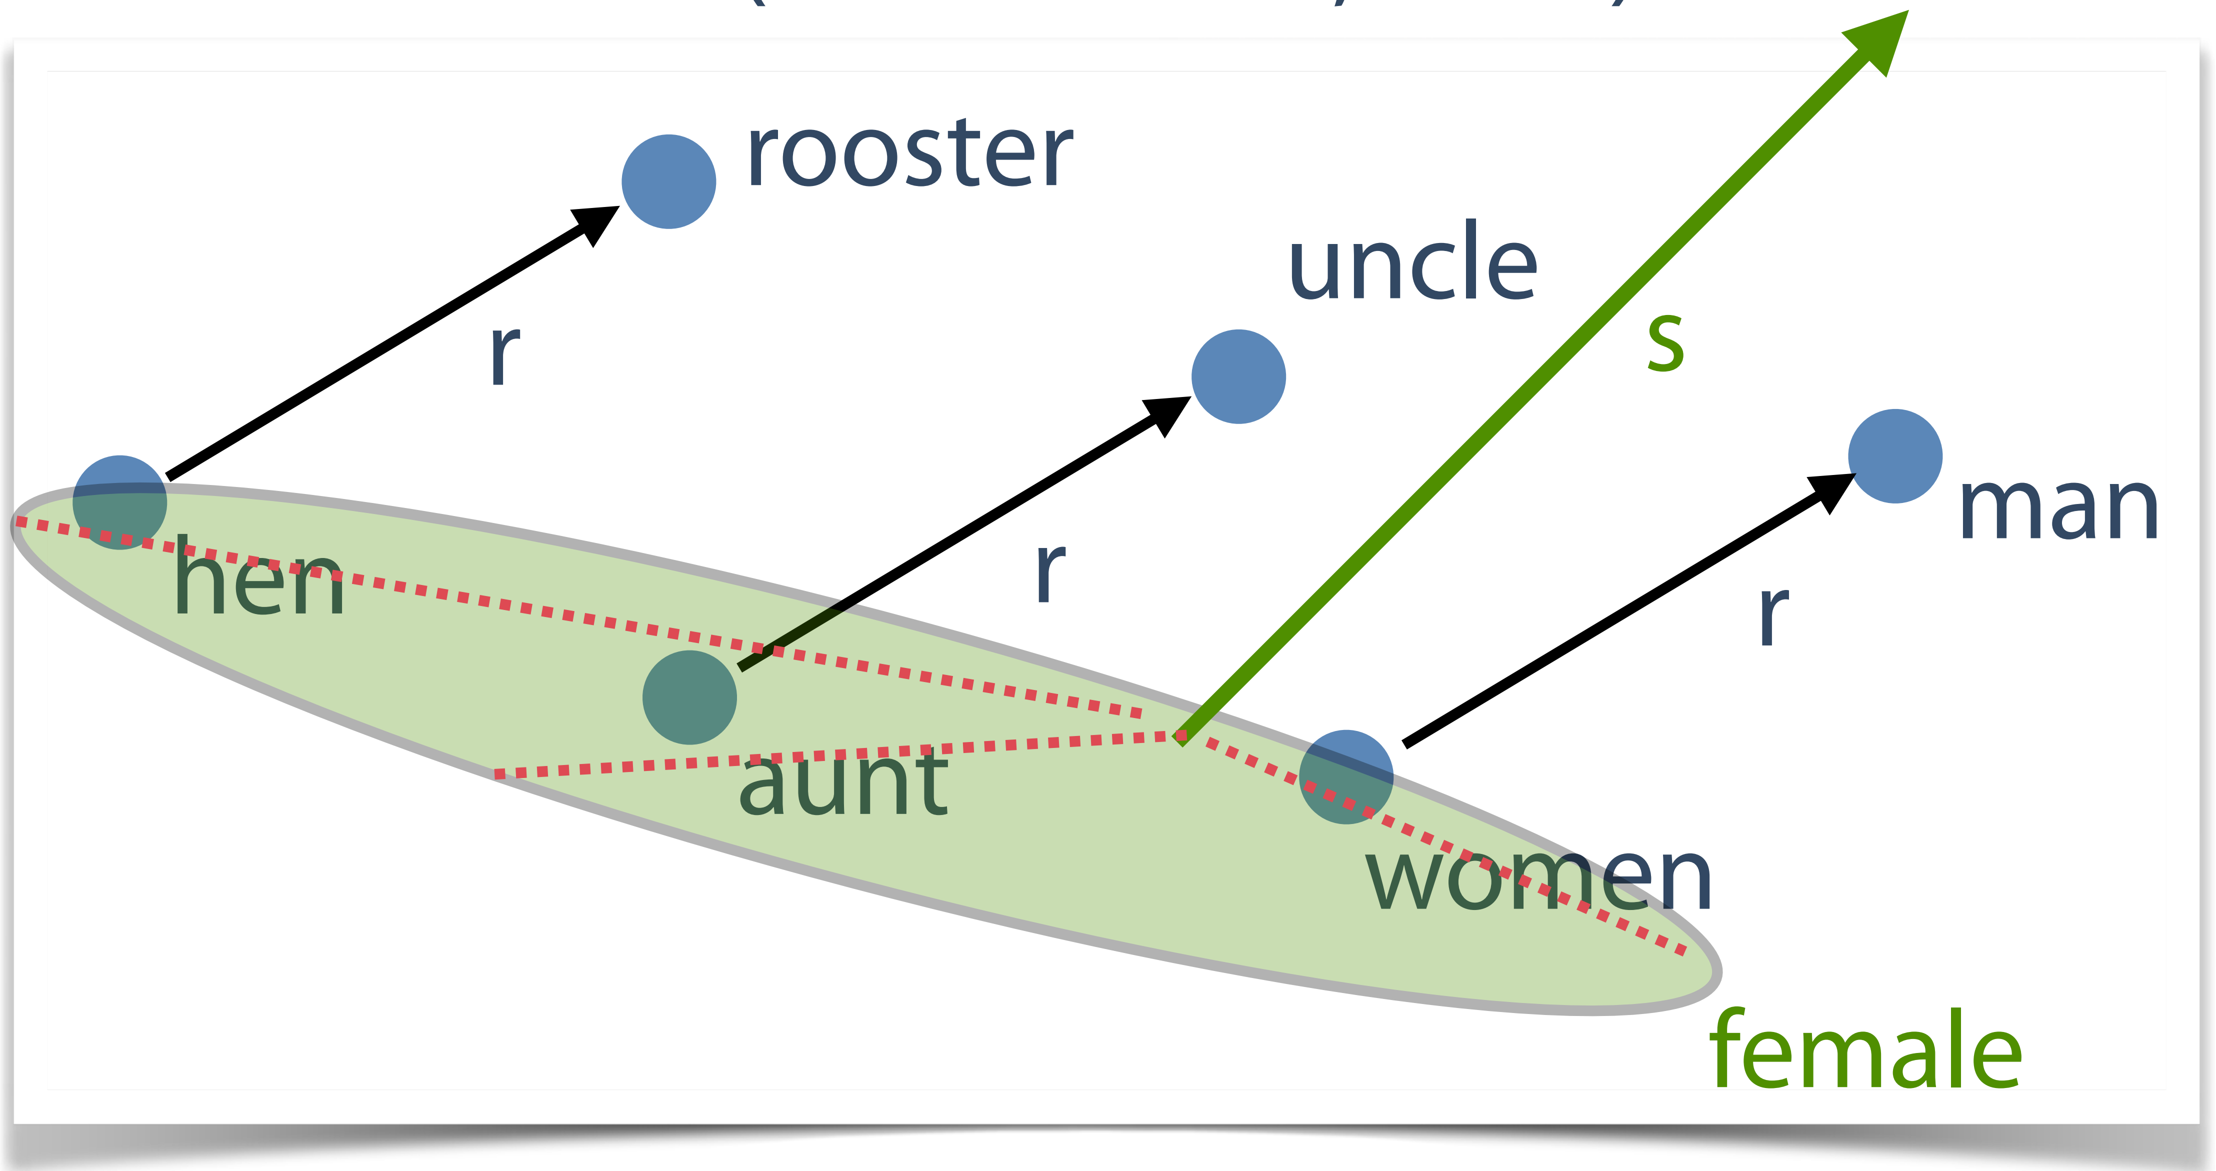

$$s_r(x, y) := ||x + r - y||_2$$

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole ’18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible

▶ **Can we obtain an exact logic characterisation?**
  - …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

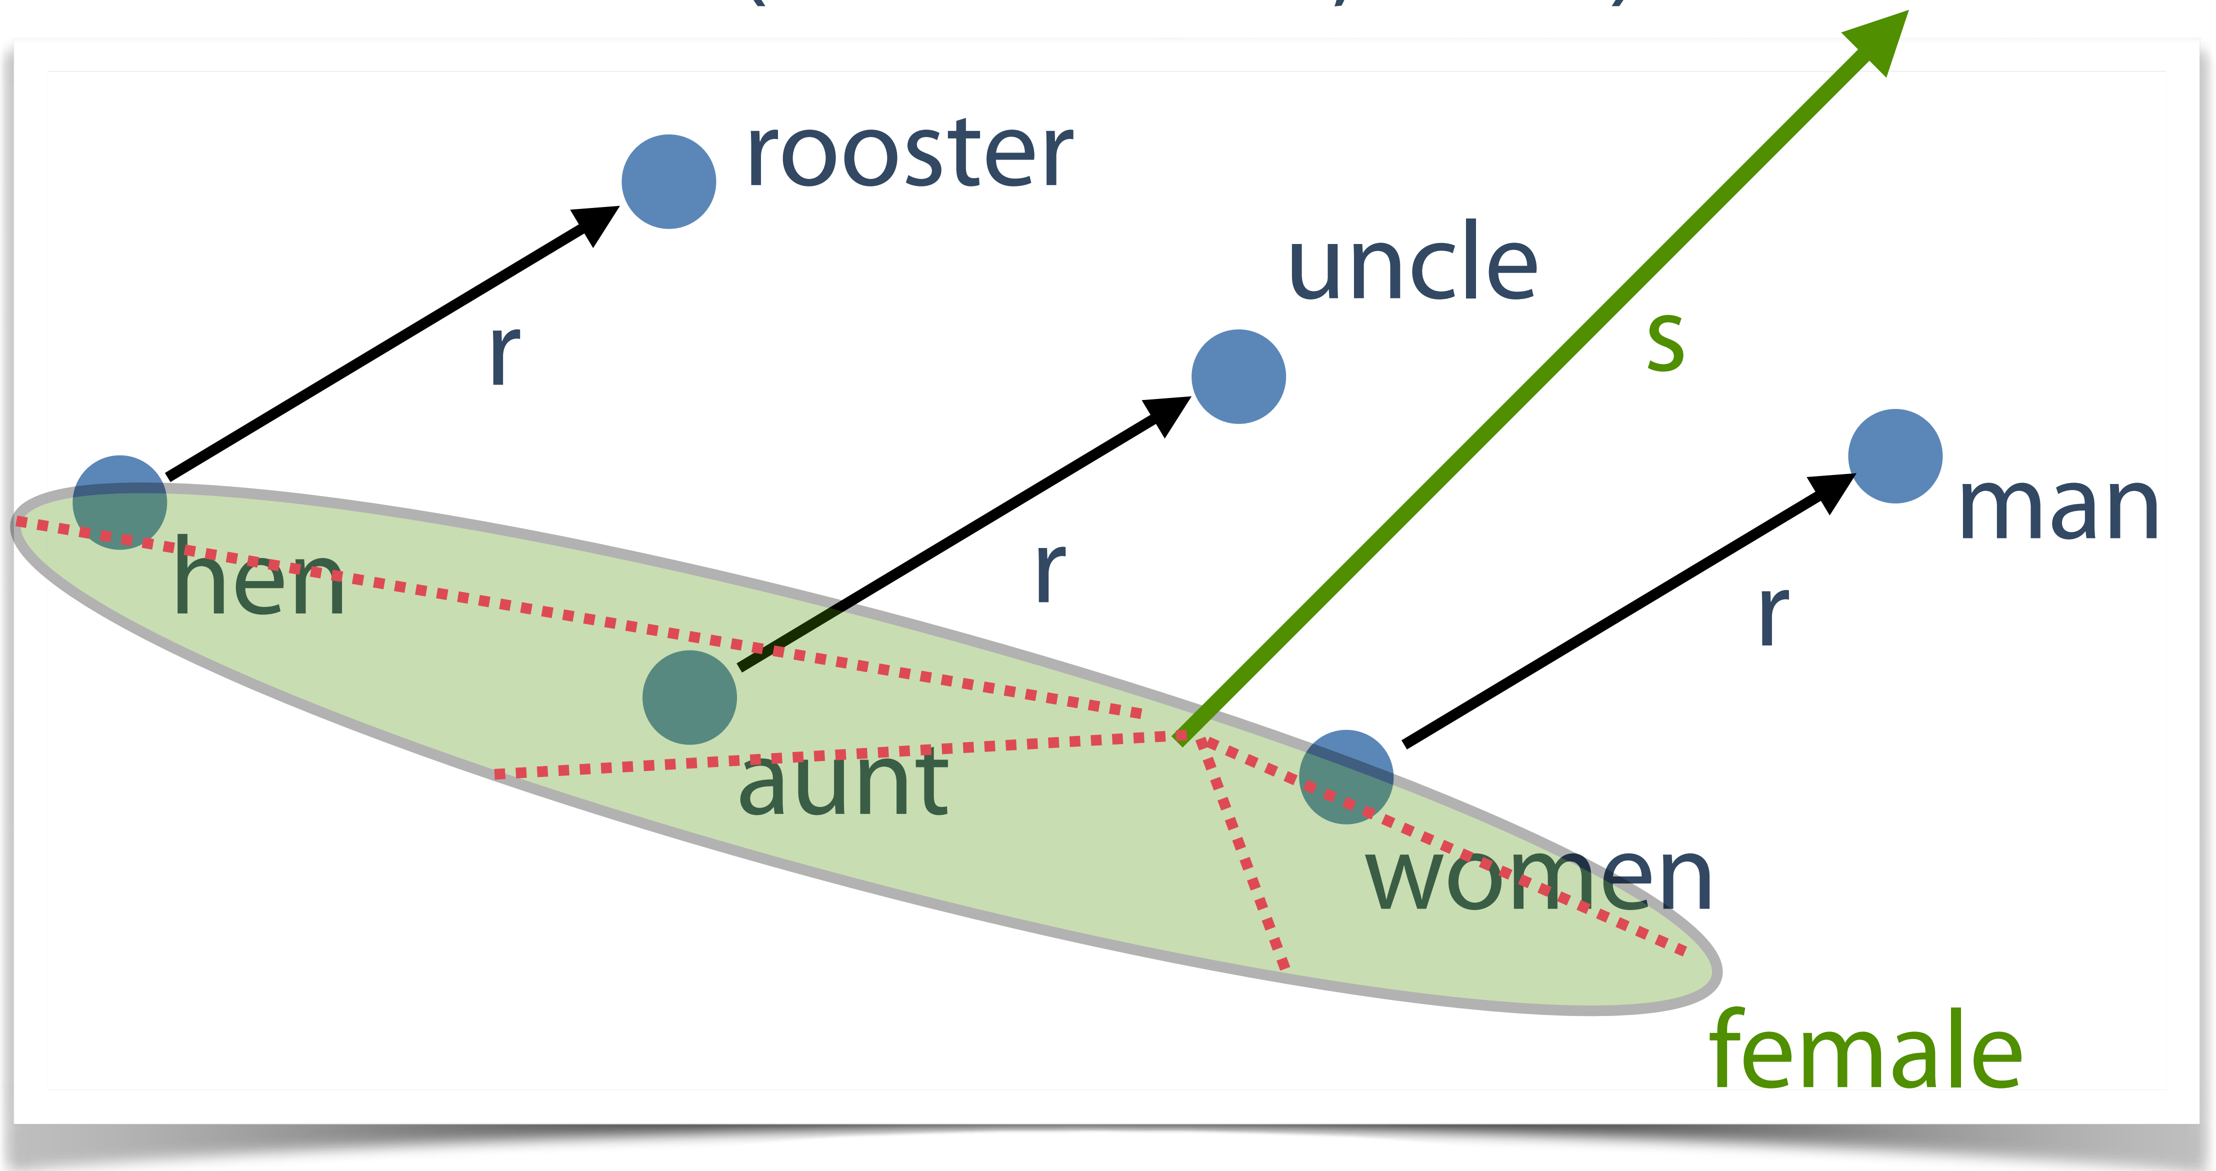

$$s_r(x, y) := ||x + r - y||_2$$

Example & Problem

- given a set of training data, KGEs determine an embedding that minimises errors from scoring functions
  - simple geometric models ease ML
- vector translations are **functional operations**
  - expressivity of relations restricted to functions
  - some compromise between likelihood semantics of scoring and geometric limitations possible
- **not fully expressive** (Kazemi & Poole '18)
  - expressivity only considered wrt. data/triples
  - deeper reasoning may still be impossible
- ▶ **Can we obtain an exact logic characterisation?**
  - …and push the envelope of expressivity?

TransE (Bordes et al., 2013)

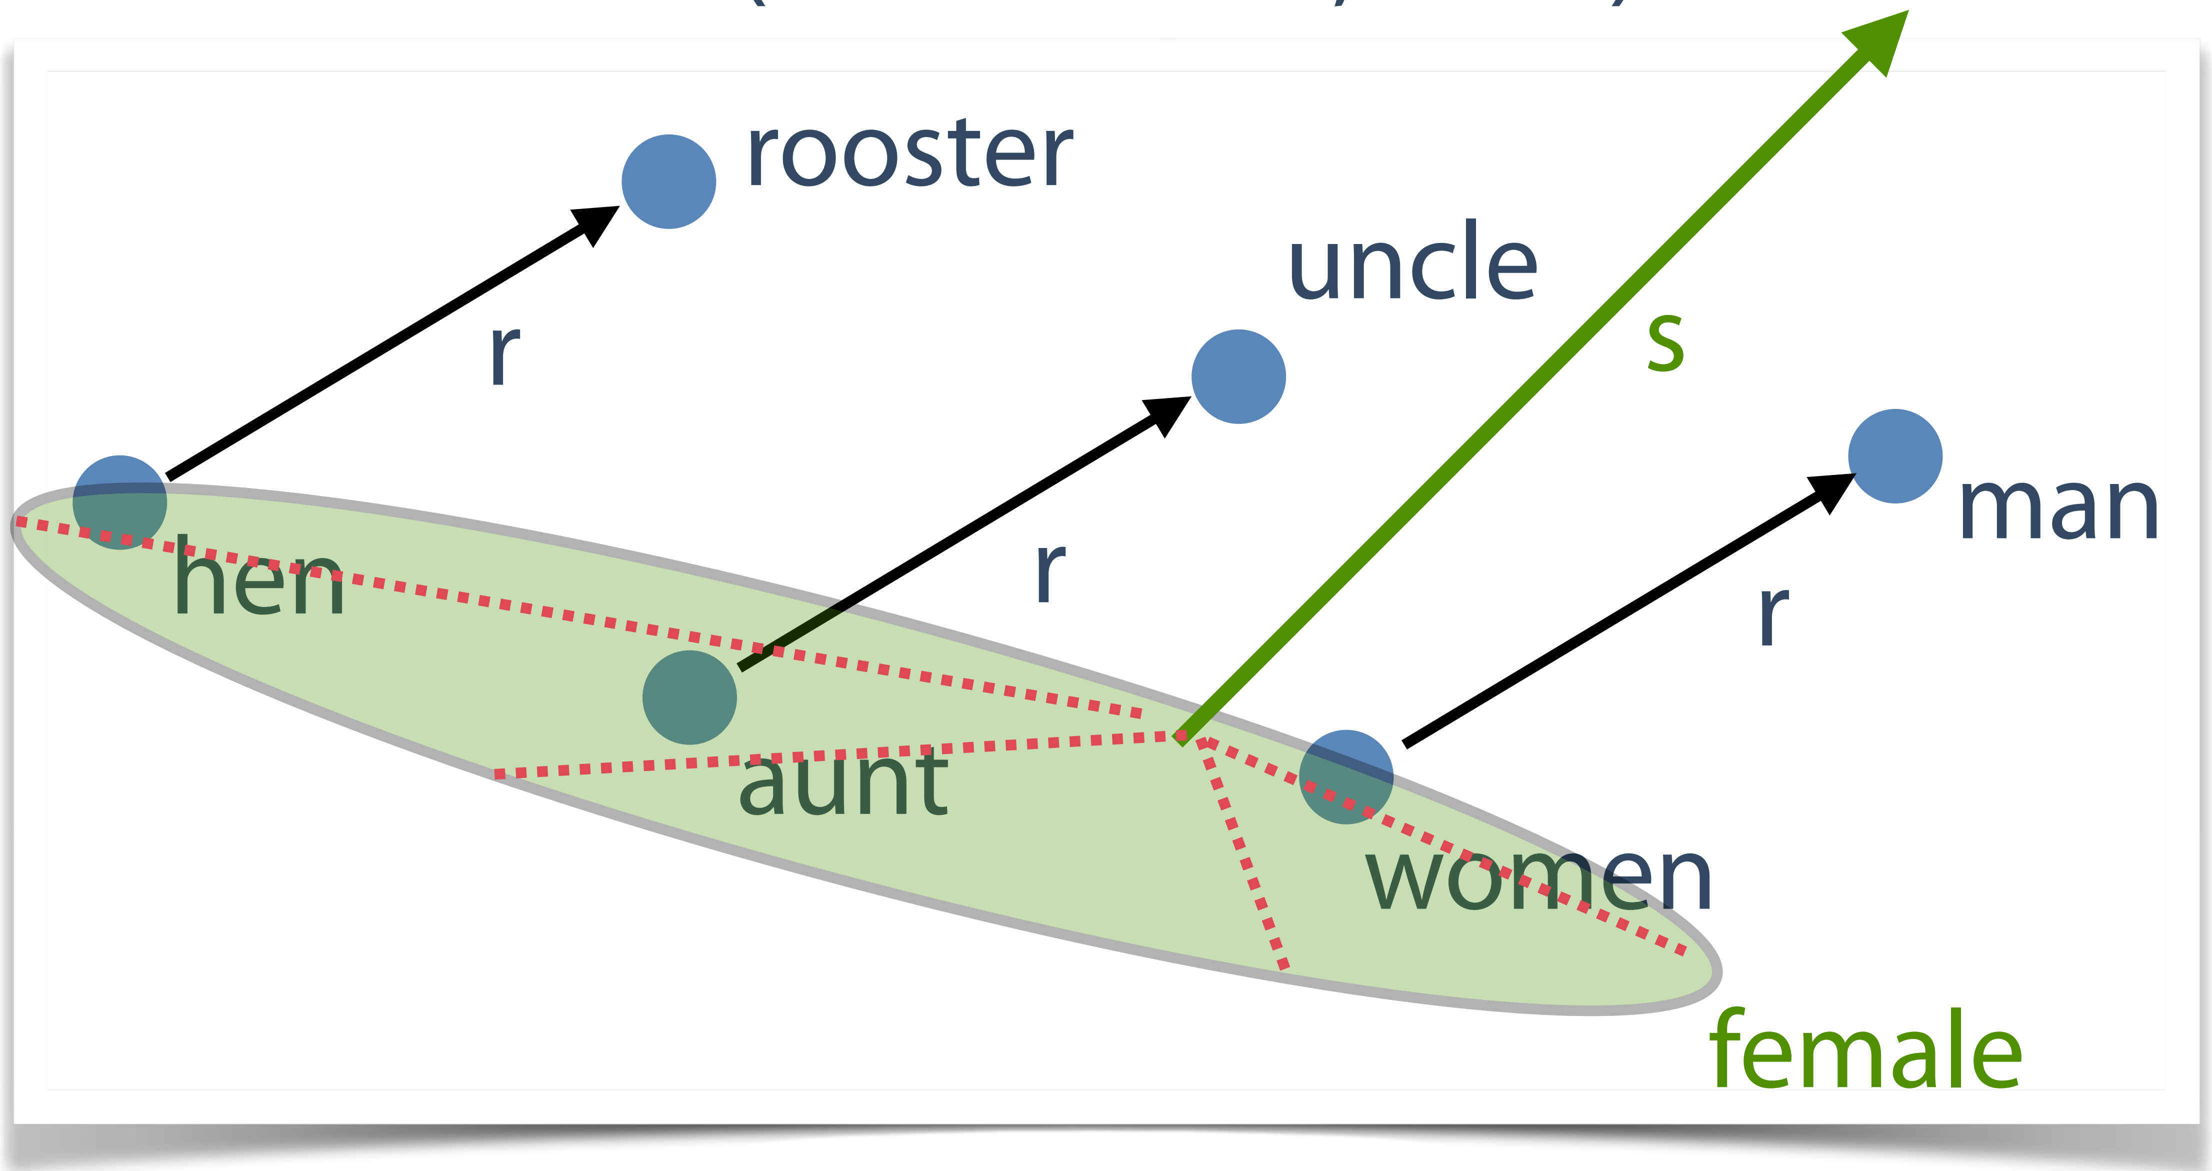

$$s_r(x, y) := ||x + r - y||_2$$

scoring value mixes uncertainty resulting from

- noise in data
- limitations of geometric model

Illustration Lack of Homogeneity

- Consider a SVM classifier using **hyperplanes** to classify concepts C,D
  - negations $\neg C, \neg D$ can be represented
  - conjunction $C \sqcap D$ cannot be represented as hyperplane (neither $C \sqcup D$)

- we say the geometric structure of the SVM is not **homogeneous** as the concepts it can represent is not closed under logic operations

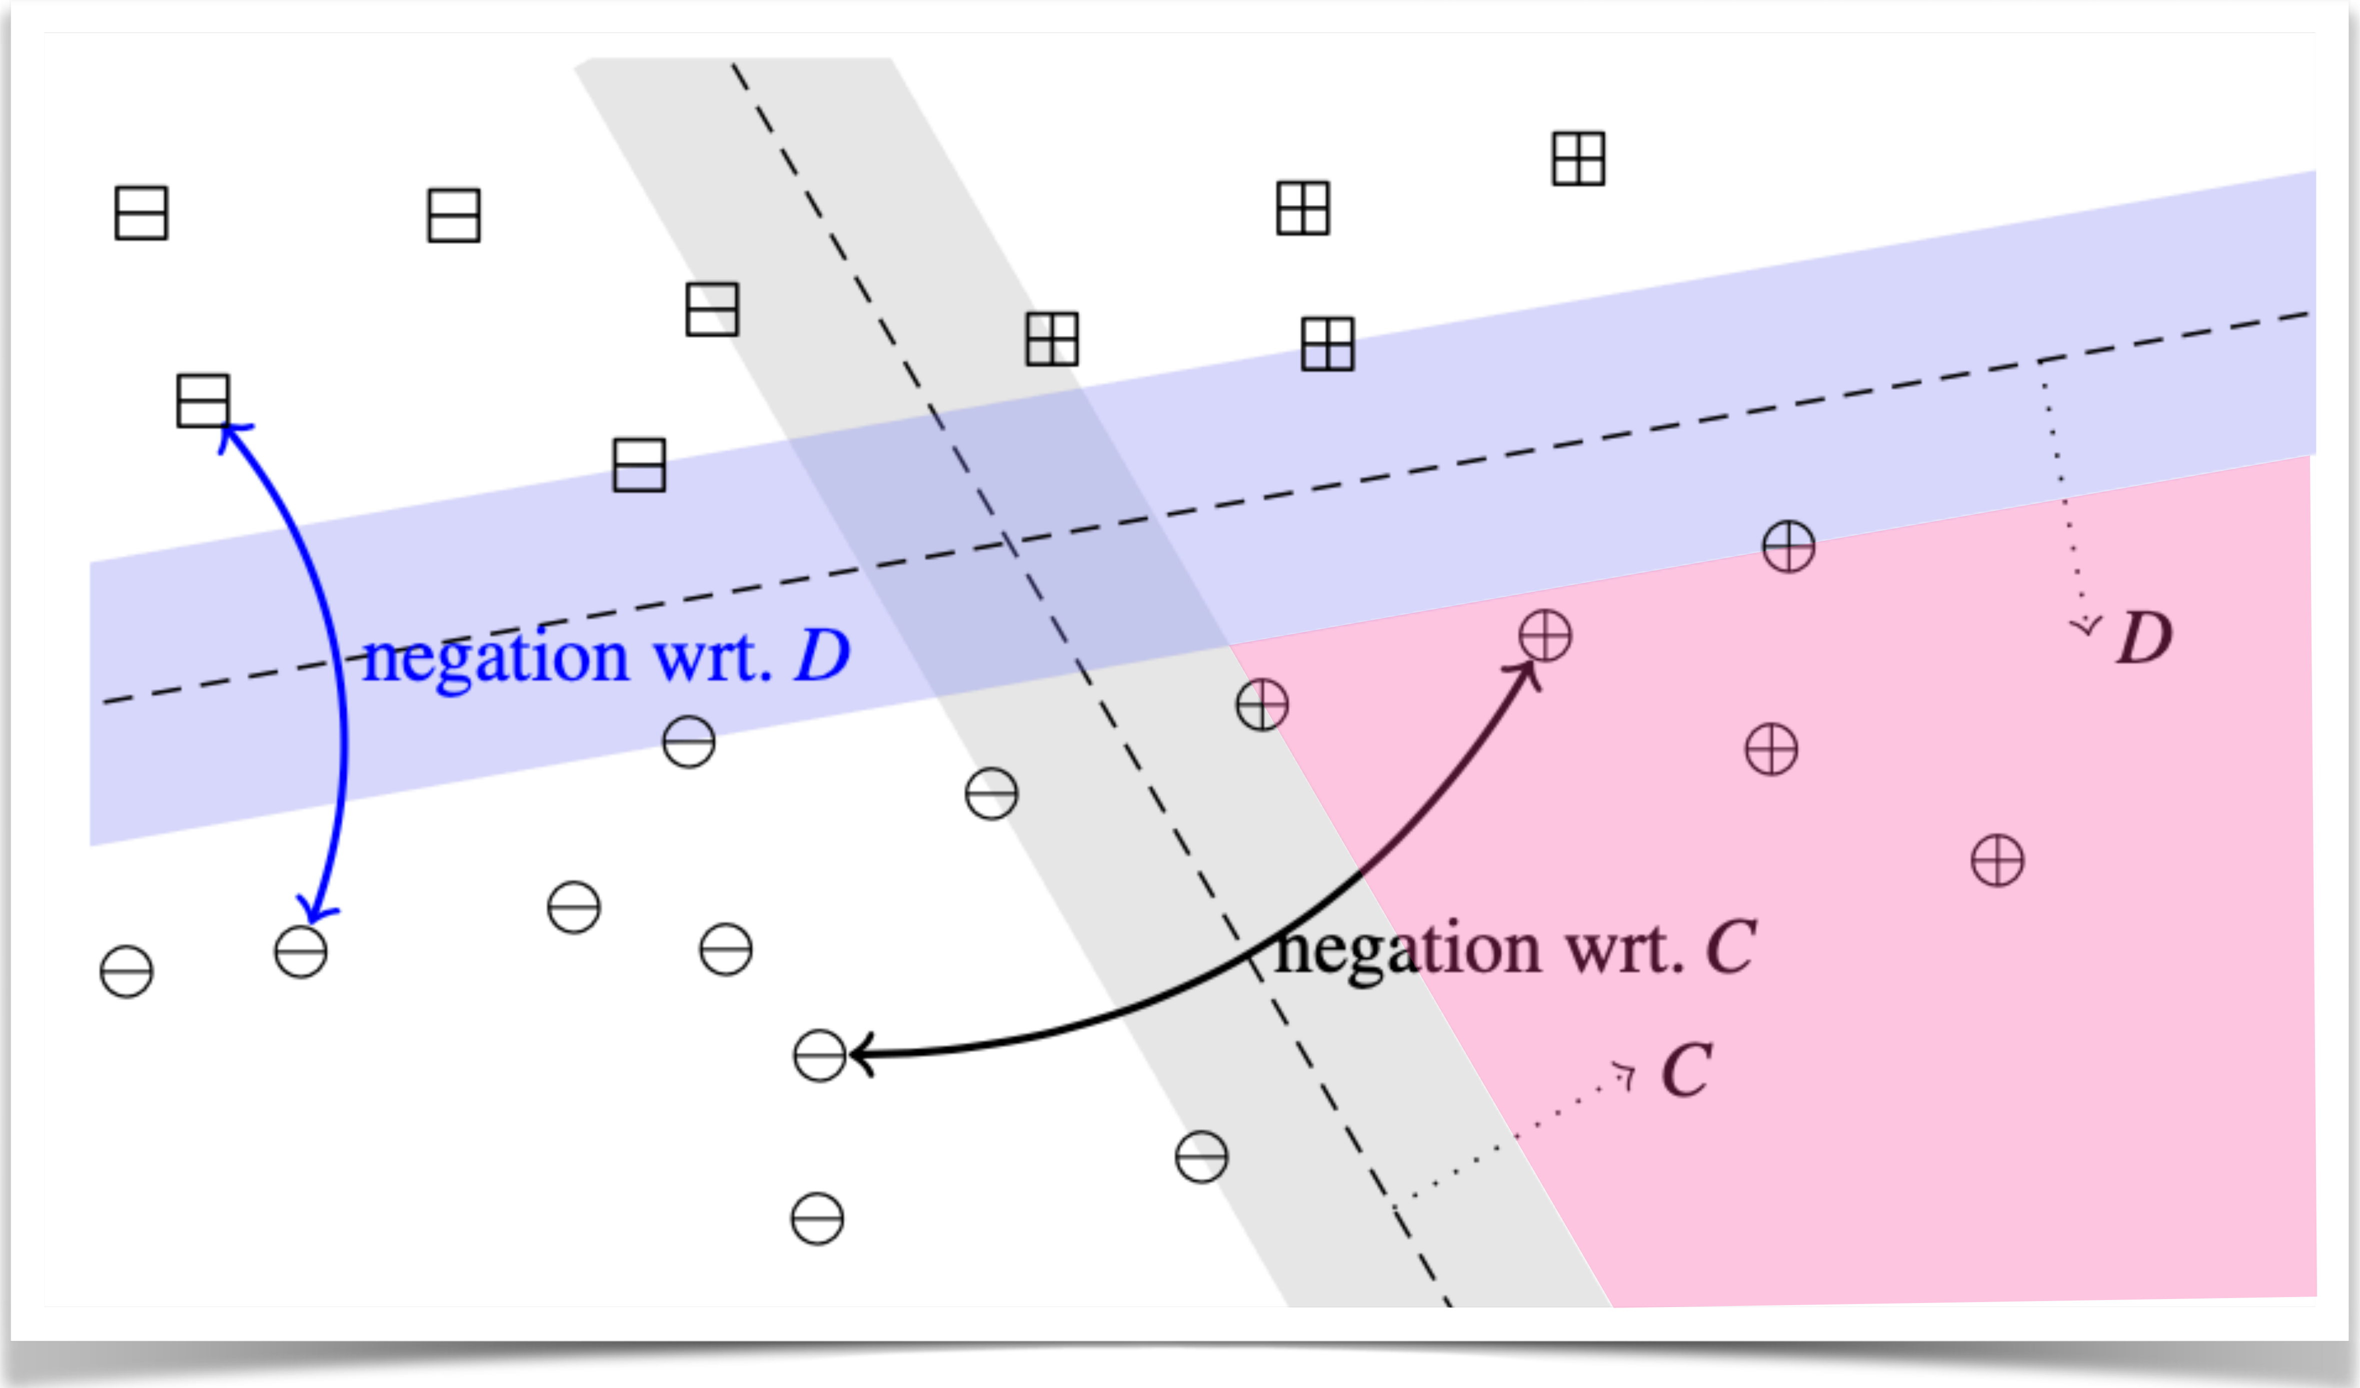

Logical & Geometrical Commitments

- designing a knowledge representation involves **ontological commitments**, i.e., how we choose to see the world (Davis et al. ’93)
- designing a KGE also induces **commitments**, some of which may be hard to identify
  - by committing to a set of geometries, we commit to a certain logic
  - widely ignored in data-driven investigations
  - some investigations on alignment of **logic and KGE** (e.g., Gutierrez-Basulto & Schockaert ’18, Kulmanov et al. ’19)

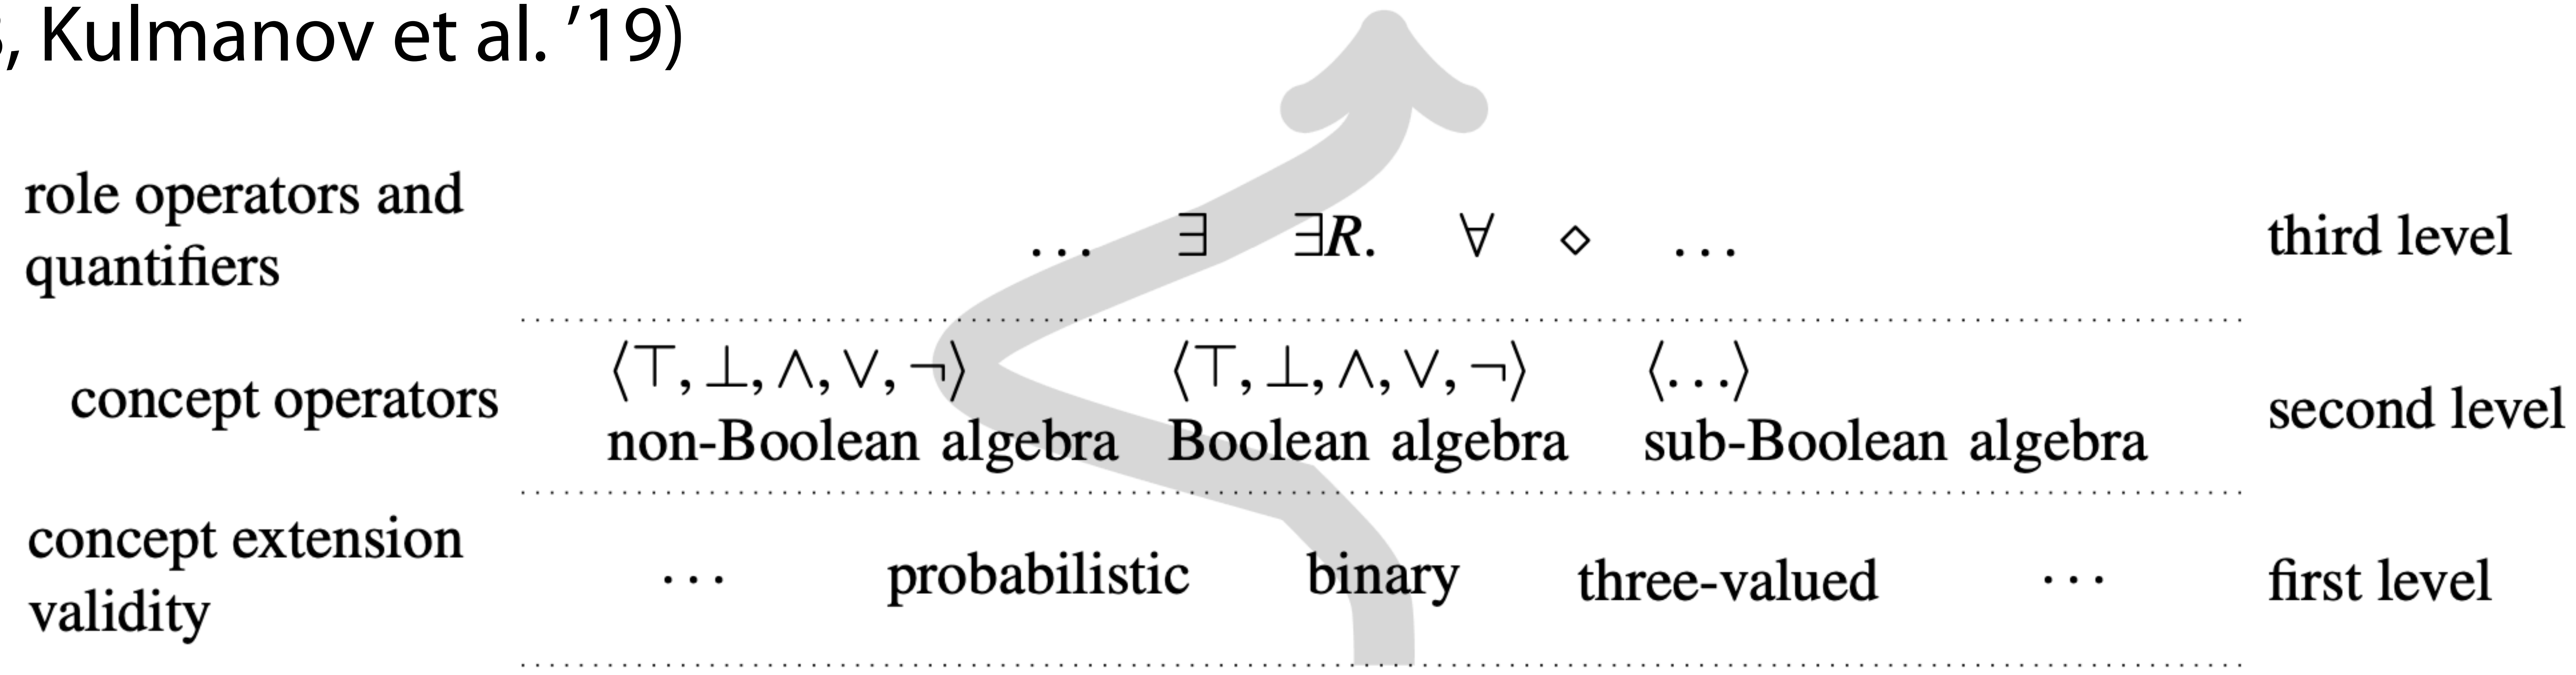

three levels of geometric-logic commitments

Aims of Logic-Based KGEs

- **full expressivity** of an underlying concept language
  - learned embedding supports deep symbolic reasoning
- enable **learning with background ontology**
  - *No multi-label learning problem is just about labels!*
  - **examples:** $\forall x \,.\, \mathrm{bird}(x) \rightarrow \mathrm{animal}(x)$, $(\mathrm{HorrorFilm} \sqcap \mathrm{FamilyFilm}) \sqsubseteq \bot$
- **idea:** concepts are not points, but geometrically shaped sets
  - relations between geometric entities
- **task:** identify pairing of concept language and geometric structure
  - expressive concept languages particular useful
  - easy geometric structures suggest better ML performance

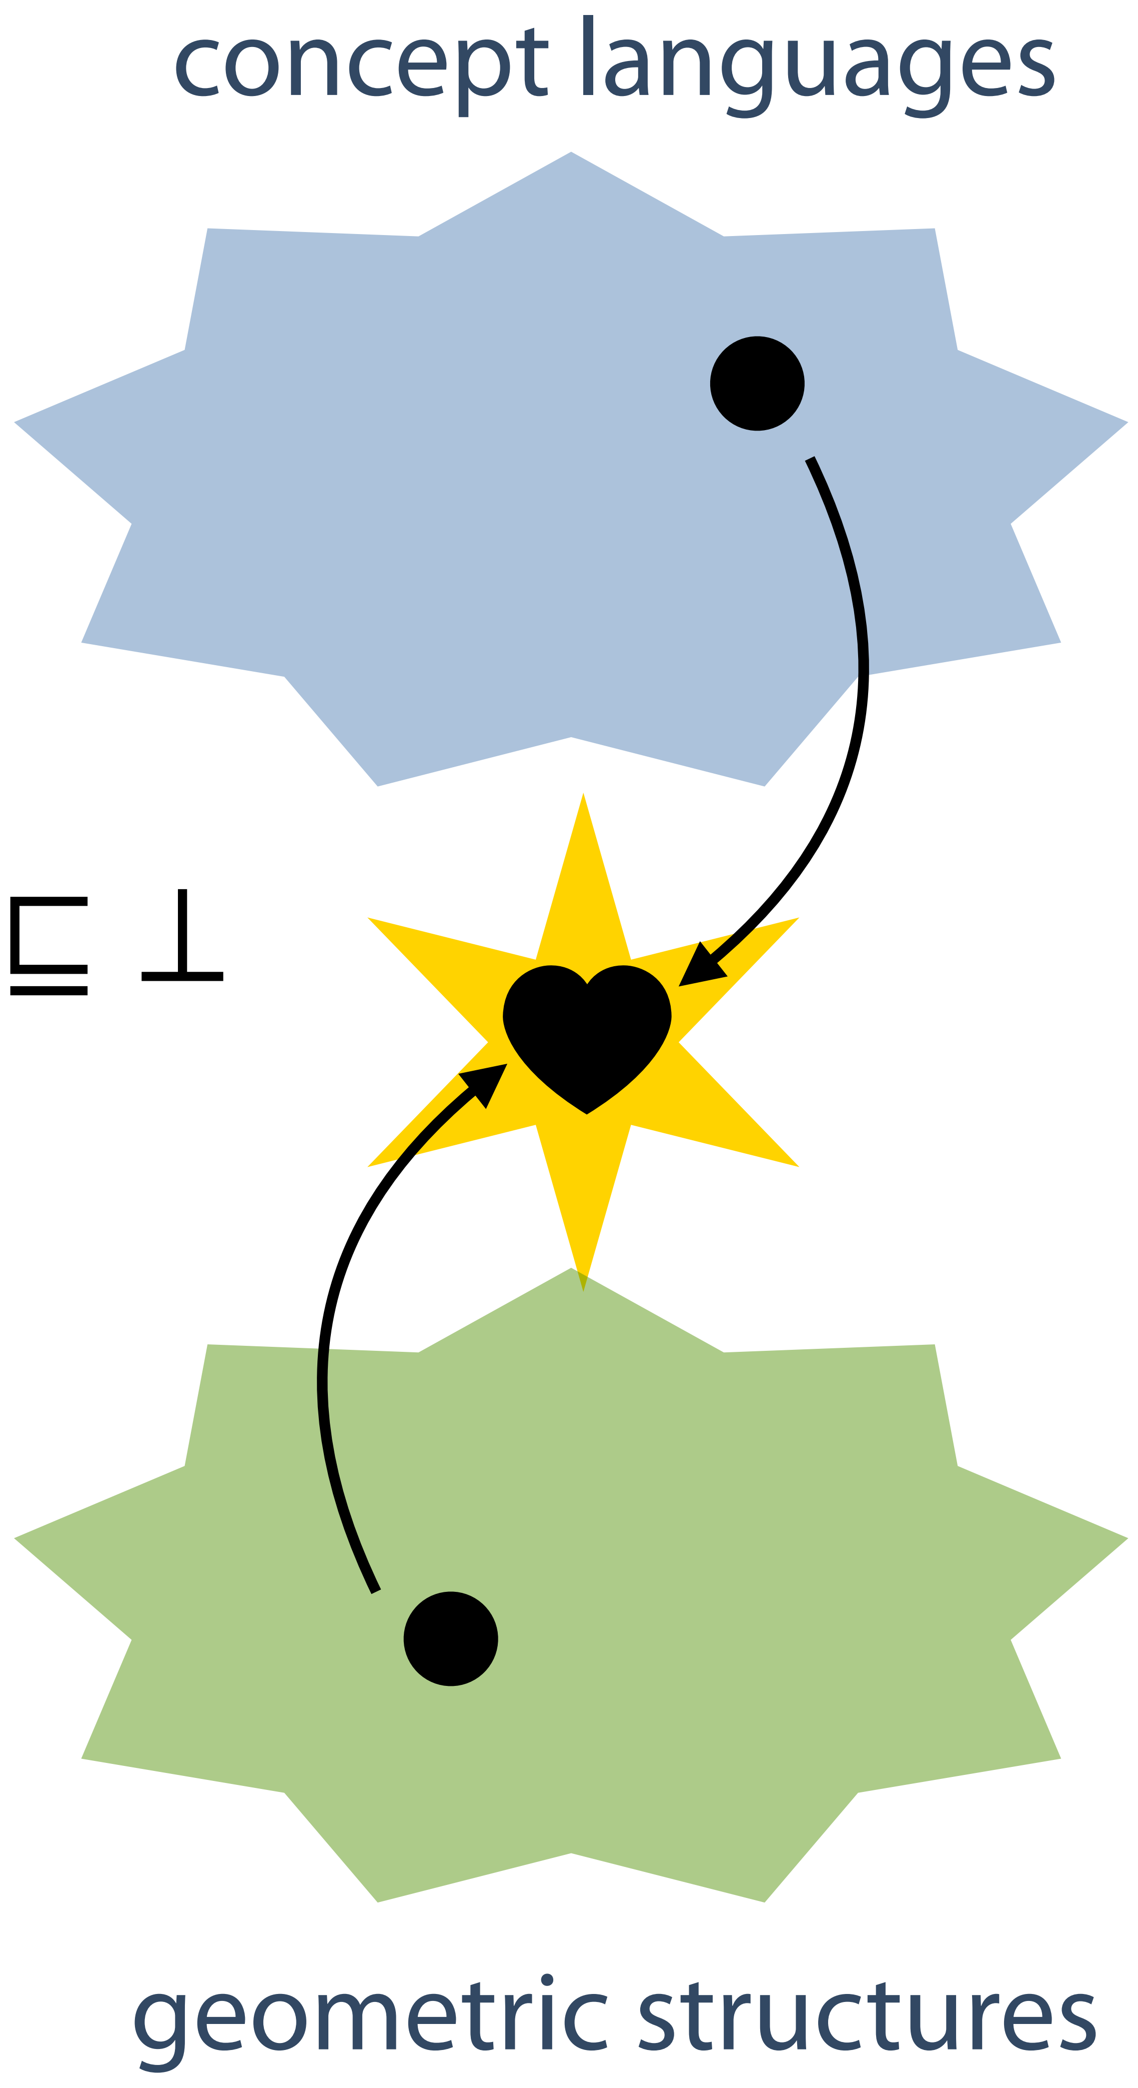

Conceptual Spaces

- Gärdenfors (2000) proposes to view concepts as embedded in a space
  - cognitively motivated theory of concept learning
  - quality dimensions spanning subspaces
- convex sets considered as natural categories
- concept-level reasoning grounded in spatial reasoning

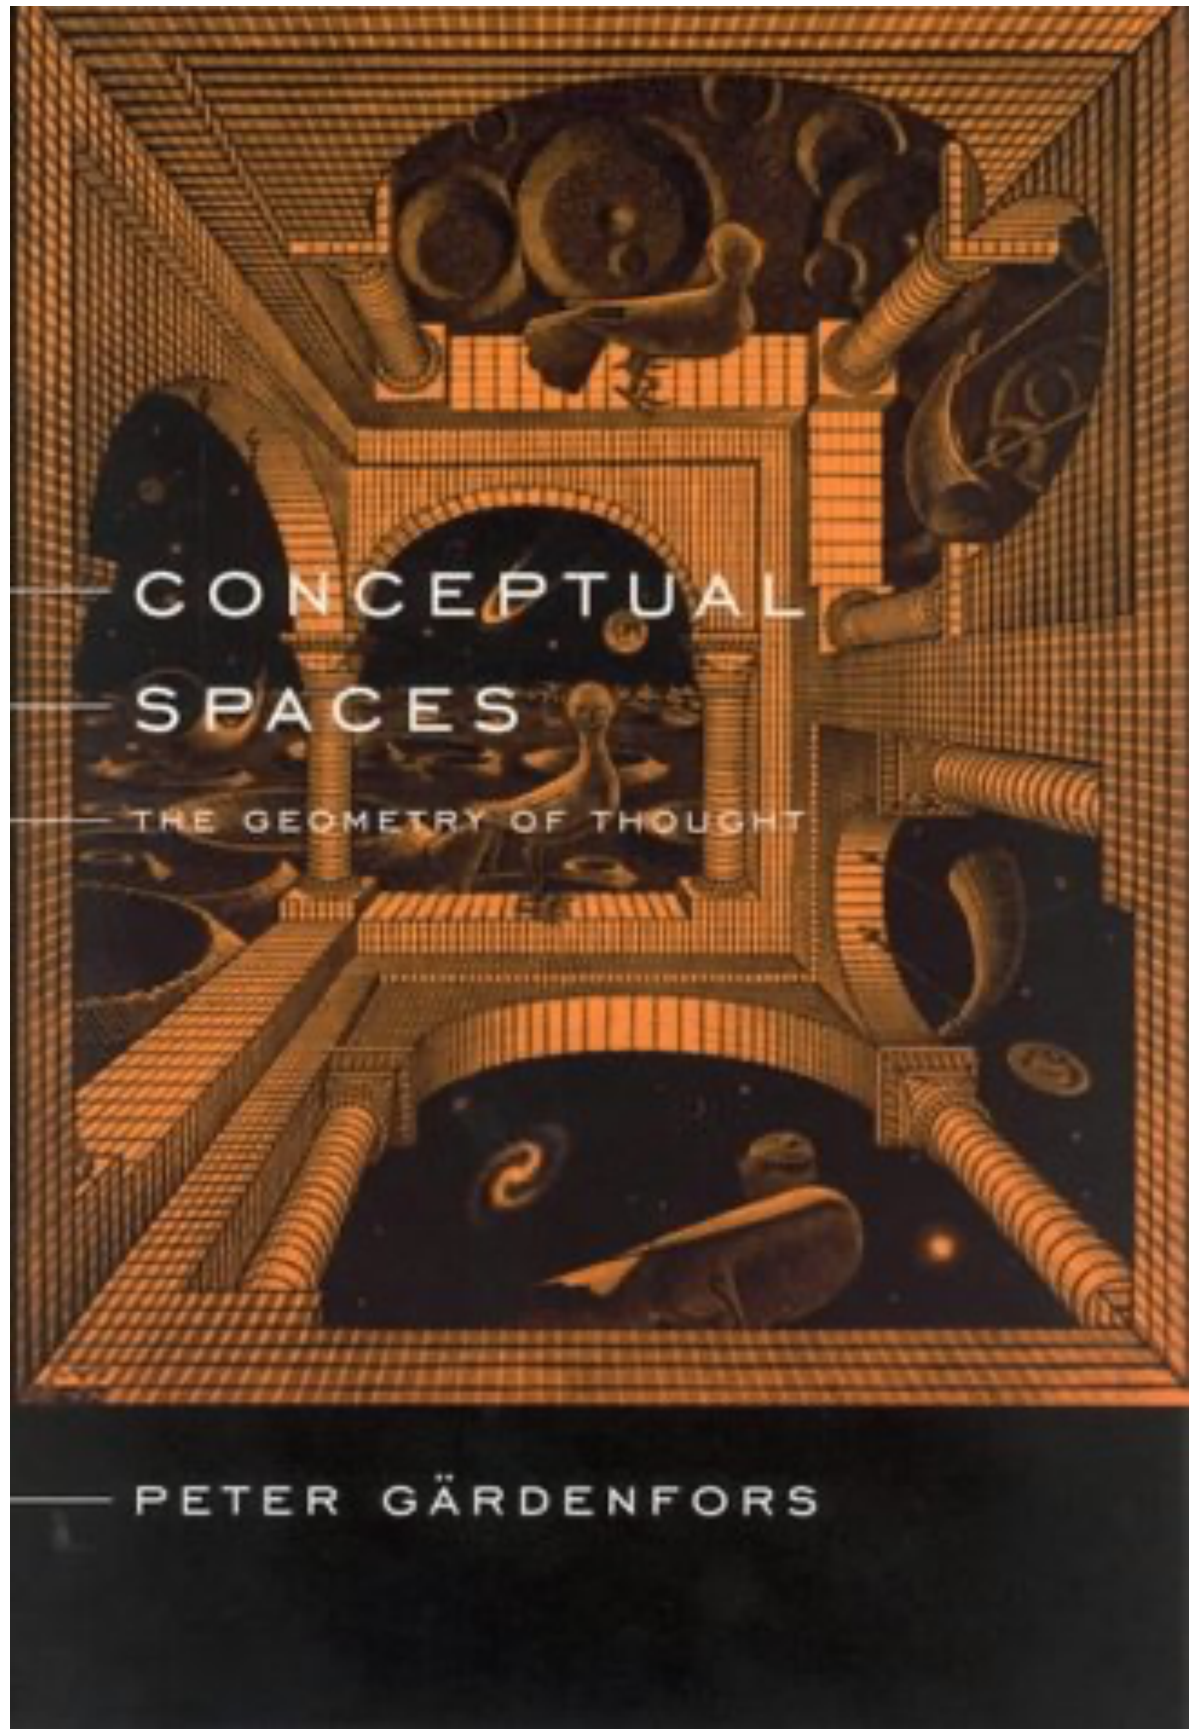

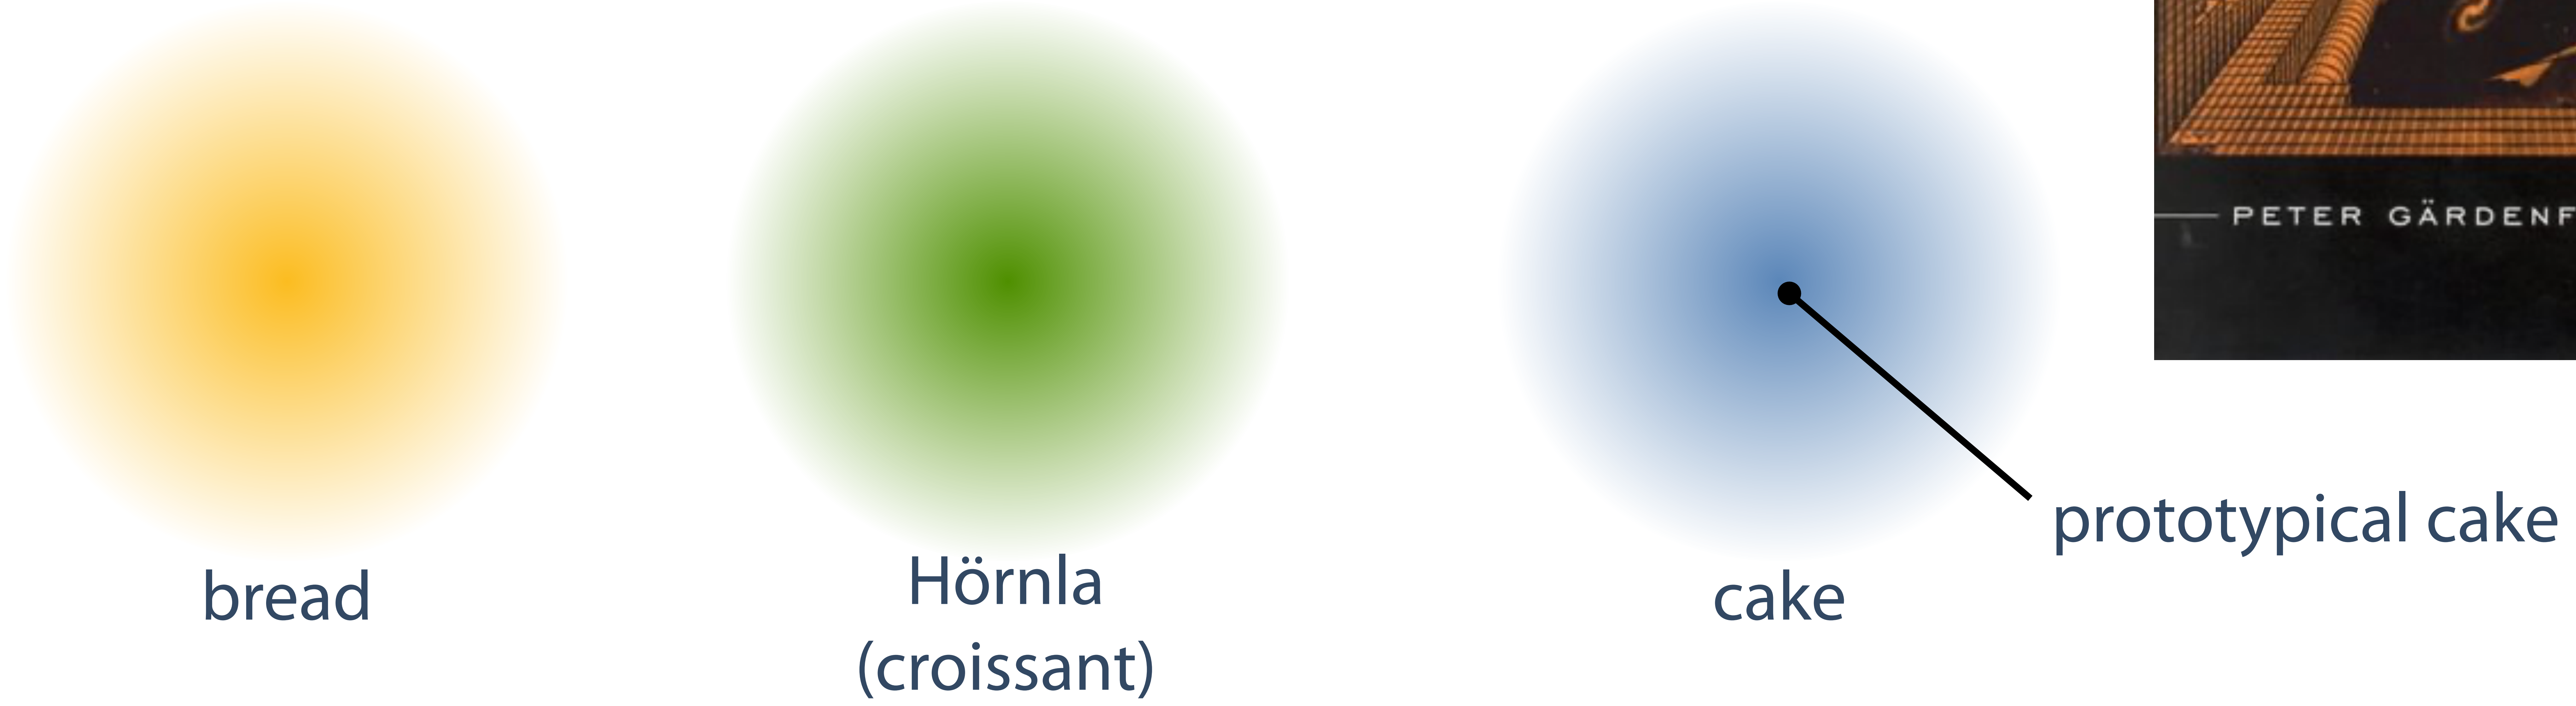

Related & Inspiring Approaches

geometrical structure	logic	concept lattice	negation	reference
convex sets	Quasi-Chained Datalog$^{\pm}$	distributive	atomic	Gutiérez-Basulto & Schockaert 2019
hyperspheres	$\mathcal{EL}$	distributive	atomic	Kulmanov et al. 2019
closed subspaces	Minimal Quantum Logic	orthomodular	orthonegation	Garg et al 2019
axis-aligned cones	$\mathcal{ALC}$	distributive	full Boolean	Özçep et al 2020

remainder of this talk

- important feature: **full Boolean concept negation**
  - negative information introduces **uncertain information**
    $\neg\mathrm{genre}(x, \mathrm{horrorFilm}) \rightarrow \mathrm{genre}(x, \mathrm{familyFilm}) \vee \mathrm{genre}(x, \mathrm{fantasyFilm}) \vee \ldots$
  - negation also allows coverage to be expressed, e.g.,
    $\neg\mathrm{inState}(\mathrm{cat}, \mathrm{hungry}) \leftrightarrow \mathrm{inState}(\mathrm{cat}, \mathrm{sleepy})$

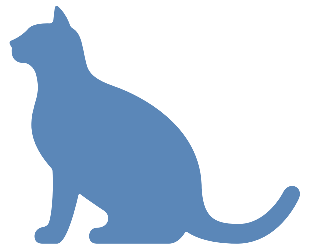

Description Logic $\mathscr{ALC}$

- **terminological (tbox) and assertional (abox) knowledge**
  - background ontology and facts
- **concepts**, represented by symbols (C, D, ...)
  - concept subsumption $\sqsubseteq, \sqsupseteq, \equiv$
  - set-theoretic operations $\sqcap, \sqcup, \neg$
  - role quantification $\exists r . C, \forall r . C$
- **assertions**, using constant and role symbols
  - represent facts

$R \sqsubseteq H$	researcher, humans
$R \sqcup \neg H$	researcher or alien
$R \sqcap \neg H \sqsubseteq \bot$	we are human!
$\exists \text{hasPaper} . \text{AITopic}$	"AI expert"
$R \sqcap \exists \text{hasPaper} . \text{AITopic}$	"AI researcher"

$$\text{human}(\text{alex}), \text{hasPaper}(\text{alex}, p_1), \cdots$$

Convex Cones

efficient convex optimisation techniques

- **idea**: interpret concepts as **convex cones**

expressive geometric/algebraic operations

  - $X \subseteq \mathbb{R}^n$ is a convex cone iff $\forall y, z \in X, \lambda, \mu \in \mathbb{R}_{\geq 0} \,.\, \lambda x + \mu y \in X$
  - as usual, individuals interpreted as points

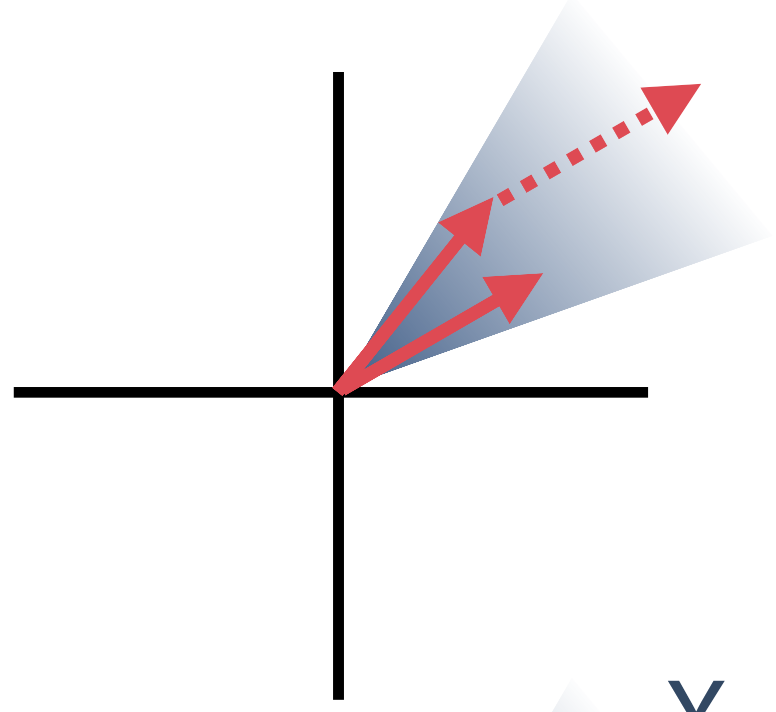

- **idea**: use **polarity as negation** (derived from scalar product)
  - polar of cone defined as $X^\circ = \{x \in \mathbb{R}^n \,|\, \forall y \in X \,.\, \langle x, y \rangle \leq 0\}$
  - using the usual scalar product $\langle x, y \rangle = x^T \cdot y$

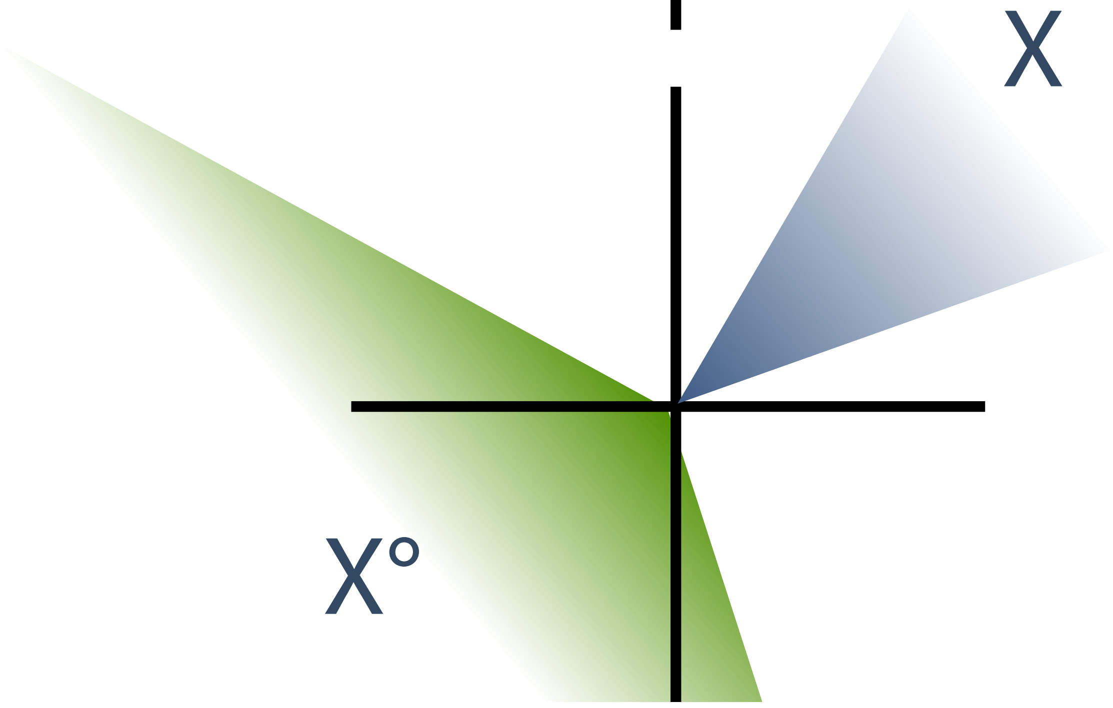

- properties
  - convex cones closed under polarity
  - $(X^\circ)^\circ = X$
  - $\mathrm{convexHull}(X \cup Y) = (X^\circ \cap Y^\circ)^\circ$

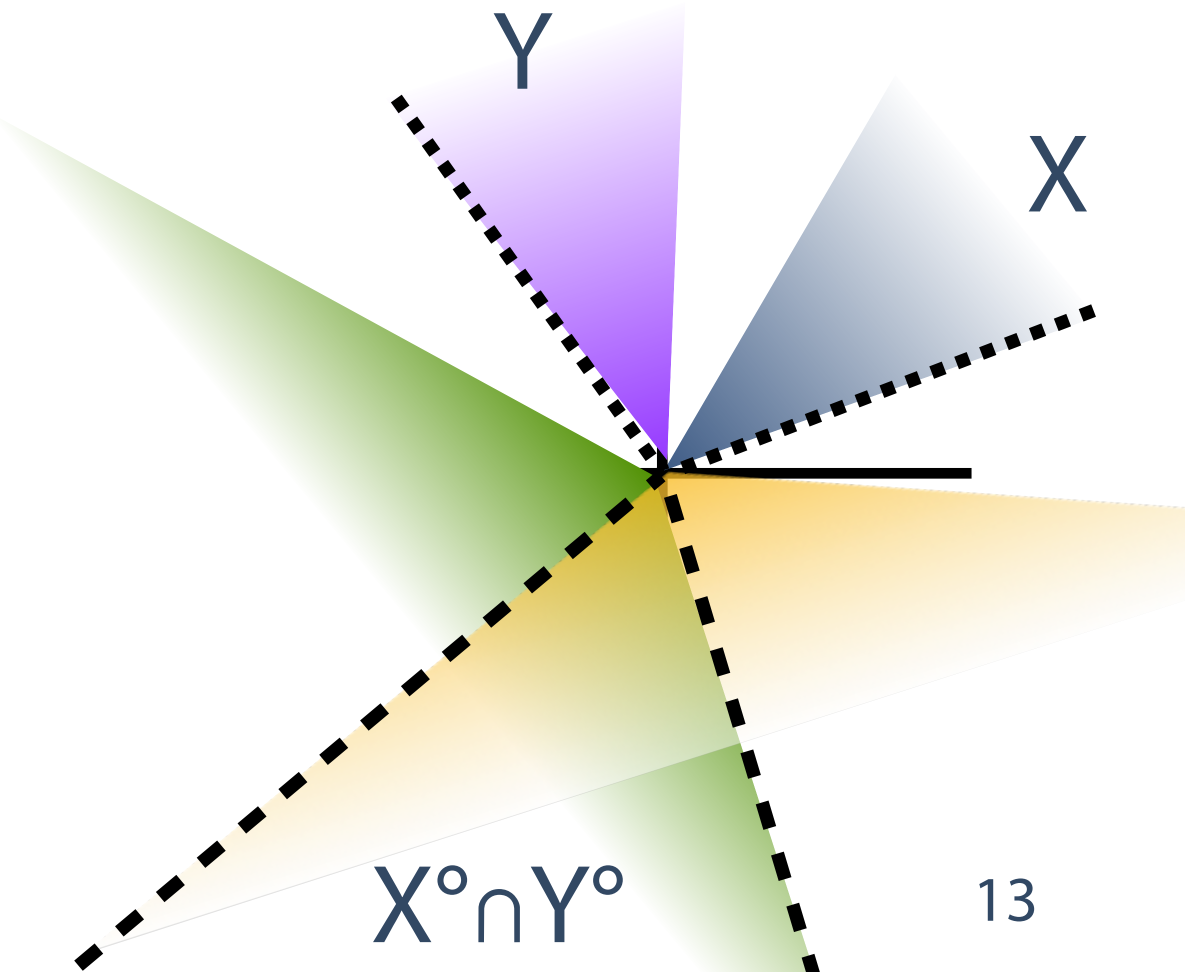

Geometric Model (for Boolean $\mathscr{ALC}$)

- consider an embedding I as interpretation function…
  - interpreting constants as points in $\mathbb{R}^n$
  - interpreting concepts as cones
  - interpreting $\top$ as $\mathbb{R}^n$ and $\bot$ as $\{(0\cdots0)^T\}$
- concept terms are defined inductively using geometric operations of intersection (conjunction), polarity (negation)
  - disjunction via De-Morgan
  - we have $(C \sqcup \neg C)^I \equiv \top$
- we say the embedding to a model in the logic sense for
  - $C(a)$ iff $a^I \in C^I$

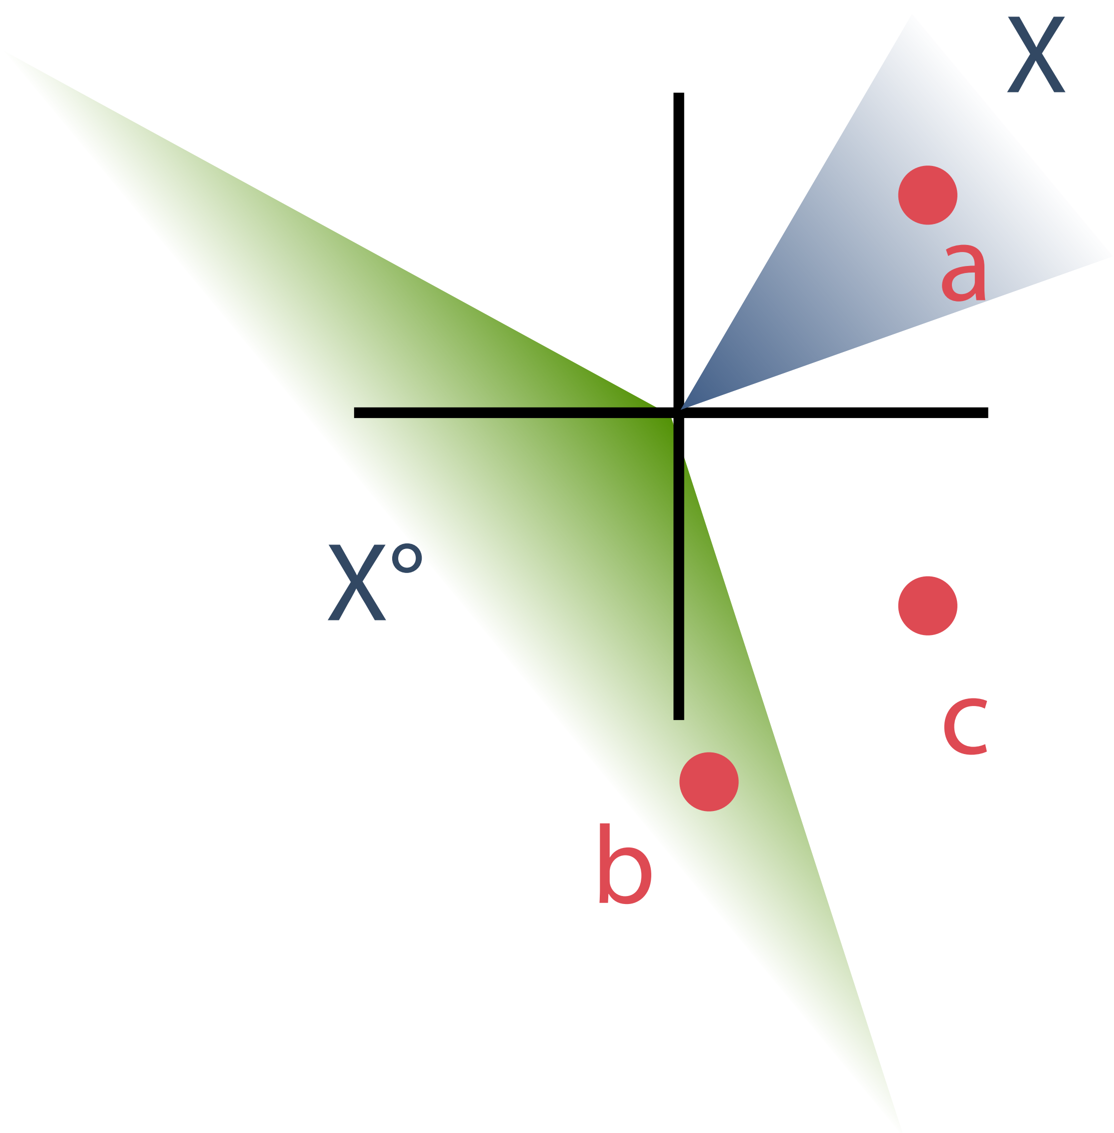

let $C^I = X\ldots$

$C(a)$

$\neg C(b)$

$(C \sqcup \neg C)(c)$

Example

- consider simple ontology $C \sqsubseteq D$
  - $C^I = \{(x,y) \in \mathbb{R}^2 \mid x \geq 0\}, (\neg C)^I = \{(x,y) \in \mathbb{R}^2 \mid x \leq 0\}$
  - $D^I = \{(x,y) \in \mathbb{R}^2 \mid x, y \geq 0\}, (\neg D)^I = \{\ (x,y) \in \mathbb{R}^n \mid x, y \leq 0\}$

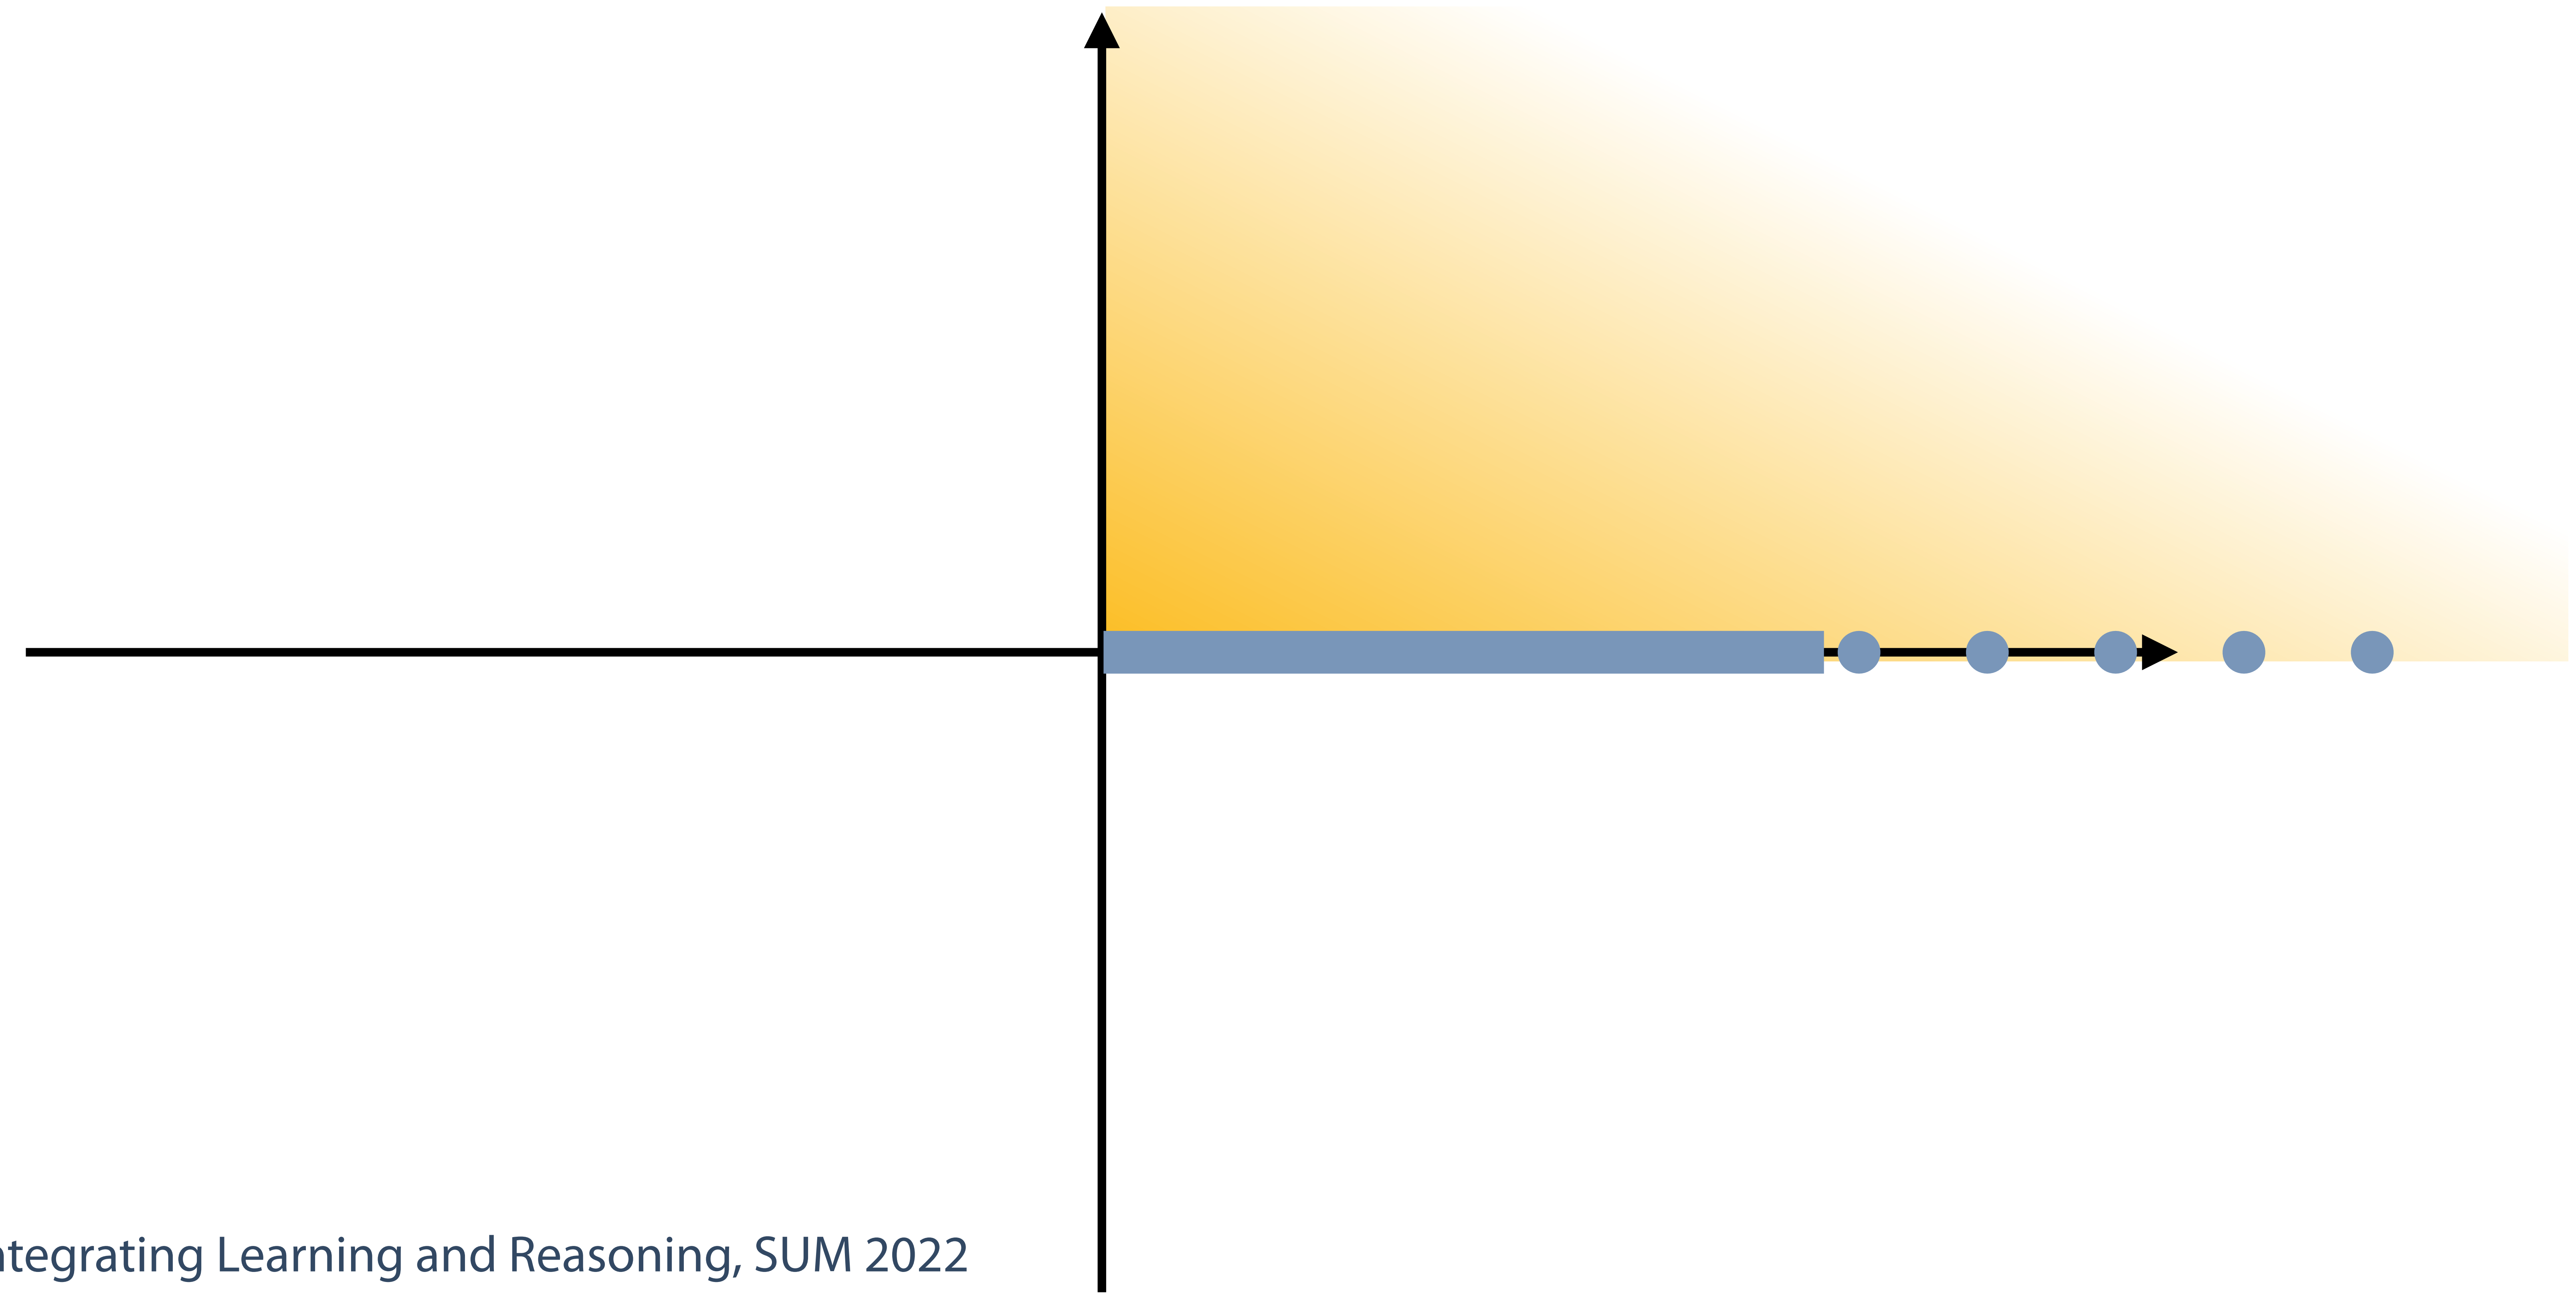

Example

- consider simple ontology $C \sqsubseteq D$
  - $C^I = \{(x, y) \in \mathbb{R}^2 | x \geq 0\}, (\neg C)^I = \{(x, y) \in \mathbb{R}^2 | x \leq 0\}$
  - $D^I = \{(x, y) \in \mathbb{R}^2 | x, y \geq 0\}, (\neg D)^I = \{\ (x, y) \in \mathbb{R}^n | x, y \leq 0\}$

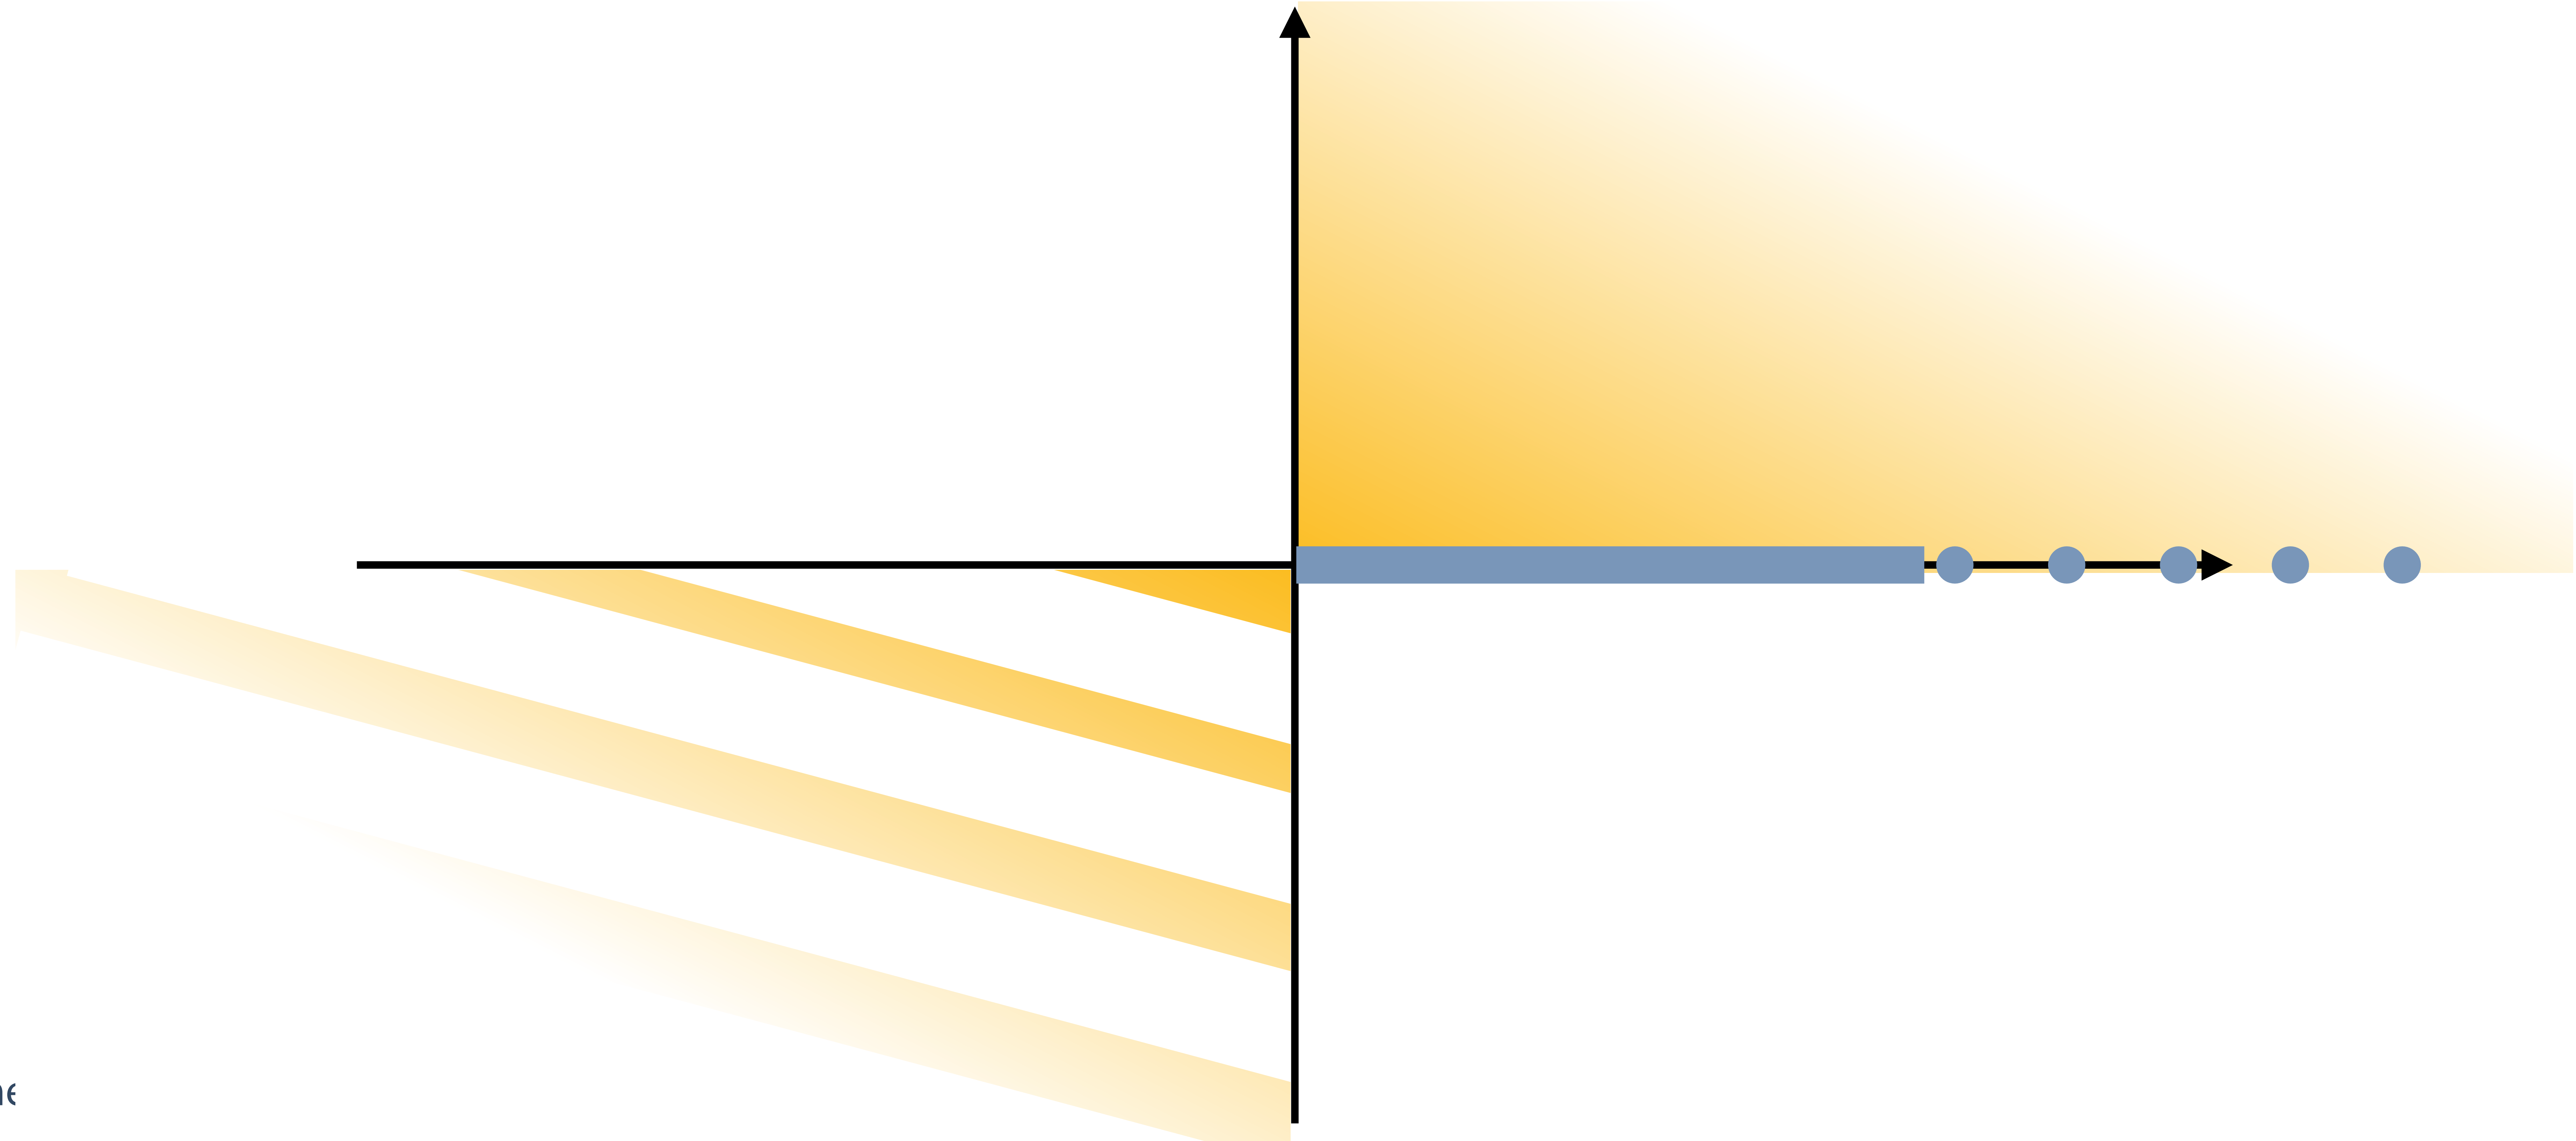

Example

- consider simple ontology $C \sqsubseteq D$
  - $C^I = \{(x,y) \in \mathbb{R}^2 \mid x \geq 0\}, (\neg C)^I = \{(x,y) \in \mathbb{R}^2 \mid x \leq 0\}$
  - $D^I = \{(x,y) \in \mathbb{R}^2 \mid x, y \geq 0\}, (\neg D)^I = \{\ (x,y) \in \mathbb{R}^n \mid x, y \leq 0\}$

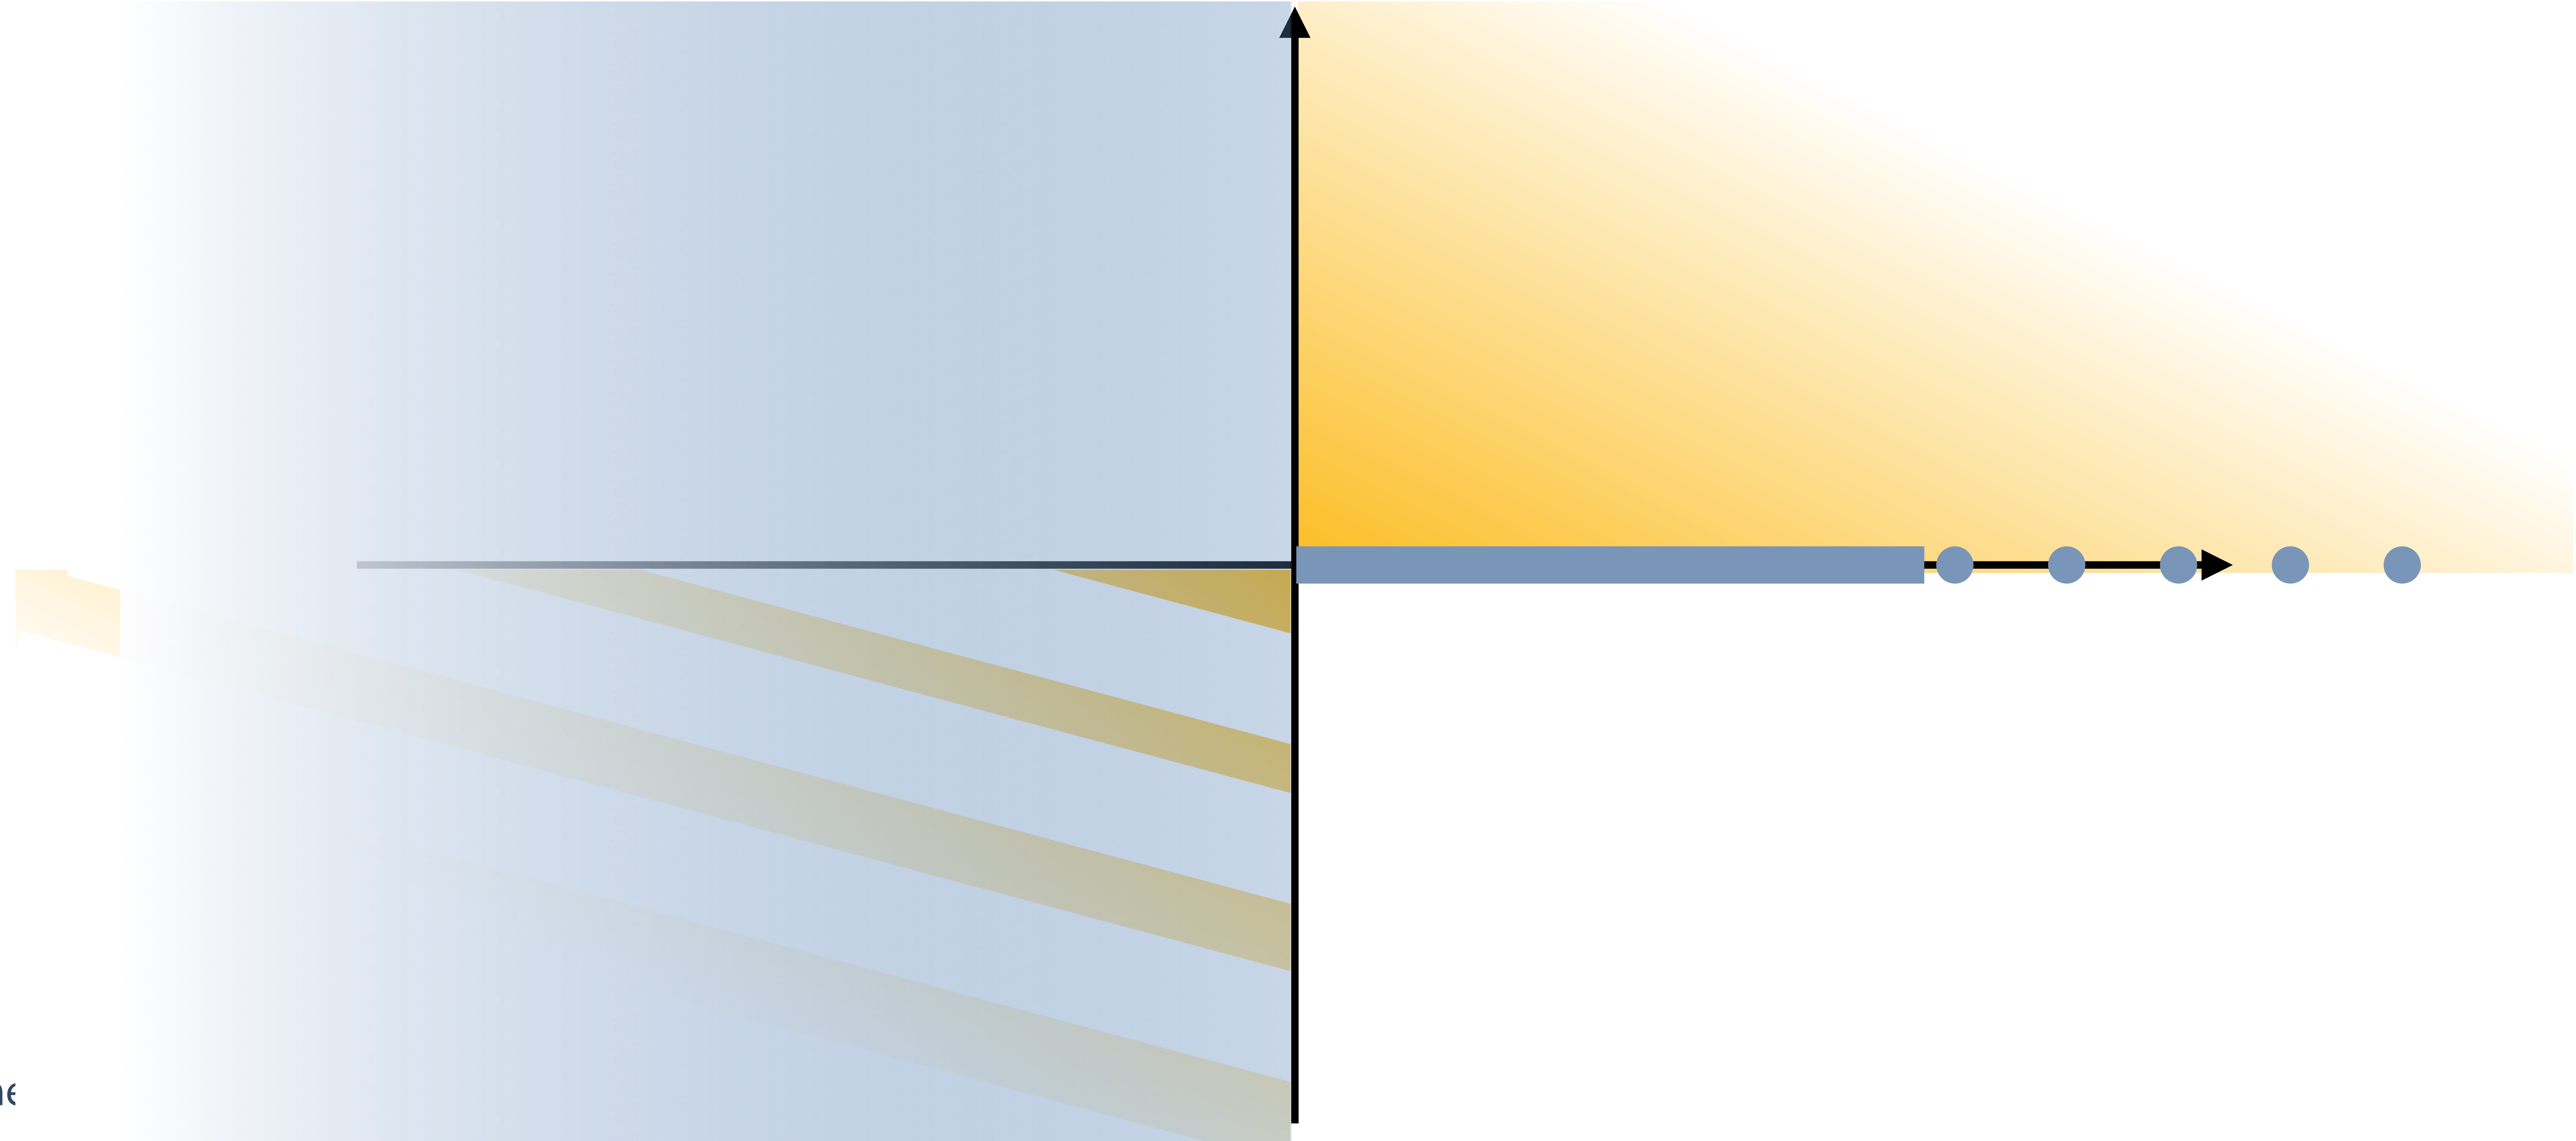

Example

- consider simple ontology $C \sqsubseteq D$
  - $C^I = \{(x,y) \in \mathbb{R}^2 \mid x \geq 0\}, (\neg C)^I = \{(x,y) \in \mathbb{R}^2 \mid x \leq 0\}$
  - $D^I = \{(x,y) \in \mathbb{R}^2 \mid x,y \geq 0\}, (\neg D)^I = \{\ (x,y) \in \mathbb{R}^n \mid x,y \leq 0\}$

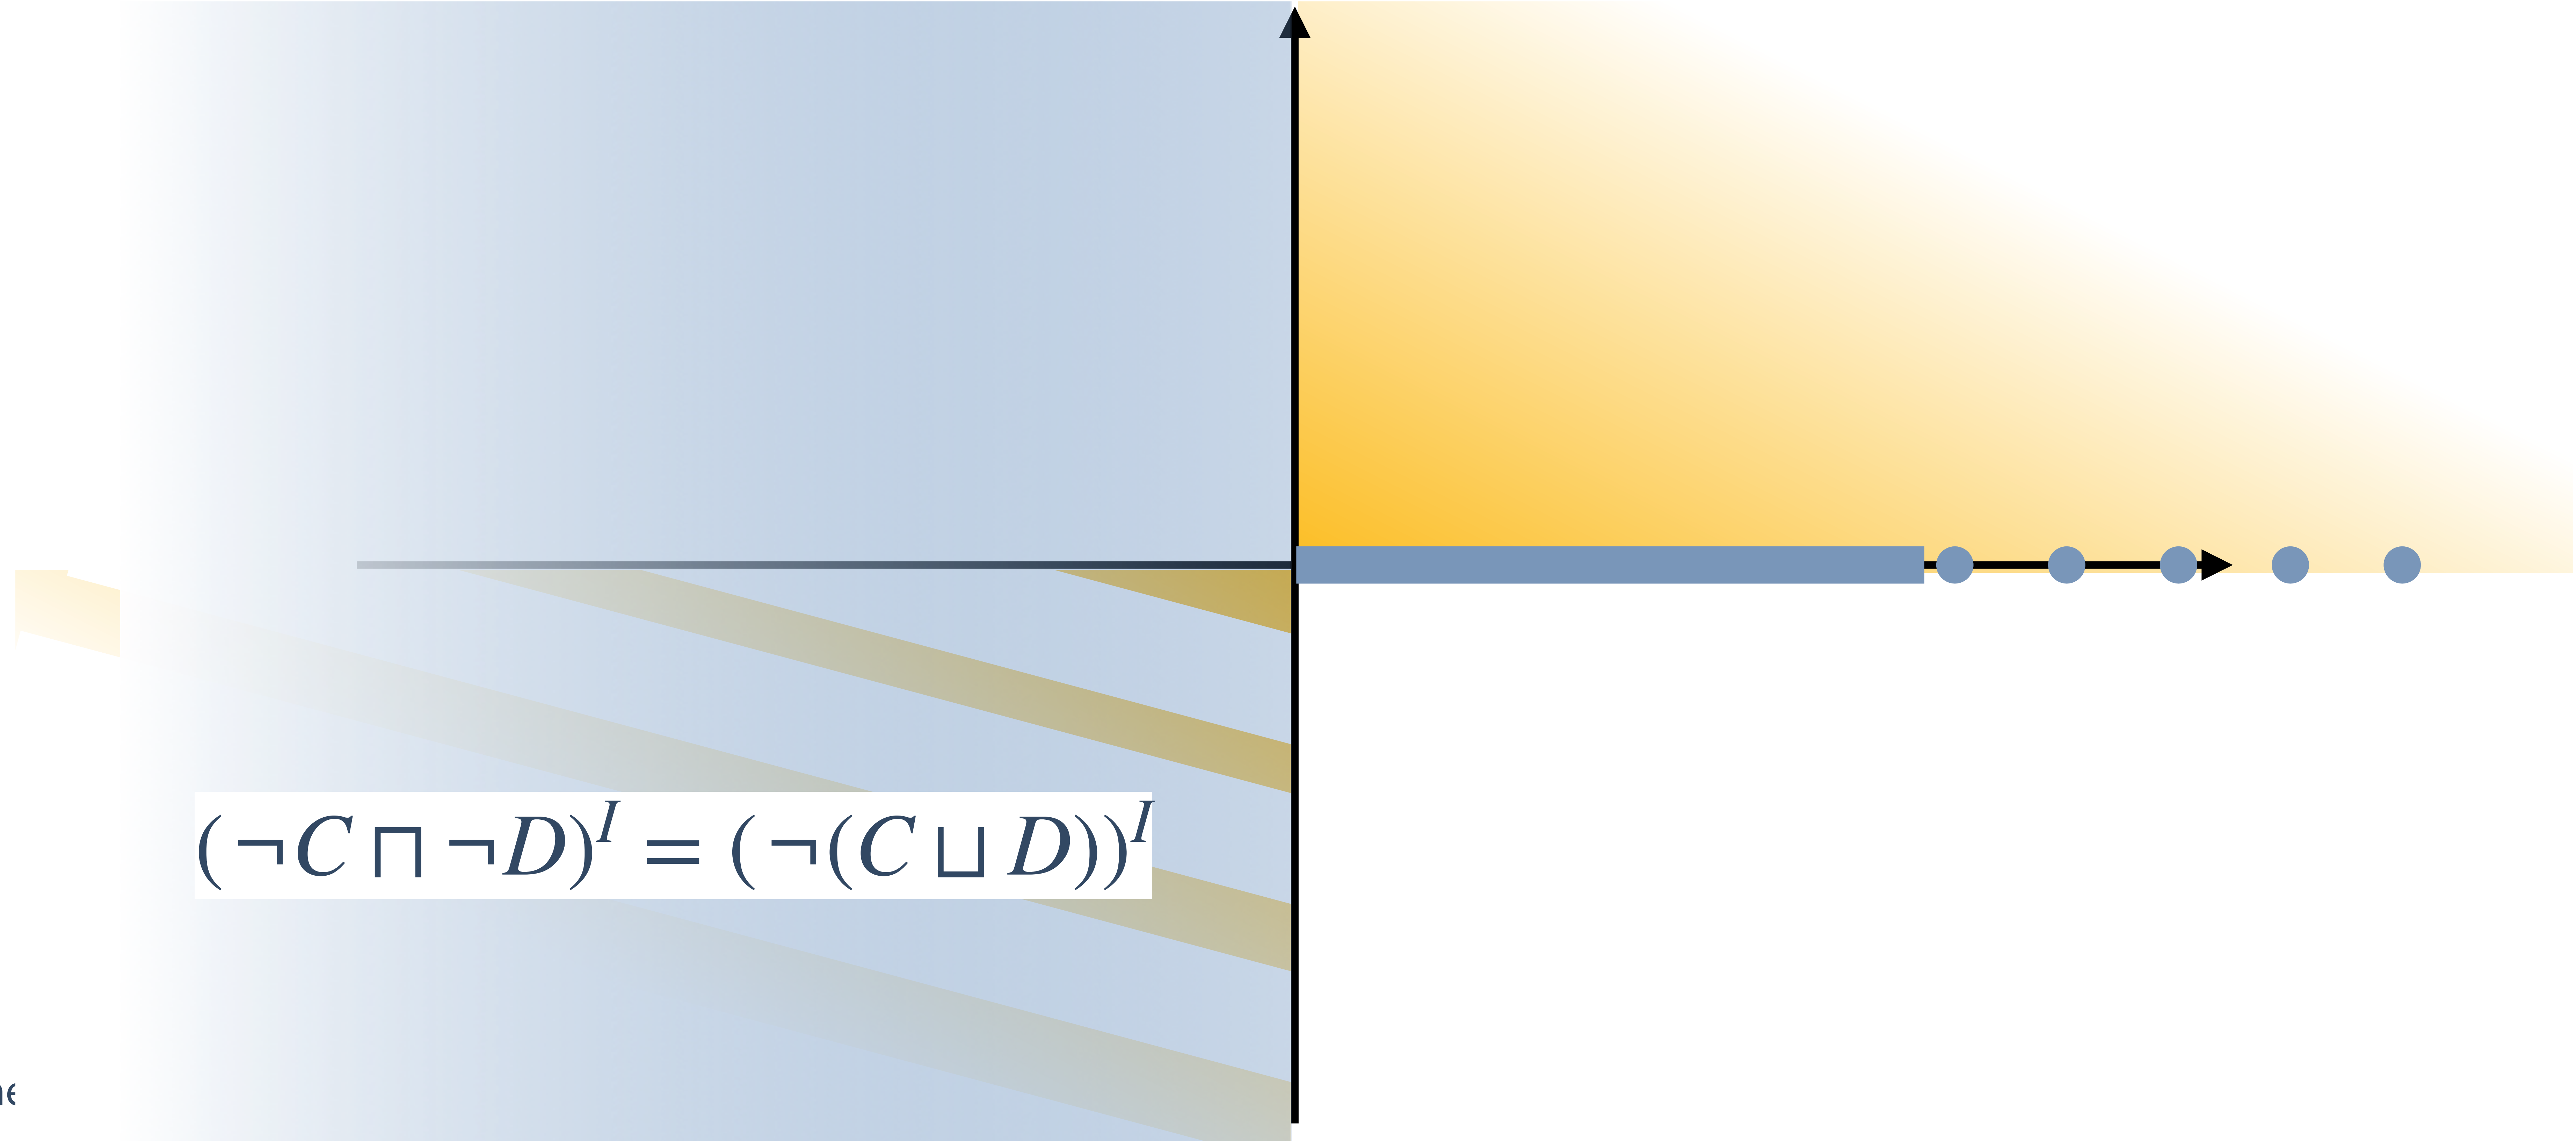

Geometrico-Algebraic Properties I

- using set intersection as conjunction, polarity as negation, and De-Morgan for defining disjunction, we arrive at an **algebra of cones**
- negation is not classical in the sense $x \in C^I$ or $x \in (\neg C)^I$
  - still it holds $(C \sqcup \neg C) \equiv \top$
  - weaker form, so-called **ortho-negation**
  - allows form of uncertainty to be captured in a geometric model
  - **example:** $\mathrm{Anmimal}(x)$, but neither $\mathrm{Mammal}(x)$ nor $\neg\mathrm{Mammal}(x)$

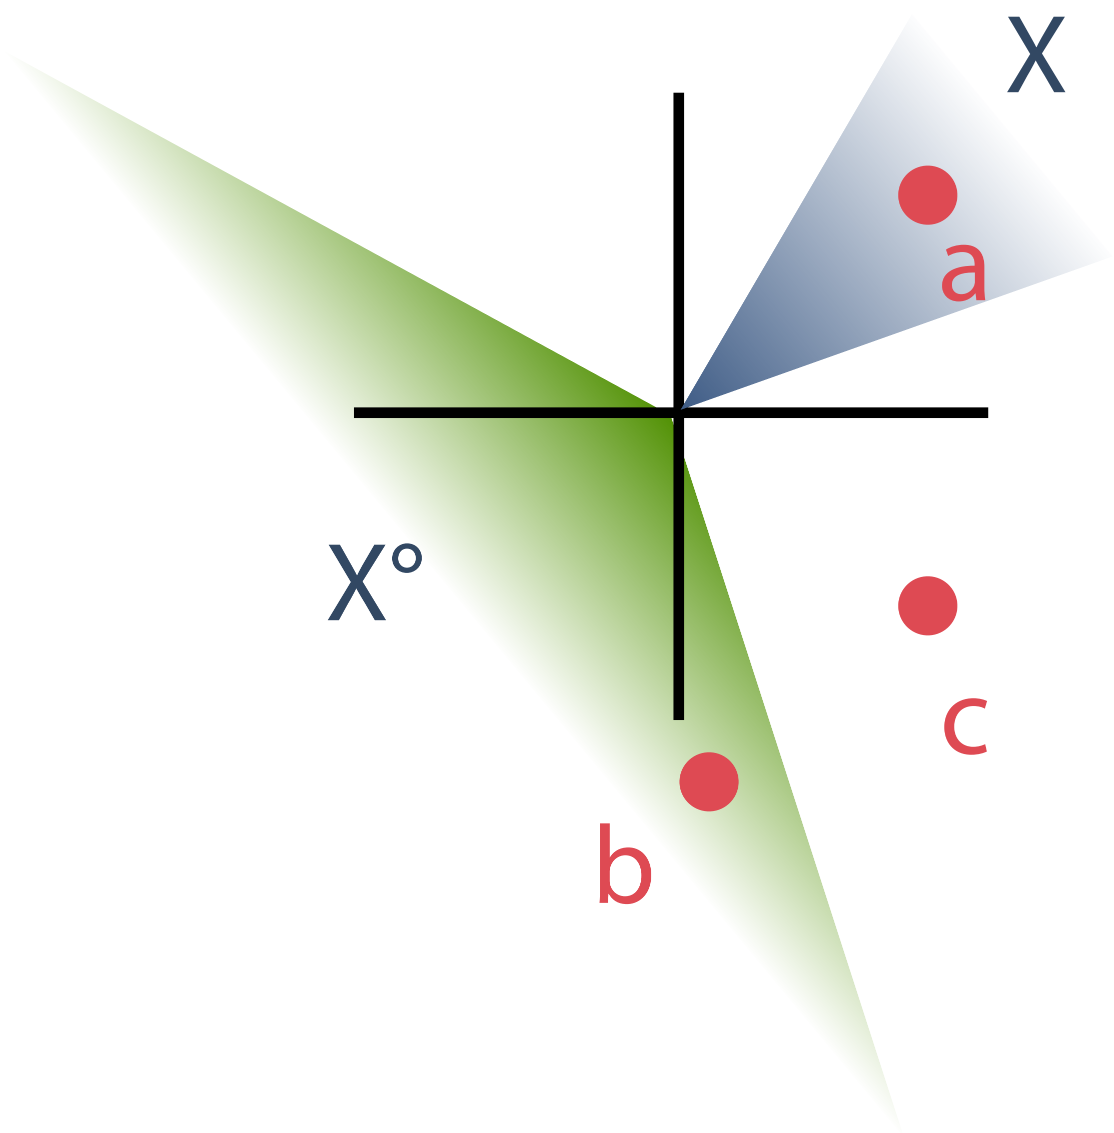

let $C^I = X\ldots$

$C(a)$

$\neg C(b)$

$(C \sqcup \neg C)(c)$

Geometrico-Algebraic Properties II

- distributivity not satisfied by *arbitrary* cones
  - $(A \sqcap B) \sqcup C \neq (A \sqcup C) \sqcap (B \sqcup C)$
- bug or feature? We opt for **bug**…
  - indeed a debatable property
  - **examples**: recall talk by Gabriele Kern-Isberner!
- one possible solution: restrict family of cones

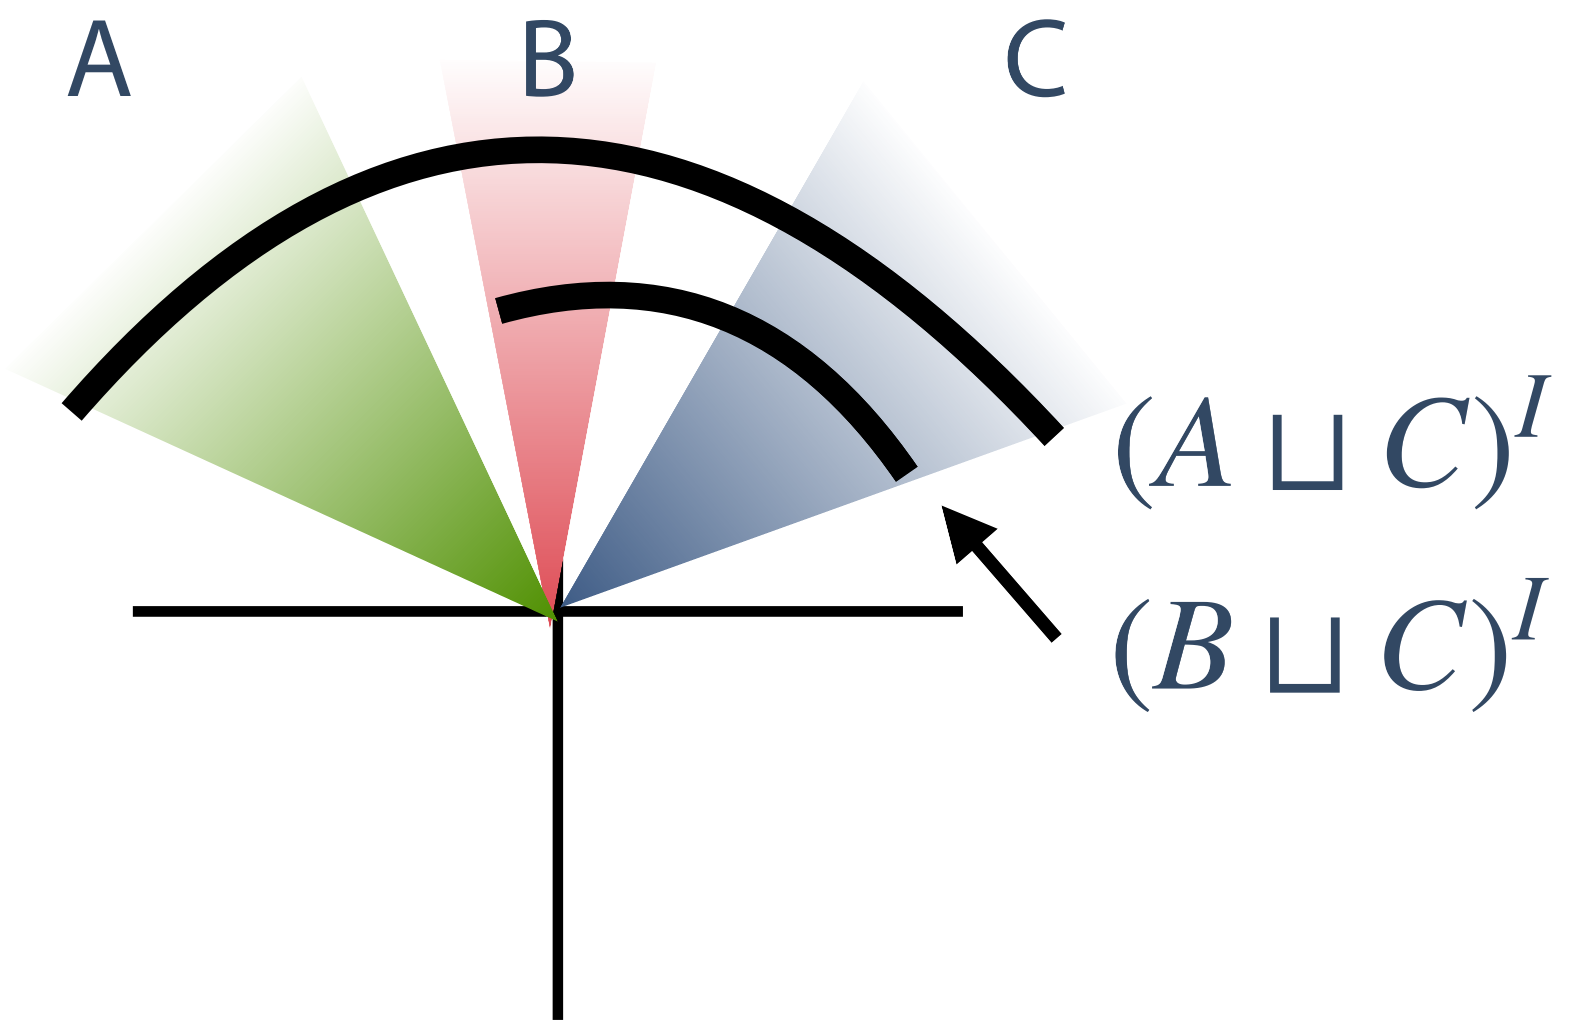

$$(A \sqcap B)^I = \{0\}$$

$$((A \sqcap B) \sqcup C)^I = C^I$$

$$B^I \sqsubset (A \sqcup C)^I, B^I \sqsubset (B \sqcup C)^I$$

$$((A \sqcup C) \sqcap (B \sqcup C))^I = C^I$$

Axis-Aligned Cones

- Definition: X is called an axis-aligned cone (al-cones) in $\mathbb{R}^n$ iff $X = X_1 \times X_2 \times \cdots \times X_n$ with $X_i \in \{\mathbb{R}, \mathbb{R}_+, \mathbb{R}_-, \{0\}\}$
  - examples for ontology $C \sqsubseteq D$ already featured al-cones

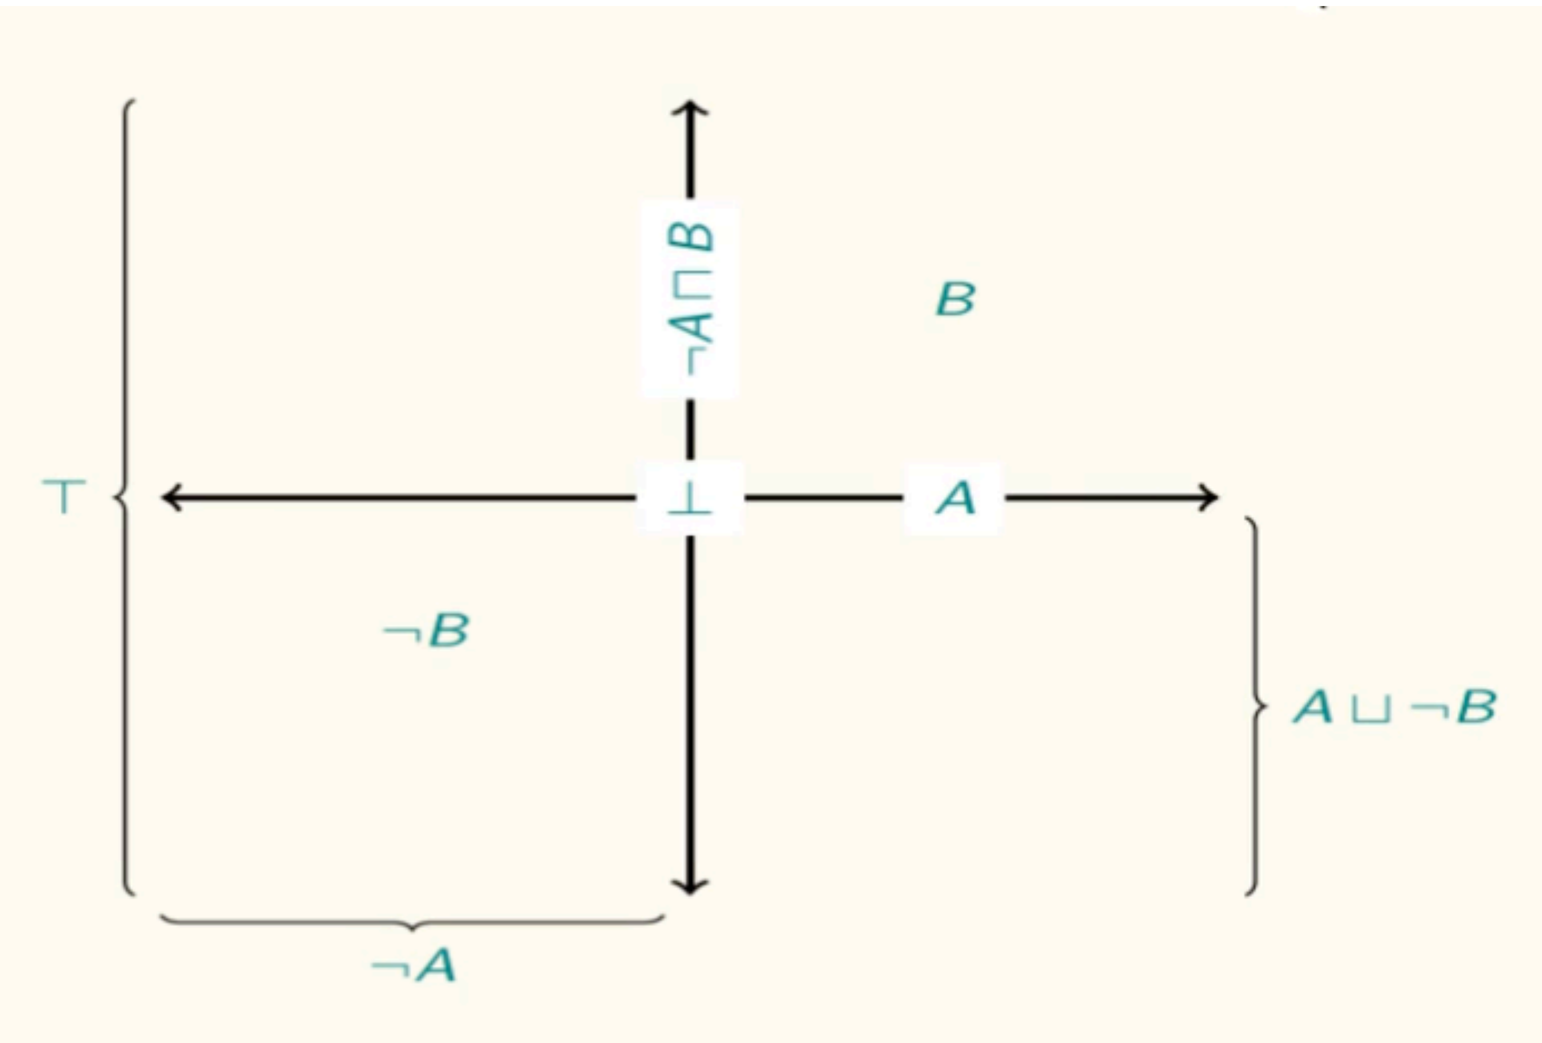

- nomen est omen: axis-aligned cones (**al-c**ones) provide a geometric model for $\mathscr{ALC}$
  - **proposition** (Özçep et al, 2020): A Boolean $\mathscr{ALC}$ ontology is satisfiable iff it is satisfiable with a faithful al-cone model
  - **faithfulness**: $a^I \in C^I$ if and only if ontology entails $C(a)$.

AL-Cone Models for Full $\mathscr{ALC}$

- **idea**: interpret relations classically as subsets of cartesian product, i.e., $R \subseteq \mathbb{R}^n \times \mathbb{R}^n$
  - for relating concepts which are al-cones, we have $R \subseteq D \times D$ with $D = D_1 \times \cdots \times D_n, D_i \in \{\mathbb{R}, \mathbb{R}_+, \mathbb{R}_+\{0\}\}$
- **problem**: $\mathscr{ALC}$ models may require infinite chains of concepts
  - **example** (Baader & Küsters 2006): $\mathrm{loves(narcis, narcis), Vain(narcis)}$

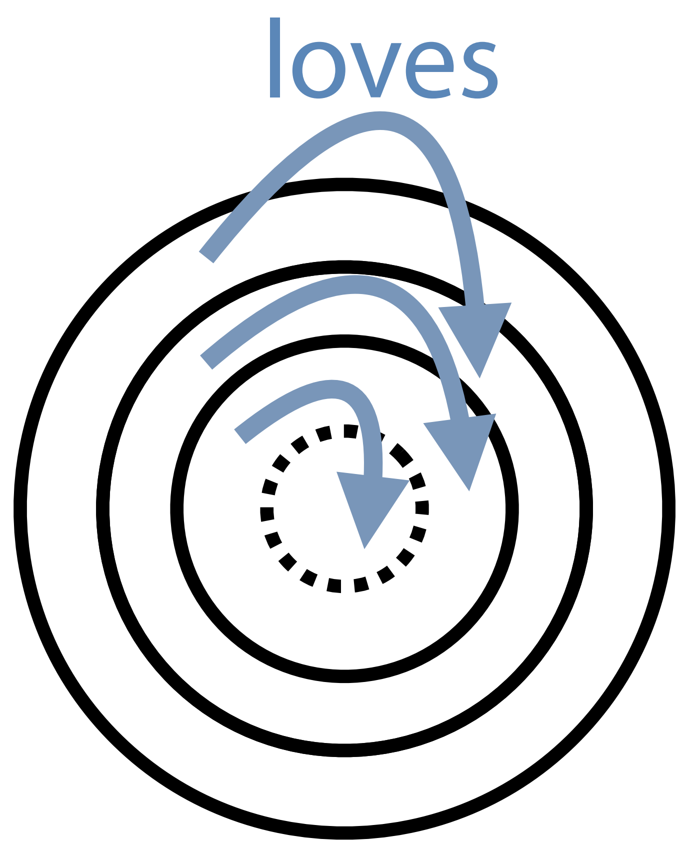

- **idea**: approximation with bound on quantifier rank
  - rationale: when querying a model, we may assume a maximum nesting of quantifiers
- **proposition** (Özçep et al., 2020): $\mathscr{ALC}$ ontologies with fixed quantifier rank k are satisfiable if and only if they are satisfiable with a faithful al-cones model
- disadvantage: relations are not first-order members entities of an embedding

(Desired) Limitations of Expressivity

- there are more general logics than can be modelled by cones
  - limitations are not necessarily a drawback
- cones constitute some restriction of Goldblatt's minimal Orthologic Omin
  - polarity satisfies orthonormality, cones constitute an **orthlattice**
- **example**: assume logic of cones does not allow MC_8 to be represented
  - sample group (b) shall be representative (d)
  - distinguishing $b^\perp$ and $d^\perp$ not sensible
- ▸ evaluate adequacy of logical commitments!

Axioms

$$A \vdash A \qquad A \,\&\, B \vdash A \qquad A \,\&\, B \vdash B \qquad A \dashv\vdash \sim\sim A$$
$$A \,\&\, \sim A \vdash B \qquad A \vee B \dashv\vdash \sim(\sim A \,\&\, \sim B)$$

Rules

$$\frac{A \vdash B,\ B \vdash C}{A \vdash C} \qquad \frac{A \vdash B,\ A \vdash C}{A \vdash B \,\&\, C} \qquad \frac{A \vdash B}{\sim B \vdash \sim A}$$

Omin

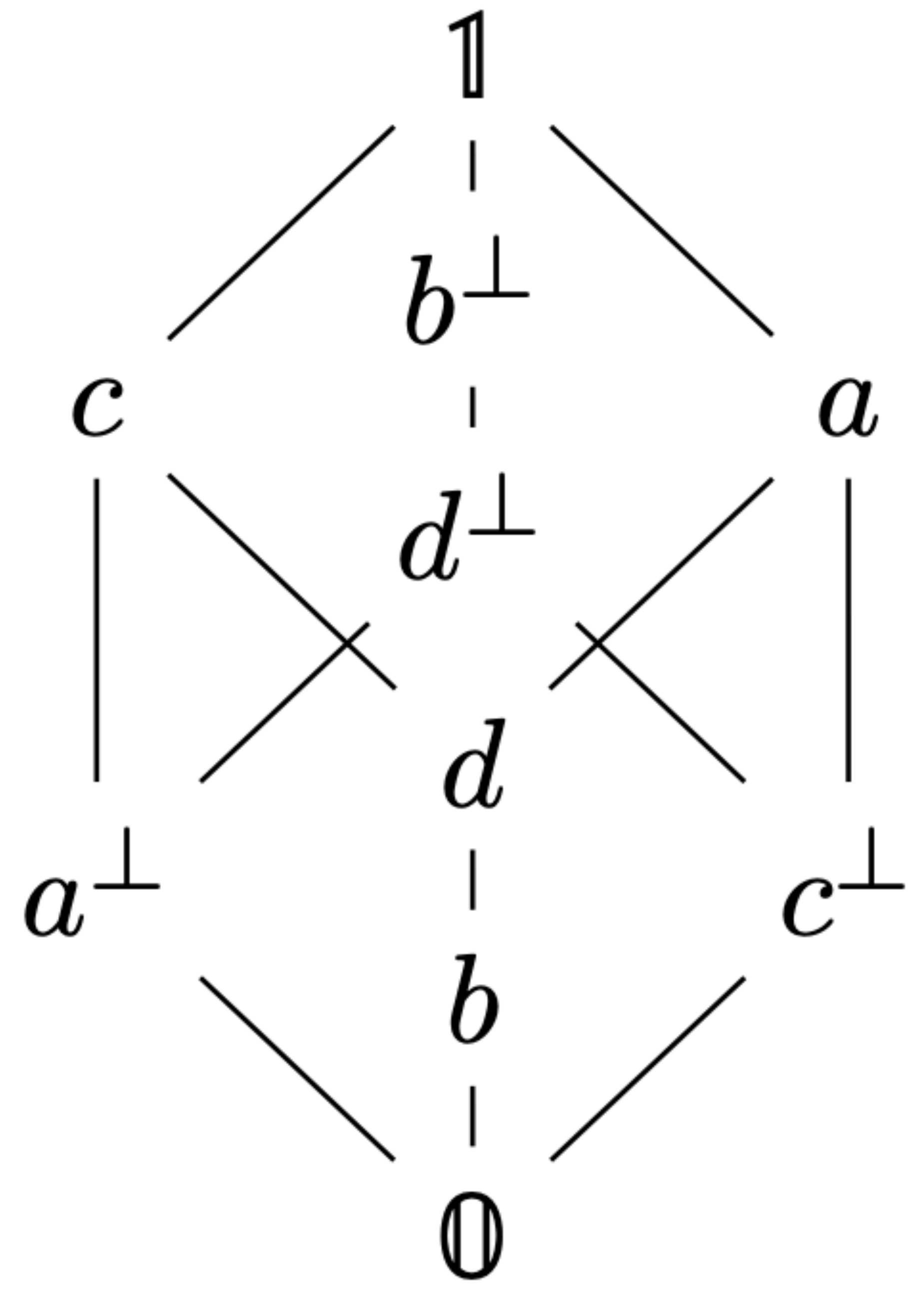

MC_8 ortholattice

(Desired) Limitations of Expressivity

- there are more general logics than can be modelled by cones
  - limitations are not necessarily a drawback
- cones constitute some restriction of Goldblatt's minimal Orthologic Omin
  - polarity satisfies orthonormality, cones constitute an **ortholattice**
- **example**: assume logic of cones does not allow MC_8 to be represented
  - sample group (b) shall be representative (d)
  - distinguishing $b^\perp$ and $d^\perp$ not sensible

▸ evaluate adequacy of logical commitments!

Axioms

$$A \vdash A \quad A \mathbin{\&} B \vdash A \quad A \mathbin{\&} B \vdash B \quad A \dashv\vdash {\sim}{\sim} A$$

$$A \mathbin{\&} {\sim} A \vdash B \quad A \vee B \dashv\vdash {\sim}({\sim} A \mathbin{\&} {\sim} B)$$

Rules

$$\frac{A \vdash B,\ B \vdash C}{A \vdash C} \qquad \frac{A \vdash B,\ A \vdash C}{A \vdash B \mathbin{\&} C} \qquad \frac{A \vdash B}{{\sim} B \vdash {\sim} A}$$

Omin

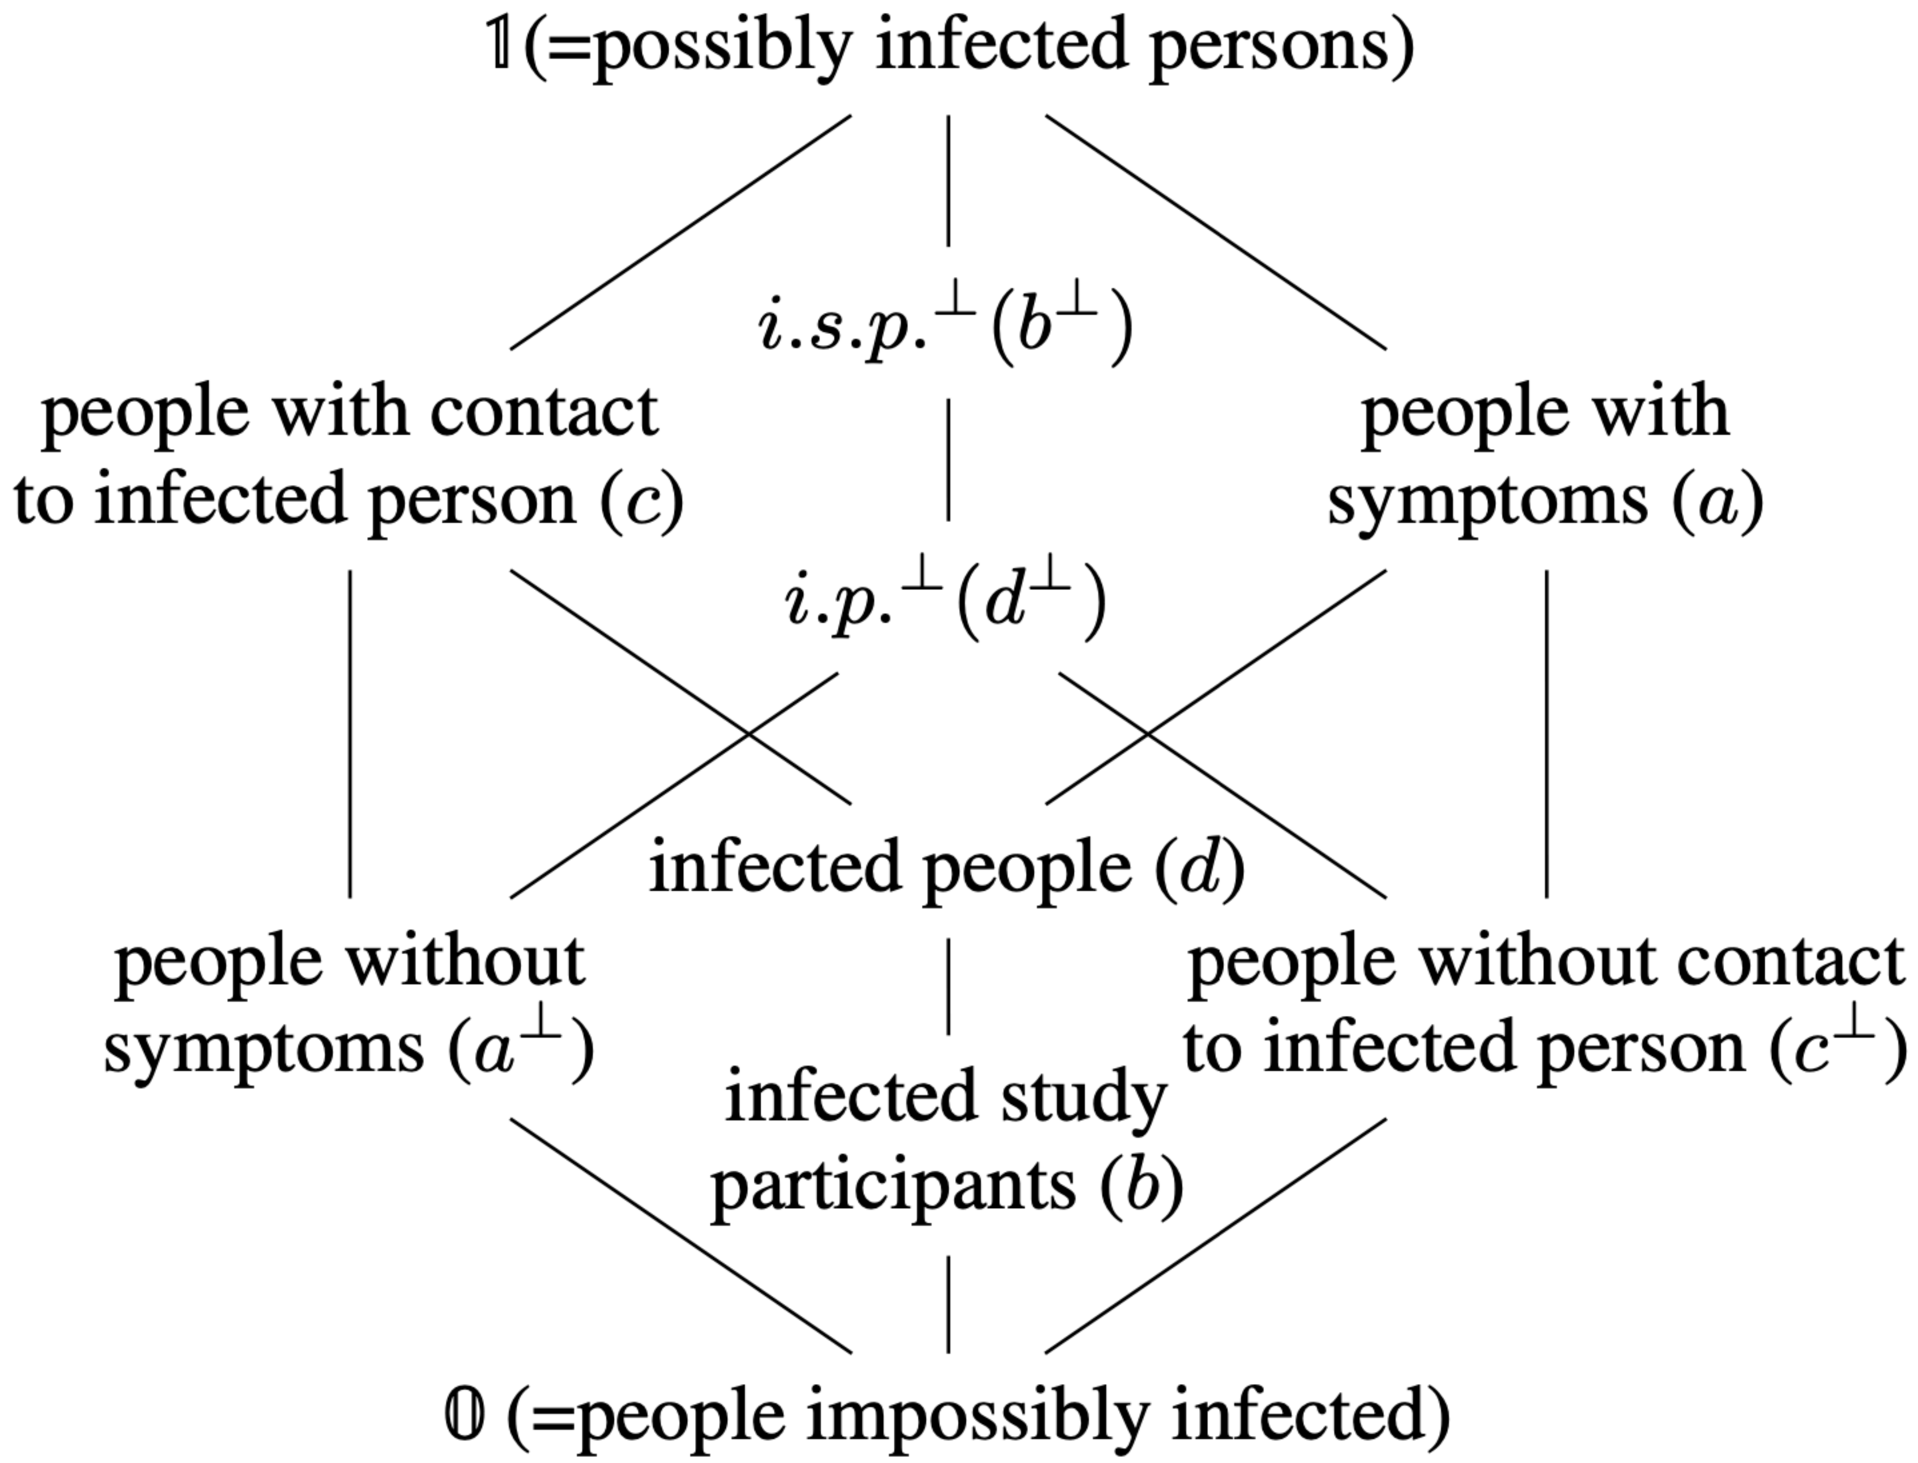

(Desired) Limitations of Expressivity

- there are more general logics than can be modelled by cones
  - limitations are not necessarily a drawback
- cones constitute some restriction of Goldblatt's minimal Orthologic Omin
  - polarity satisfies orthonormality, cones constitute an **ortholattice**
- **example**: assume logic of cones does not allow MC8 to be represented
  - sample group (b) shall be representative (d)
  - distinguishing $b^\perp$ and $d^\perp$ not sensible
- ▸ evaluate adequacy of logical commitments!

Axioms

$$A \vdash A \qquad A \,\&\, B \vdash A \qquad A \,\&\, B \vdash B \qquad A \dashv\vdash \sim\sim A$$

$$A \,\&\, \sim A \vdash B \qquad A \vee B \dashv\vdash \sim(\sim A \,\&\, \sim B)$$

Rules

$$\frac{A \vdash B,\ B \vdash C}{A \vdash C} \qquad \frac{A \vdash B,\ A \vdash C}{A \vdash B \,\&\, C} \qquad \frac{A \vdash B}{\sim B \vdash \sim A}$$

Omin

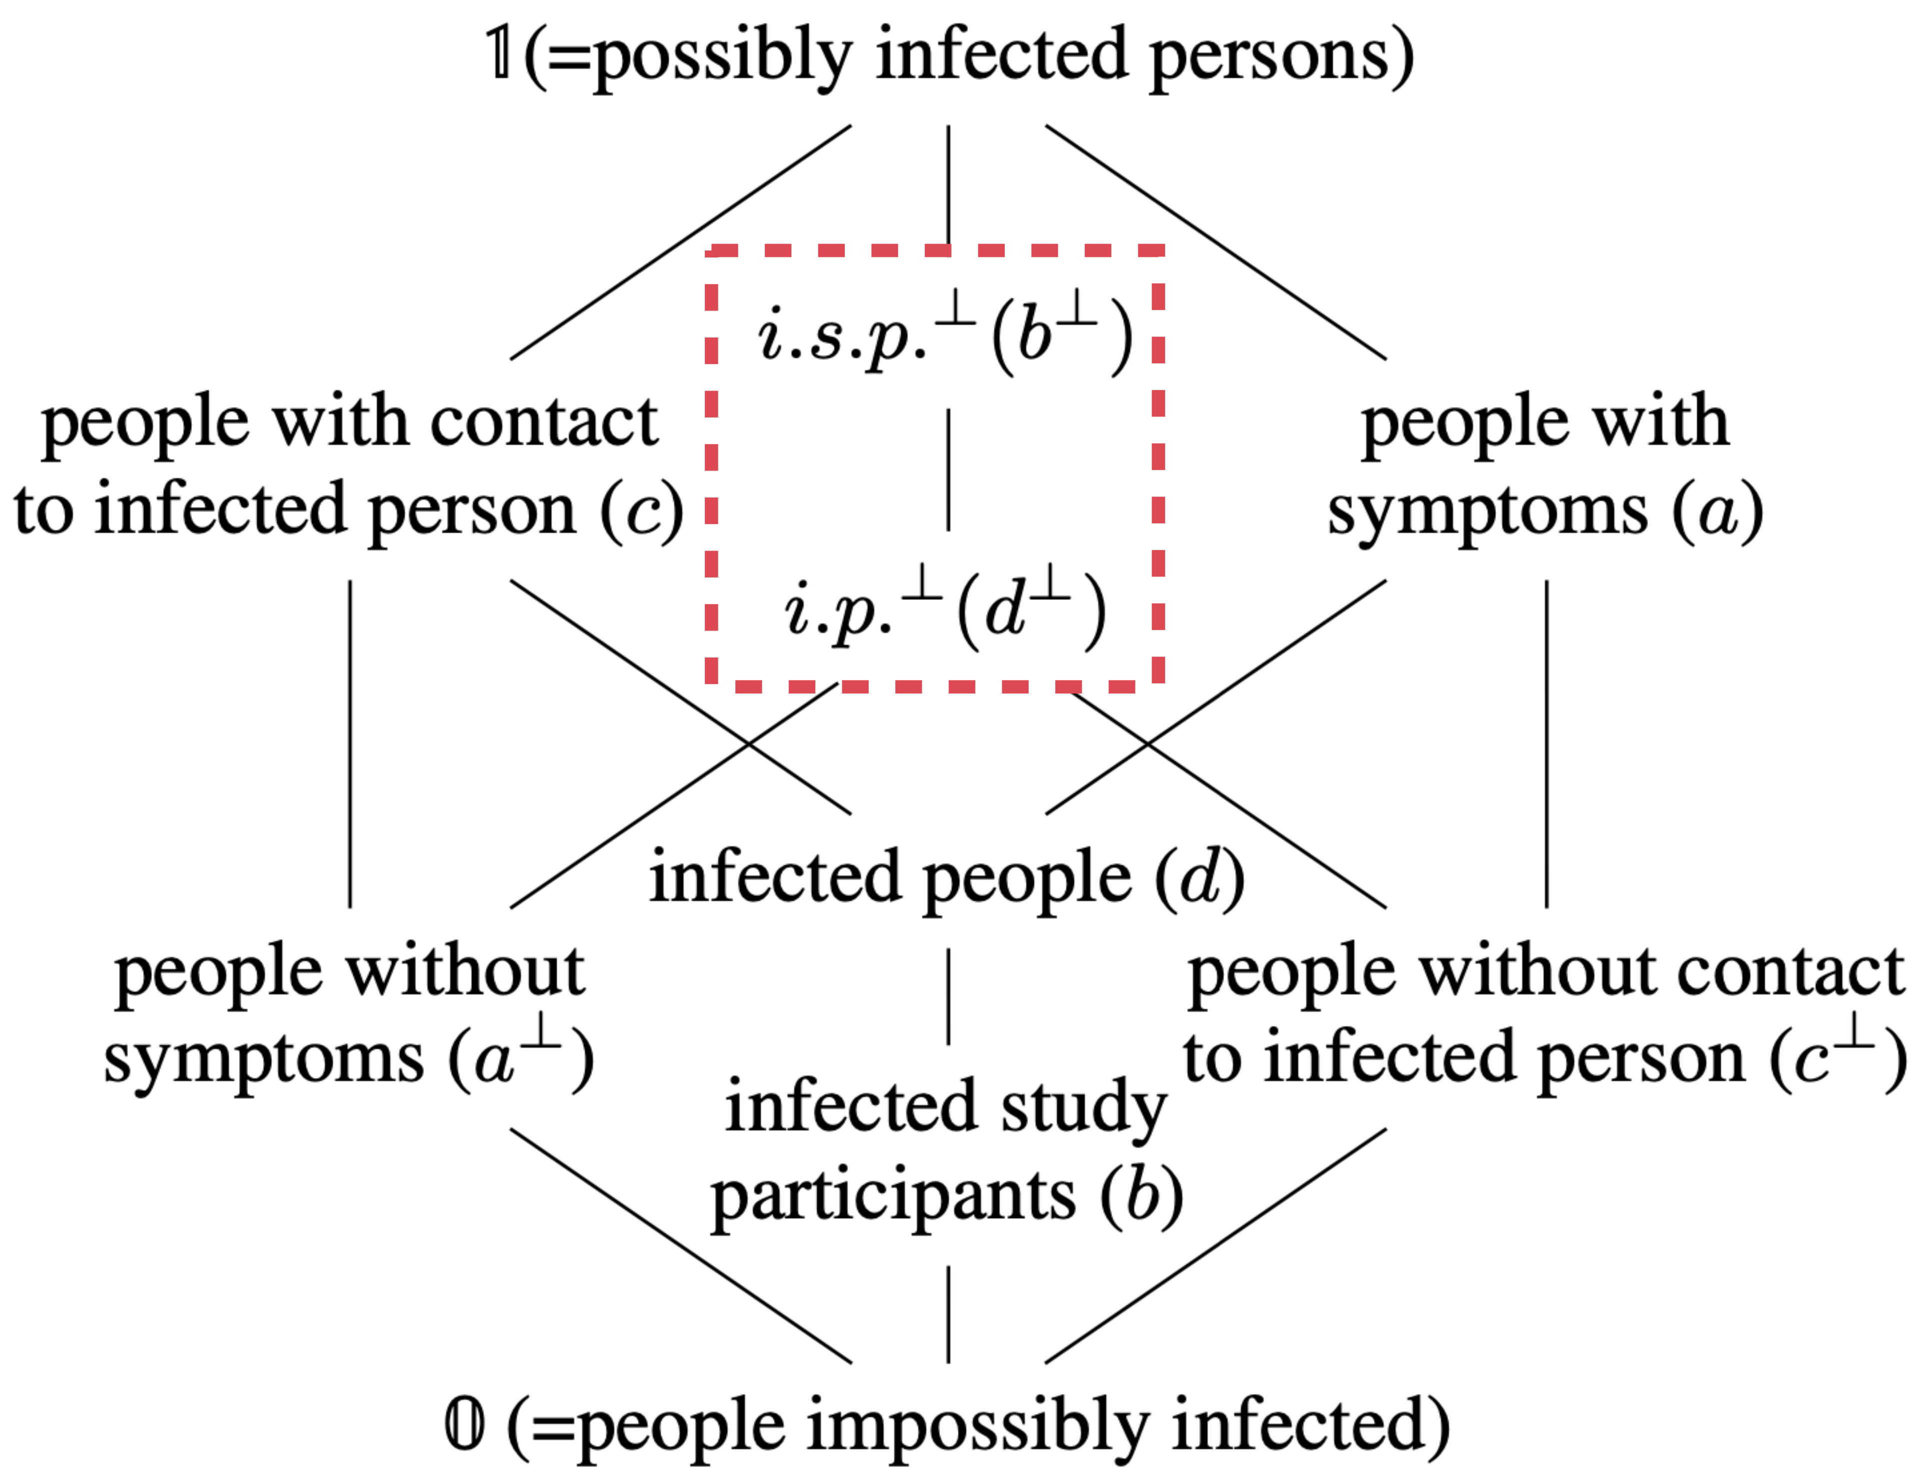

Learning with Cones: First Results

- general idea: given background ontology, learn embedding
- learning al-cone embedding: rather search than ML
  - component-wise discrete
- learning arbitrary cones can be achieved by SVMs
  - cones defined as intersection of hyperplanes
  - allows kernel trick to be used
- example with AWA2 dataset ("animals with attributes", Zero-Shot Learning) (Leemhuis et al. 2022)
  - problem: current datasets do not involve negation

baseline may guess labels

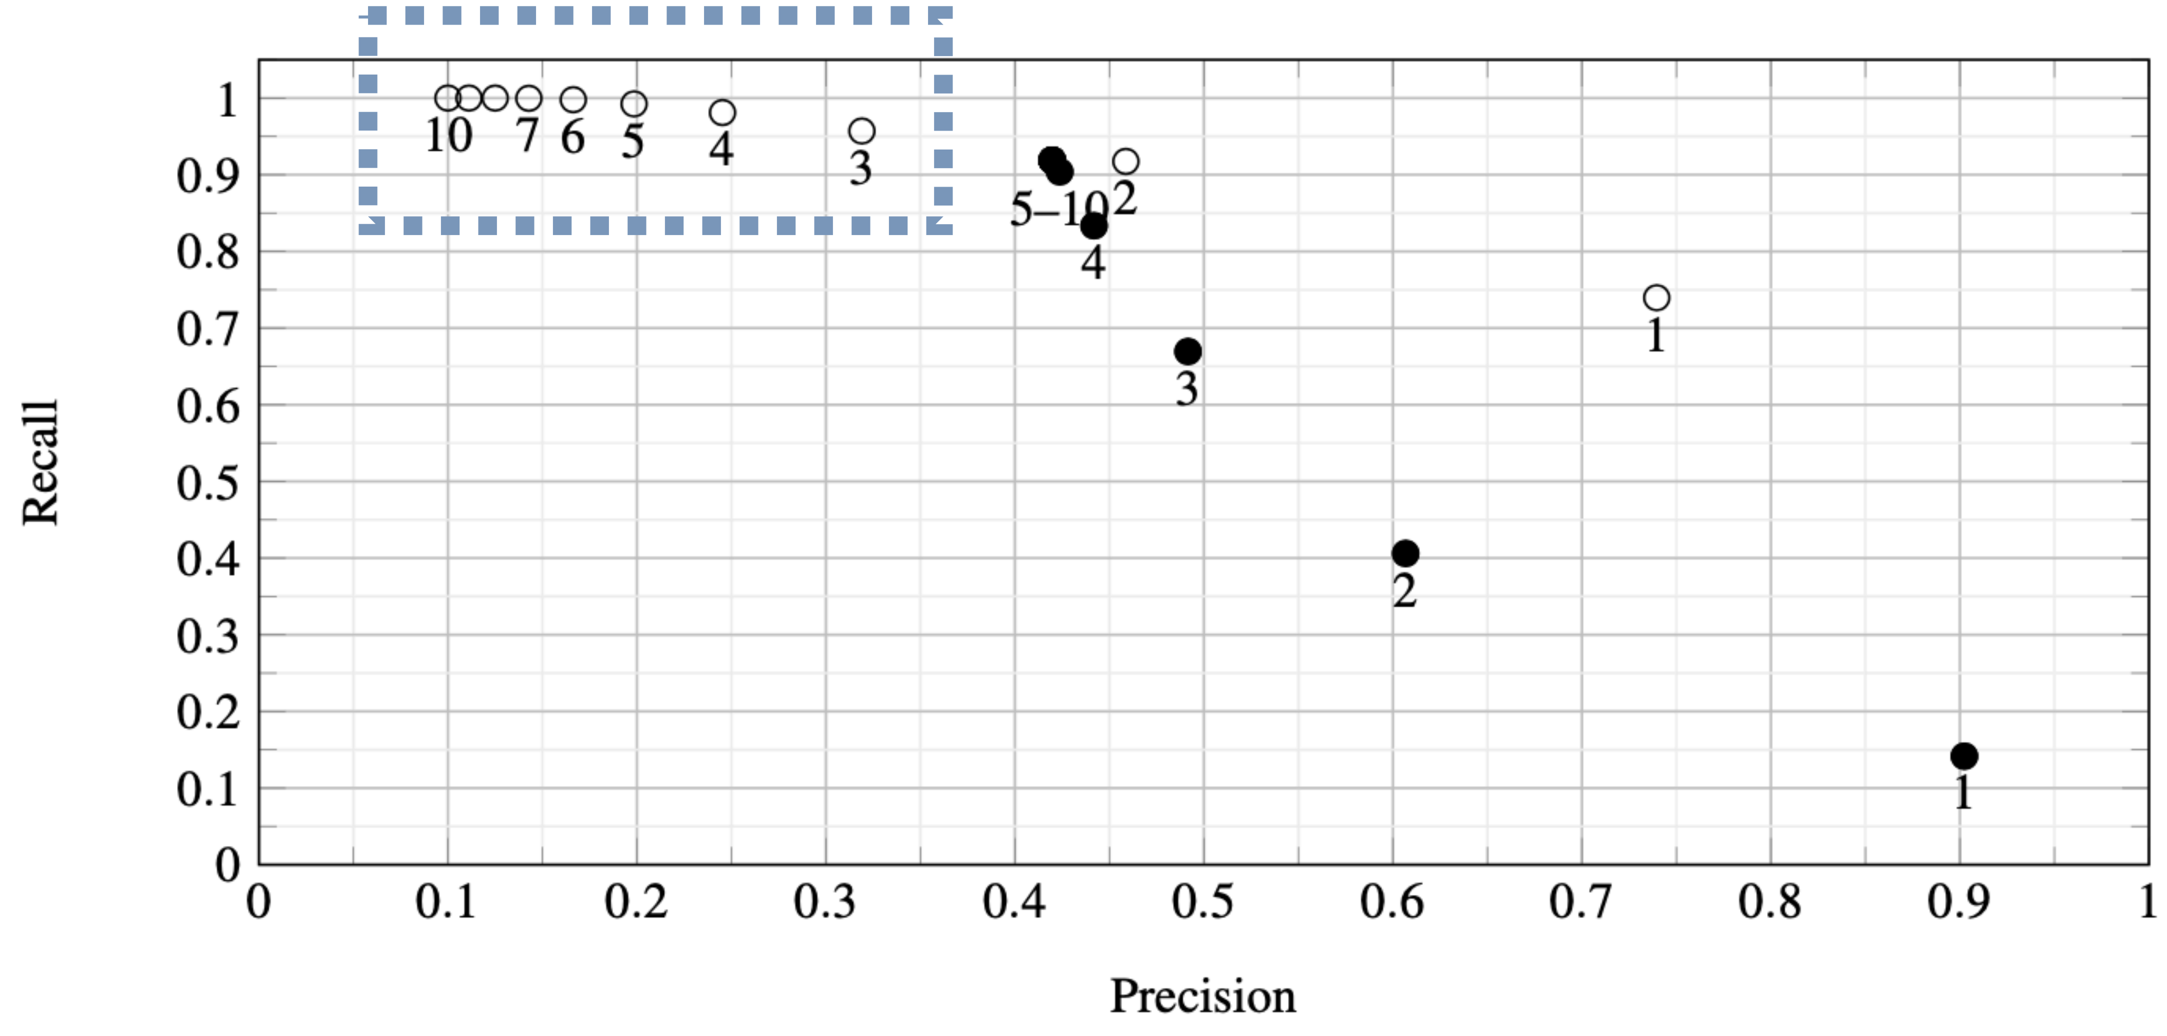

cone model gears towards higher precision
- cannot learn (wrt. ontology) wrong

Reification

- **idea**: treat relations geometrically like concepts (Leemhuis et al, 2022)
  - functions map to domain and co-domain
  - let $(\exists R\,.\,C)^I = \pi_{1,R}\left(\pi_{2,R}^{-1}(C^I) \cap H\right)$
- reification allows non-functional relations to be represented using two functions
  - assume $\pi_{i,R}$ to be projections, $\pi_{2,R}^{-1}$ can span a subspace
  - applicable to other KG embeddings!
- **question**: Will reification also lead to better performance of KG embeddings?

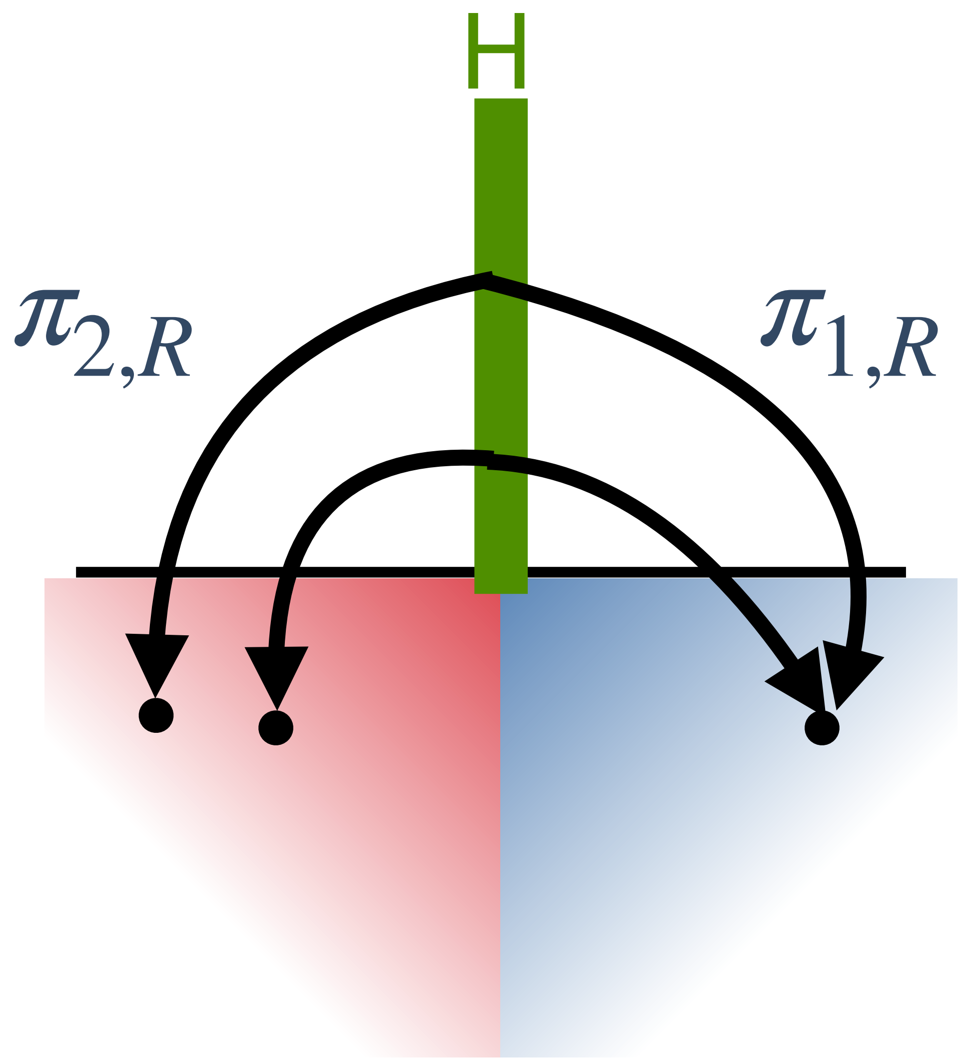

Summary

- knowledge graph embeddings connect machine learning and symbolic reasoning
  - semantics of embeddings not well-understood
  - classic embeddings mix uncertainty resulting from noisy data with uncertainty arising from poor semantic alignment
- geometric models can retain uncertainty in data
  - beyond prototypicality/likelihood
  - cones with polarity and intersection constitute algebra
  - al-cones as example for a model for description logic $\mathscr{ALC}$

```
  L E A R N I N G
+ R E A S O N I N G
-------------------
  E M B E D D I N G
```

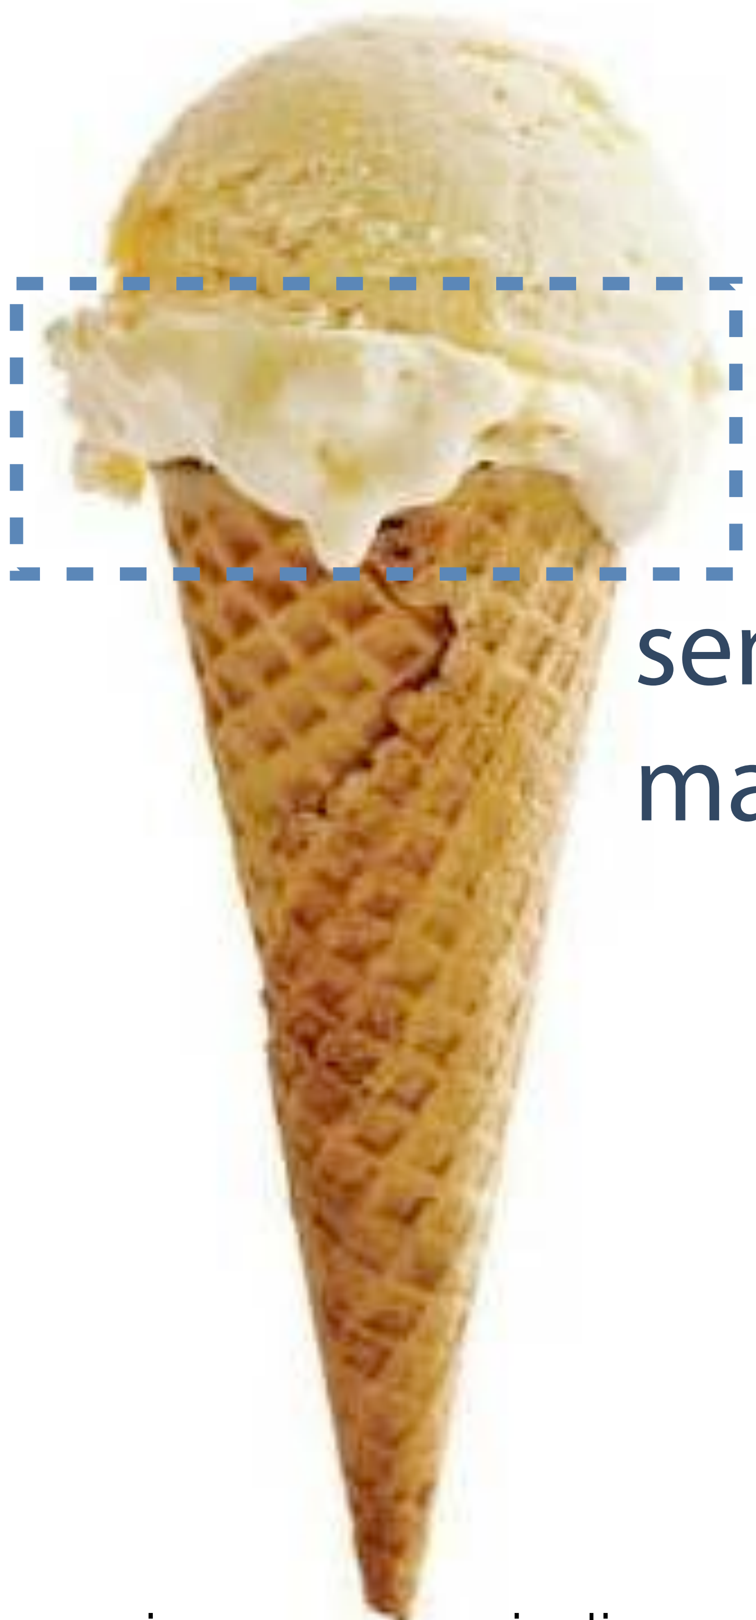

semantic alignment matters

image: www.indiamart.com

Conclusion and Outlook

- **geometric and logic commitments constitute important design decisions**
- interesting combinations of geometric models and concept languages can be found
  - (al-)cones may just be the beginning
  - semantically proper treatment of desired logic features is possible
  - find a good balance between feasibility of learning and expressivity of concept languages
- geometric models are still under-explored
  - and sometimes puzzling
- learning uncertain models: How can we gear learning to making concrete commitments?

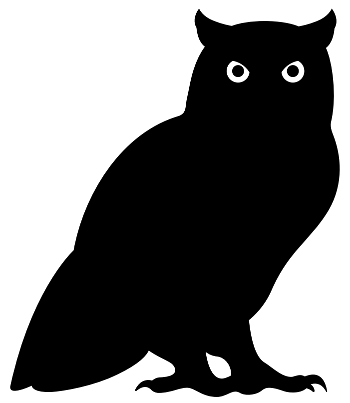 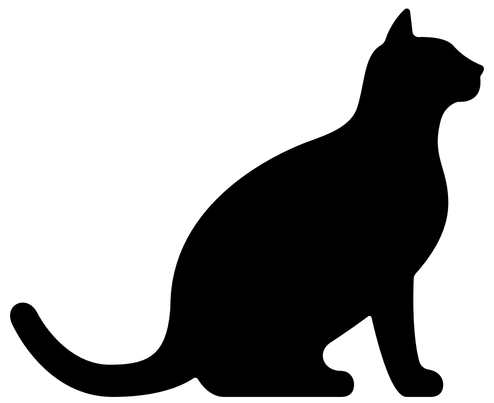 $\rightsquigarrow \top$

References

this presentation was based on…

- Ö. L. Özçep, M. Leemhuis, D. W. (2020). Cone Semantics for Logics with Negation, In Proc. of IJCAI
- M. Leemhuis, Ö. L. Özçep, D. W. (2022). Learning with Cone-Based Geometric Models and Orthologics, *Annals of Mathematics and Artificial Intelligence*, 90:7–9
- M. Leemhuis, Ö. L. Özçep, D. W. (2022). Knowledge Graph Embeddings with Ontologies: Reification for Representing Arbitrary Relations, In Proc. of KI: Advances in Artificial Intelligence, Springer

references

- F. Baader, R. Küsters (2002). Unification in a Description Logic with Inconsistency and Transitive Closure of Roles. Proc. of DL
- S. Badreddine, A. A. Garcez, L. Serafini, M. Spranger (2022). Logic tensor networks, AIJ, 303: 103649
- J. Bischofberger & M. Ragni (2020). Improving Cognitive Models for Syllogistic Reasoning. Proceedings of the 42th Annual Meeting of the Cognitive Science Society - Developing a Mind: Learning in Humans, Animals, and Machines, CogSci 2020
- Gärdenfors (2000). Conceptual Spaces: The Geometry of Thought, MIT press
- D. Garg, S. Ikbal, S. K. Srivastava, H. Vishwakarma, H. Karanam, L V. Subramaniam (2019). Quantum Embedding of Knowledge for Reasoning, Proc. of NeurIPS
- V. Gutiérrez-Basulto & S. Schockaert (2018). From knowledge graph embedding to ontology embedding? An analysis of the compatibility between vector space representations and rules, Proc. of KR
- S. M. Kazemi, D. Poole (2018). SimplE Embedding for Link Prediction in Knowledge Graphs, Proc. of NeuIPS
- M. Kulmanov, W. Liu-Wei, Y. Yan, R. Hoehndorf (2019). EL Embeddings: Geometric Construction of Models for the Description Logic EL++
- R. Riegel, A. G. Gray, F. P. S. Luus, N. Khan, N. Makondo, I. Y. Akhalwaya, H. Qian, R. Fagin, F. Barahona, U. Sharma, S. Ikbal, H. Karanam, S. Neelam, A. Likhyani, S. K. Srivastava (2020): Logical Neural Networks. CoRR abs/2006.13155
- P. Sen, B. W. S. R. de Carvalho, R. Riegel, A. G. Gray (2022). Neuro-Symbolic Inductive Logic Programming with Logical Neural Networks, Proc. of AAAI
- L. Serafini, A. S. d'Avila Garcez (2016). Learning and reasoning with logic tensor networks, Proc. of the Italian Association for Artificial Intelligence