BORO – General Context

C-FORS Summer School in Foundational Ontology
(C-FORS 2025)

23 May 2025, University of Oslo, Norway

Chris Partridge, Chief Ontologist, BORO Solutions

Business Object Reference Ontology

Schedule

Morning Sessions		
	9:00 - 9:05	Session 0 – Introduction
	9:05 - 9:45	**Session 1 – Context**
	10:00 - 10:45	Session 2 – BORO Ontology
	11:00 - 12:00	Session 3 – Analysis Tools
Afternoon Sessions		
	1:15 - 3:30	Session 1 – Practical Examples
	3:30 - 5:00	Session 2 – Examples Discussion / Presentation

Session 1 – Context: Structure

- Overall approach
- Context
  - ontologies
  - foundational ontologies – purpose
  - foundational ontologies – nature

Overall approach

Overall practice-focused approach

- A common way of looking at engineering,
  - including ontological engineering,
    - is as a **practice**, a way of doing things
    - engineering as a discipline emerges from and supports the practice
- Foundational ontology engineering, or
  - ontological engineering using foundational ontologies,
    - is then also a **practice**
- Aim is to give some idea of the **practice** of using the BORO Foundational Ontology

Business Object Reference Ontology

First step

- Start by providing a general context to situate the BORO practice
  - frame this with three questions
    1. why bother with ontologies?
    2. why bother with foundational ontologies?
    3. what kind of thing are foundational ontologies?
  - the way we answer these questions provides a frame for understanding BORO’s approach to foundational ontology

Ontologies

Why bother with ontologies?

Business Object Reference Ontology

Why bother with ontologies?

- A common (almost universal?) answer
  - (which we subscribe to)
    - is
      - **they are useful/needed for (the practice of) improving interoperability**
        - **particularly semantic interoperability**

        - “However, knowledge-based systems pose special requirements for **interoperability**.
          …
          For such knowledge-level communication, we need conventions at three levels: representation language format, agent communication protocol, and specification of the content of shared knowledge.
          …
          Ontologies can be used for conventions of the third kind: content-specific specifications”
          - » Gruber, T. R. (1993). A translation approach to portable ontology specifications. Knowledge Acquisition, 5(2), Article 2. https://doi.org/10.1006/knac.1993.1008

Business Object Reference Ontology

What is interoperability?

- Term only started appearing in the 1960s
  - recognised as a capability
- Genealogically: the definitions reveal a shift in emphasis
  - in the early days, interoperability,
    - particularly semantic interoperability,
      - was theoretically defined in terms of a capability to understand
        - talk about ‘meanings’
  - however, as people studied the problem,
    - this has evolved into seeing it as an operational capability
      - talk about being able to do things
- Operational capability definition is more common nowadays among practitioners

- Early (theory):
  - interoperability among components of large-scale, distributed systems is the ability to exchange services and data with one another. … Semantic interoperability ensures that these exchanges make sense—that the requester and the provider have a common understanding of the “meanings” of the requested services and data.")
    - Heiler, S. (1995). Semantic interoperability. ACM Computing Surveys (CSUR), 27(2), 271–273.
- Evolved (pragmatic – operational):
  - within the past decade, it has been commonly defined as “The ability of systems, units, or forces to provide services to, and accept services from other systems, units or forces and to use the services so exchanged to enable them to operate effectively together.”
    - Ford, T. C., Colombi, J. M., Graham, S. R., & Jacques, D. R. (2007). A survey on interoperability measurement. Gateways, 2(3)

If interoperability is recent, why now?

- The genealogy of (the term) interoperability shows,
  - as noted, it is relatively recent
    - term only started appearing in the 1960s
    - emerged around the same time as use of computing in enterprises/institutions
      - closely linked to computers/machines working together
- This implies a connection with computing
  - not the same kind of concern prior to computing

What kinds of computing (machine) interoperability?

- If the new concern is
  - how does communication work in the age of computing machines?
- This naturally gives rise to a three-way distinction
  - based on the types of systems doing the exchange
    1. Machine-Machine
    2. Machine-Human
    3. Human-Human
  - of these three – the first (and second) seem novel
    - the third is plainly not novel!
  - we will argue that as automation increases
    - more and more information is communicated from 'Machine-Machine'

Visualising the three-way distinction

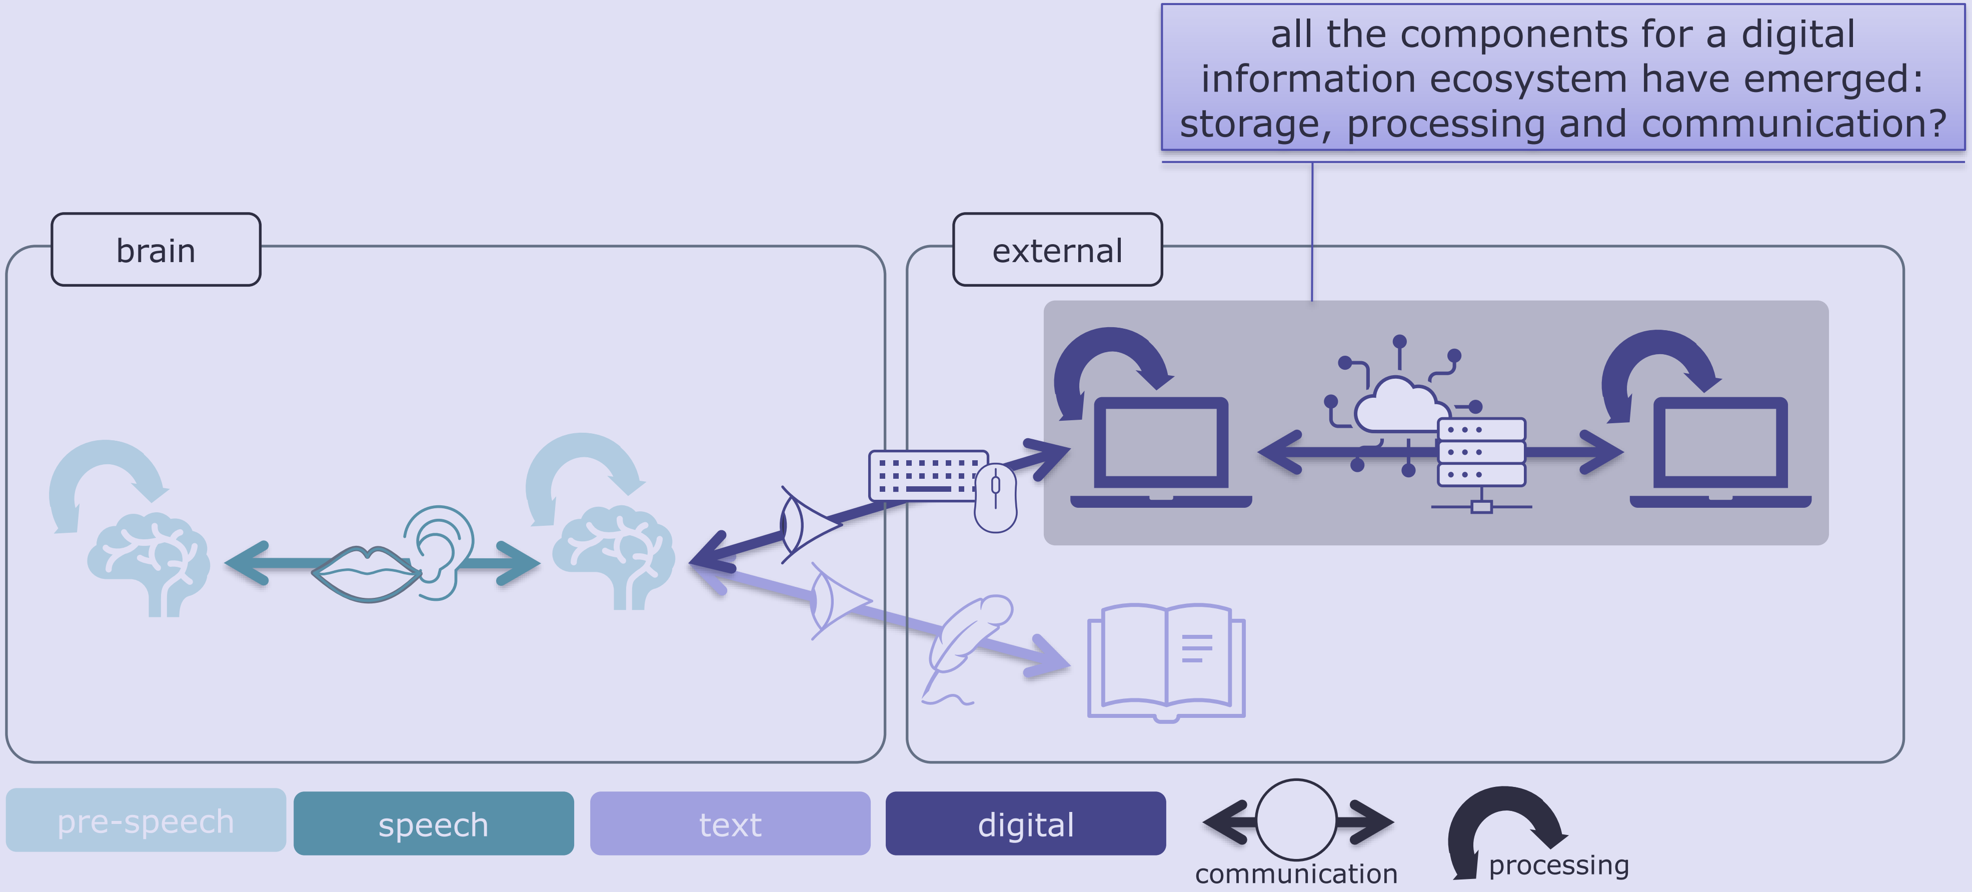

Ancient Human-Human lack of interoperability

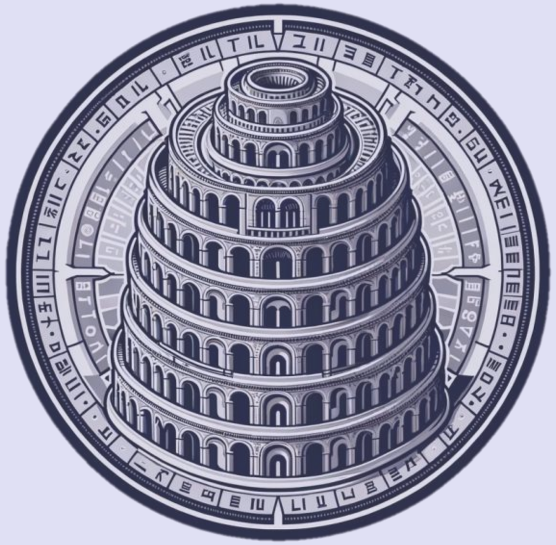

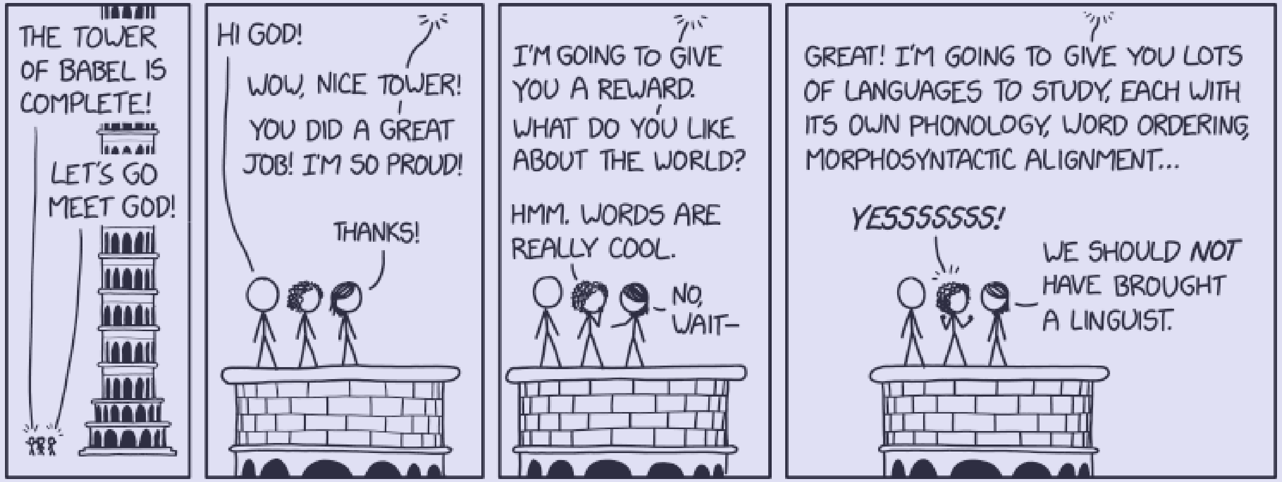

Tower of Babel narrative in Genesis 11:1–9 speaks to the power of interoperability and to the curse of not having it.

> *And the LORD said, "Look, they are one people, and they have all one language, and this is only the beginning of what they will do;* ***nothing that they propose to do will now be impossible for them****. Come, let us go down and confuse their language so they will not understand each other."*

In this 'myth', implausibly, the journey is 'backwards' from being interoperable to not being interoperable

Opportunities for improving (machine-machine) interoperability capability

- Different ways to 'lack' (machine-machine) interoperability capability
  - failing (existing) capability
    - common example
      - failures in 'accuracy' of existing interoperability capability
        - Mars Reconnaissance Orbiter – imperial versus metric
  - (new) capability gaps
    - not in the current (existing) capability
      - common example
        - curation requirements in data lakes
          - » raw data does not interoperate – it needs prior curation
          - » curation is expensive, so limited

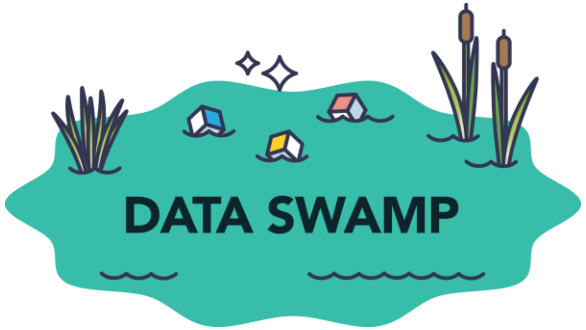

- From an opportunity exploitation perspective
  - capability gaps offer more opportunity for machine-machine exploitation
    - obvious area for improvement is reducing cost of 'curation' (i.e. interoperability)

A Coasian ‘metric’ for the cost of interoperability

a simple synchronous interoperability test

We can use intra-operability as a benchmark to assess interoperability.
The goal is for the costs of intra- and inter-operability to be roughly comparable.
It should be (roughly) as easy to communicate within as across systems.

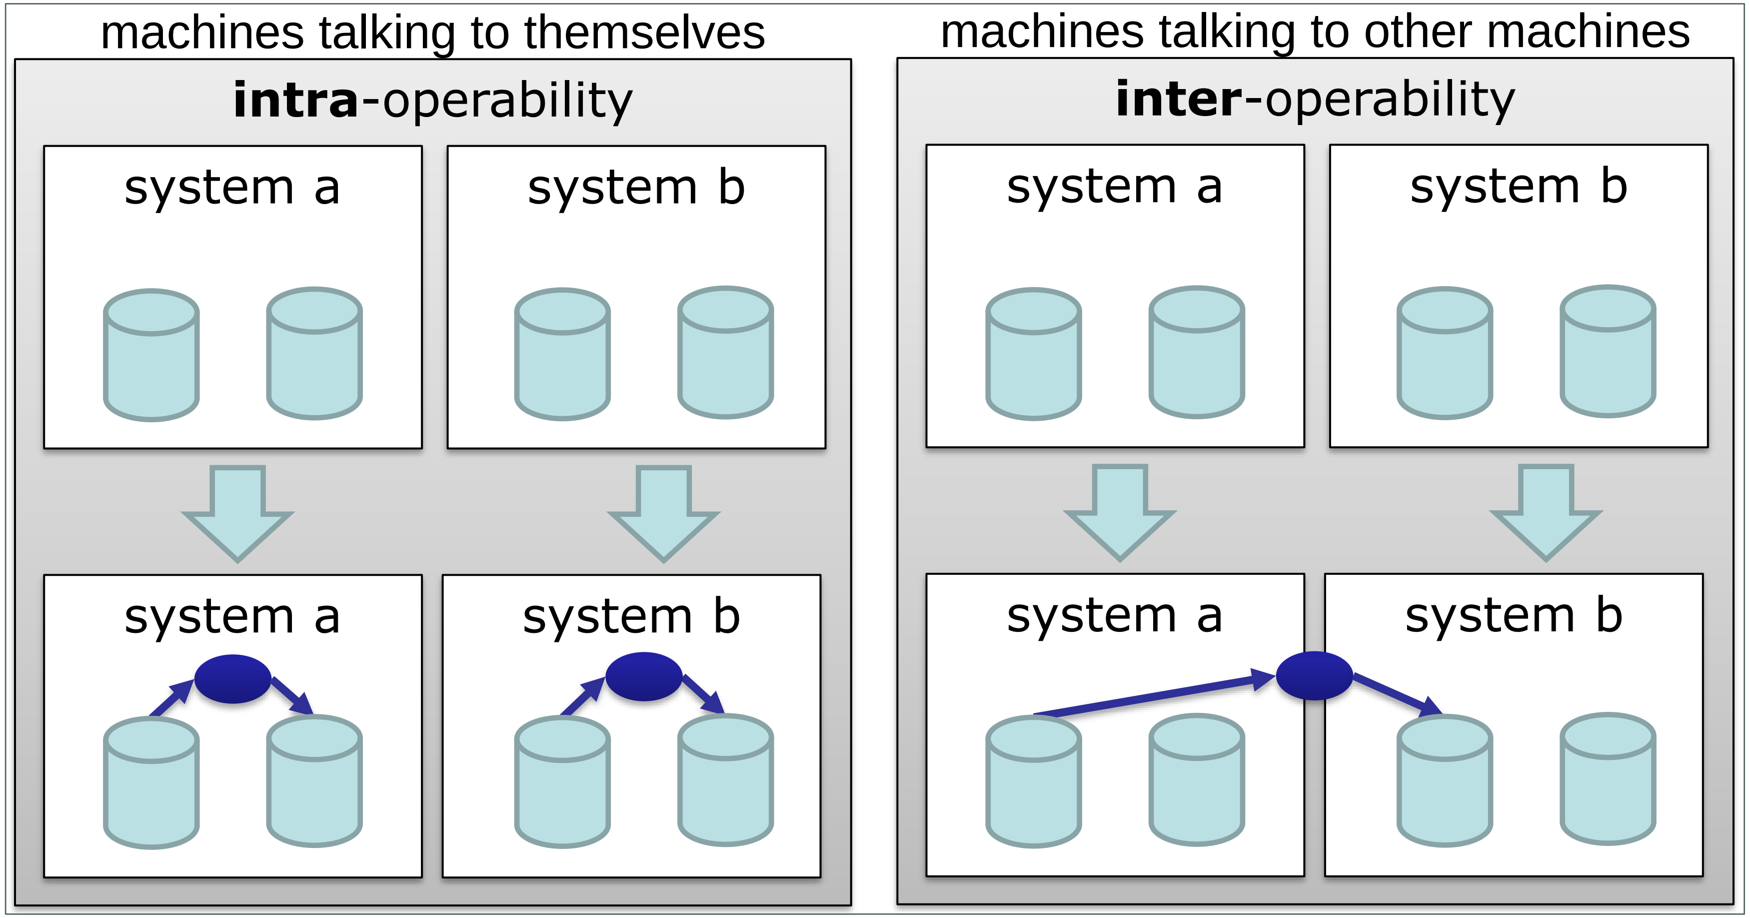

Coase, R. H. (1937). The nature of the firm. Economica

The rise of machine-machine interoperability

Metaphor from the 1980s: islands of automation in construction

Islands are data (digitised)
Sea is pre-data (manual)

This version is from1996

Height of the island shows
when data emerged:
nothing before the 60s

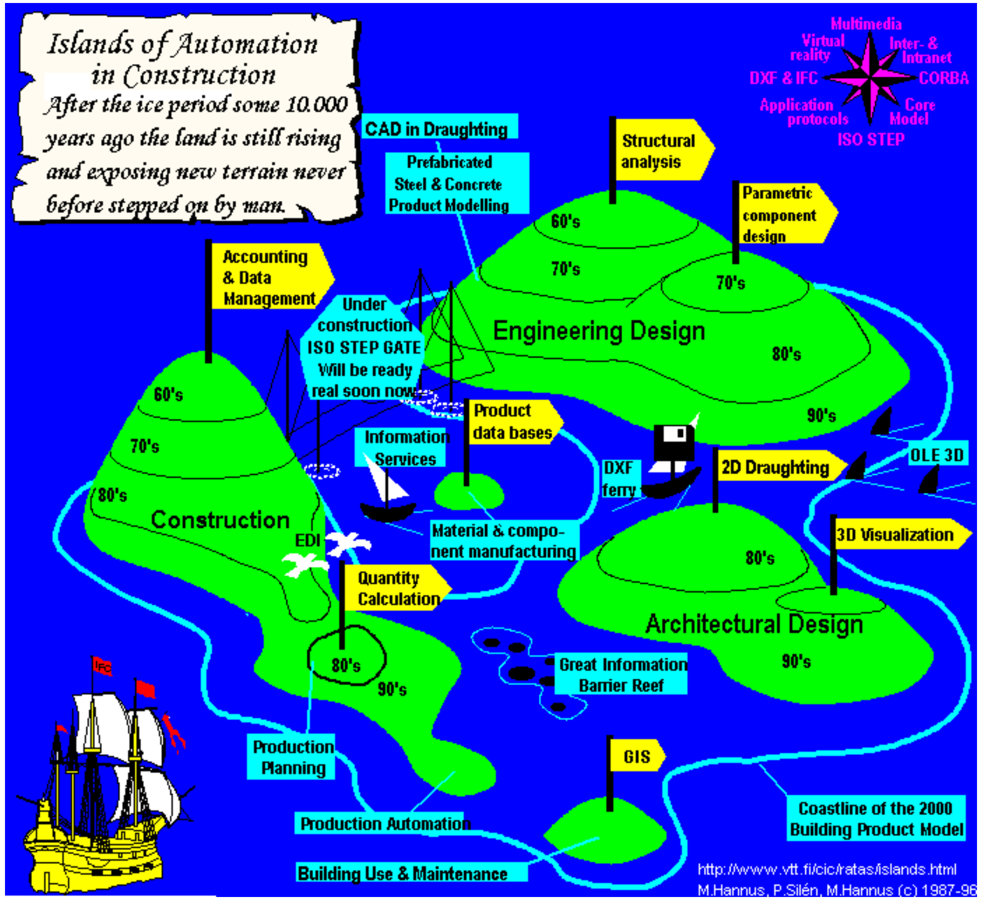

https://en.wikipedia.org/wiki/Islands_of_automation

https://en.wikipedia.org/wiki/Islands_of_automation
see also: Bjork, B. (1987) The integrated use of computers in construction - The Finnish experience, ARECCAD 87, Barcelona,

Expanding – over time - islands of automation

tl;dr – as systems grow, they abut and so need to communicate

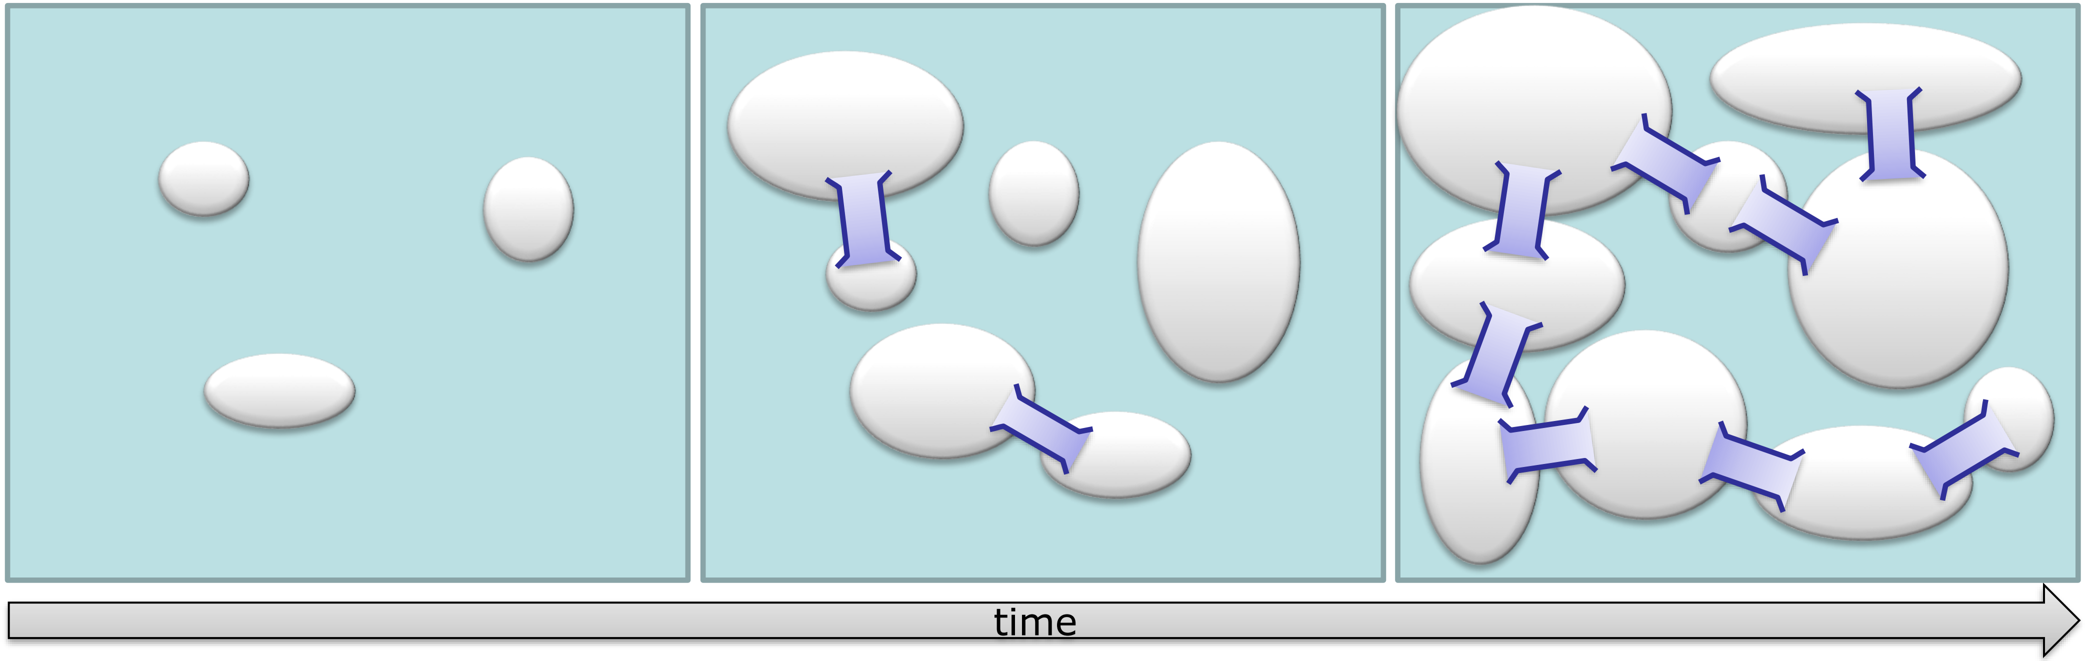

As the islands expand (**automation** increases), the systems abut – and so, the need to communicate (**interoperability - data exchange**) arises and increases.

The typical enterprise is now so automated that it no longer makes sense to talk about ‘expanding islands’ of automation … maybe, given the size, ‘**continents** of automation’. However, it does still make sense to look at the bridges between the islands/continents.

4D perspective reveals vertical and horizontal data exchange

System migrations, and their integral data migrations, are diachronous.

Where the time axis goes up the page, these can be thought of as vertical migrations.

By the same reasoning, system to system APIs are synchronous and so can be thought of as horizontal exchanges.

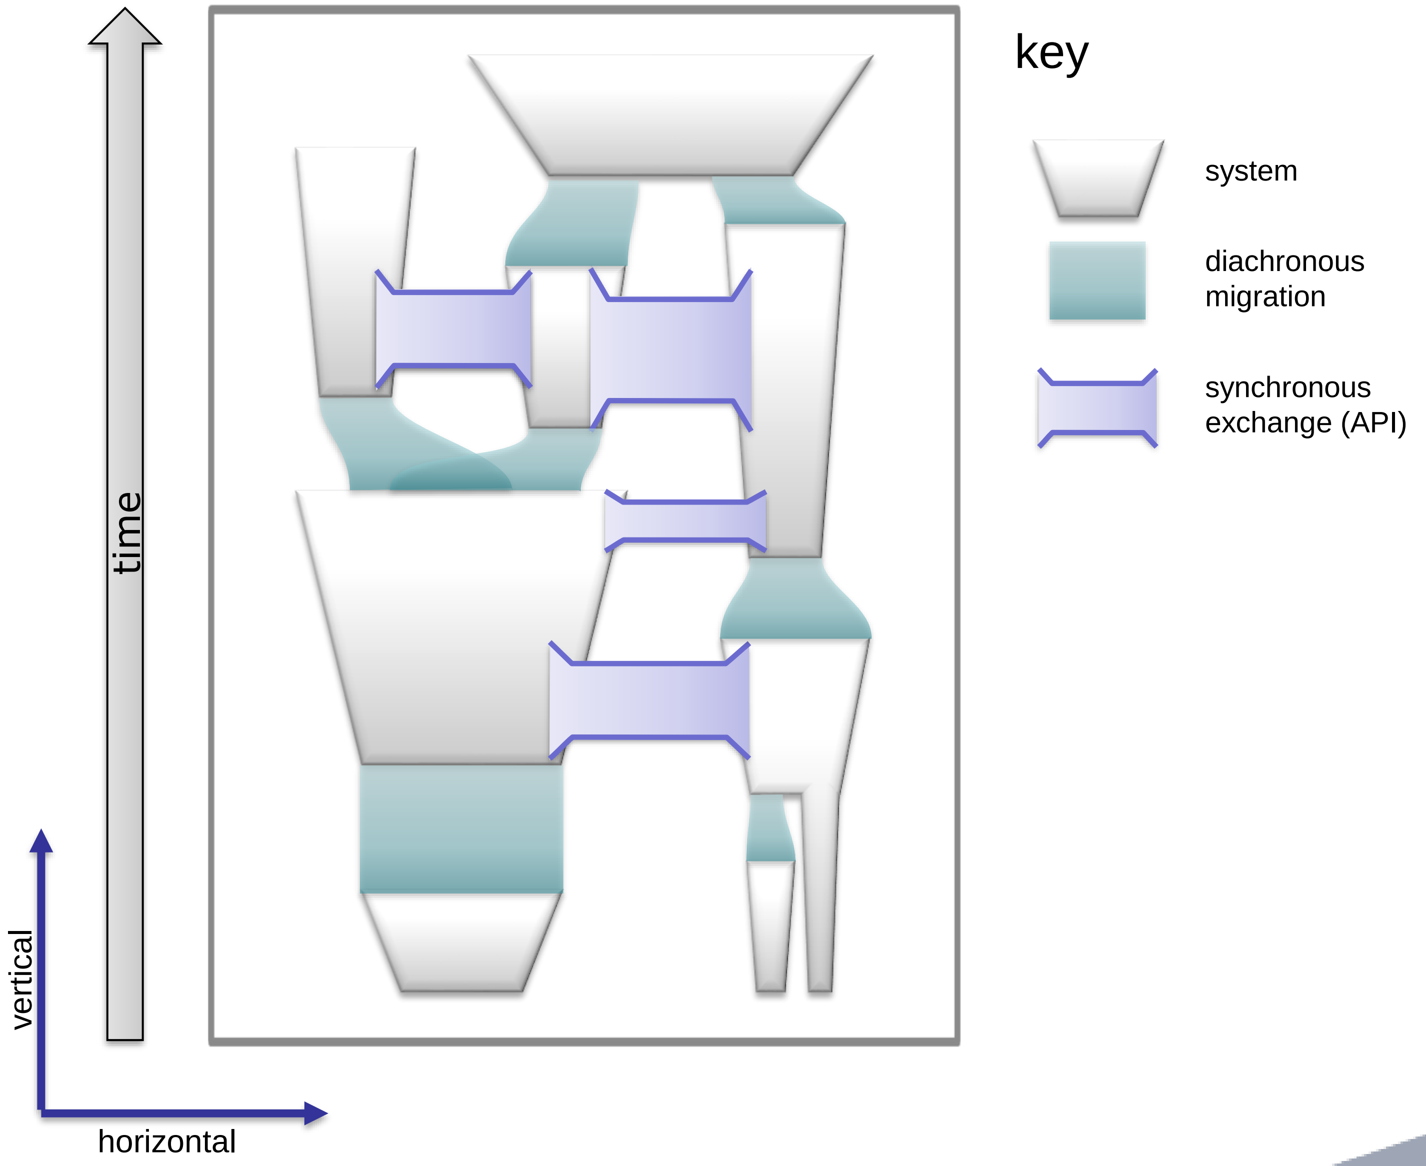

The significant scale of interoperability

- As systems get bigger, and
  - the types of data typically increase
    - the size of the task of ensuring interoperability increases
- Systems **are** now quite big
  - some idea of the scale
    - take as a reference a common enterprise application:
      - for example, SAP ERP (https://en.wikipedia.org/wiki/SAP_ERP)
        - a typical system may have 40,000 tables with over 1 million columns
      - anecdotal evidence that systems on average get replaced every decade or so
        - so, an enormous amount of migration going on

Business Object Reference Ontology

Narrow alleyways of machine interoperability

Returning to the ‘capability gaps’ way to ‘lack’ (machine-machine) interoperability capability

When one looks more closely at current interoperability, it becomes clear that the data scope is very narrow.

- In other words, the capability gaps are wide
- This becomes particularly apparent when we attempt to integrate data
  - for example, when multiple systems’ raw data is merged into data lakes
  - the data usually needs curating (transforming) before it is interoperable
    - the cost of this is significant, prohibitive.

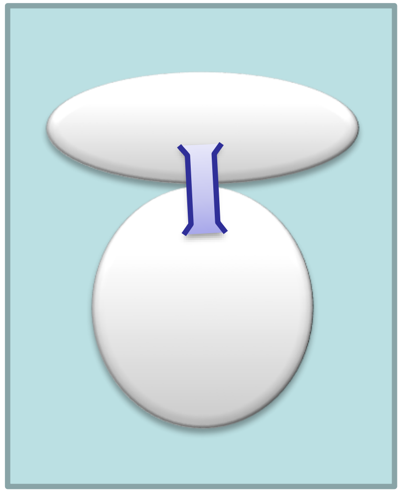
narrow

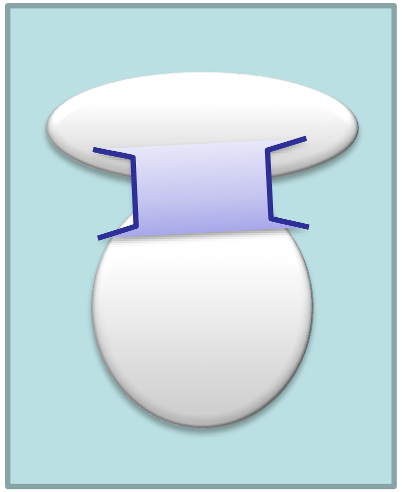
broad

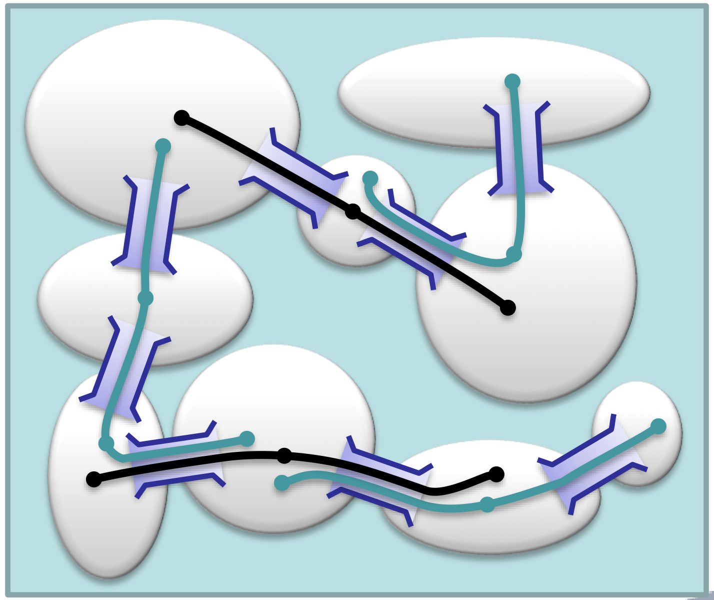

Lack of integrated interoperability

- The narrow data scope of the interoperability is a marker of missed opportunities – a ‘capability gap’
  - it is a symptom of our inability to fully exploit the information in our systems
    - it is ubiquitous and endemic
      - every application system seems to exhibit this feature
    - it grows as our systems grow
      - each new system creates another barrier to integrated interoperability
  - better exploitation requires better tools
    - ones that significantly reduce costs

Process versus the end product

Process versus the end product

- There appears to be a natural general tendency in engineering to focus on the end product and not take much account of its life cycle.
- In the context of top-level ontologies, this results in a focus upon the ontology (even the top ontology)
  - and significantly less attention, sometimes no attention,
    - on the ontologisation process from which this emerges.
  - this is clearly reflected in the current status of work on computational top-level ontologies.

- Chris Partridge, C., Mitchell, A., de Cesare, S., Beverley, J. (2025). "Broadening Ontologization Design: Embracing Data Pipeline Strategies". https://www.academia.edu/129382330

Process versus the end product in logic

- As (Dutilh Novaes, 2015) notes, in the context of logic, the process of formalisation is often neglected, and attention is focused on the formal results.
- She notes the importance of the former for real life application of the latter.
- She discusses two historical examples of processes of logical formalisation:
  - Aristotle’s syllogistic theory from the “Prior Analytics”, and
  - medieval theories of supposition,
  - to both illustrate and illuminate how to formalize logical arguments.
- Dutilh Novaes, C. (2015). "The formal and the formalized: The cases of syllogistic and supposition theory". Kriterion: Revista de Filosofia, 56, 253–270.
- See also:
  - Dutilh Novaes, Catarina. 2012. Formal Languages in Logic: A Philosophical and Cognitive Analysis

Business Object Reference Ontology

The current formalisation process

- Formalisation is:
  - an iterative, (discovery,) empirical exercise
    - should aim to 'let the data speak'
    - in evolutionary terms,
      - exploring the paths through the fitness landscape

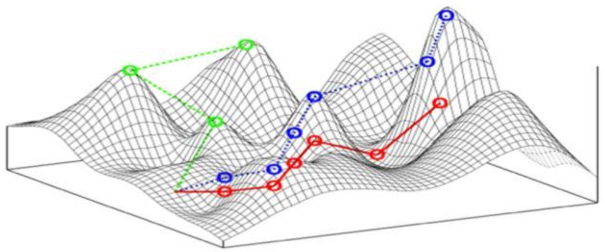

- Current situation
  - the formalisation process
    - is largely unexamined (AKA unmanaged)
    - typically, uncritically gives control of the form to the syntax

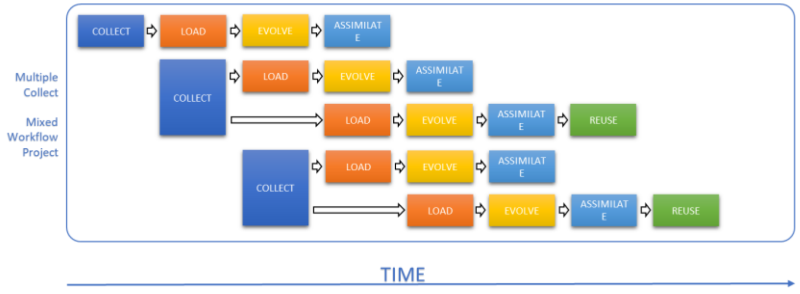

Exploitation opportunity

- This suggests an exploitation opportunity
  - improving the end product
    - by better engineering the process
  - the improved engineering is a way to reduce the capability gap
    - increasing interoperability with a better engineered process

A closer look at a ‘capability gap’: in the ‘data exchange process’

Look more closely at a ‘data exchange process’

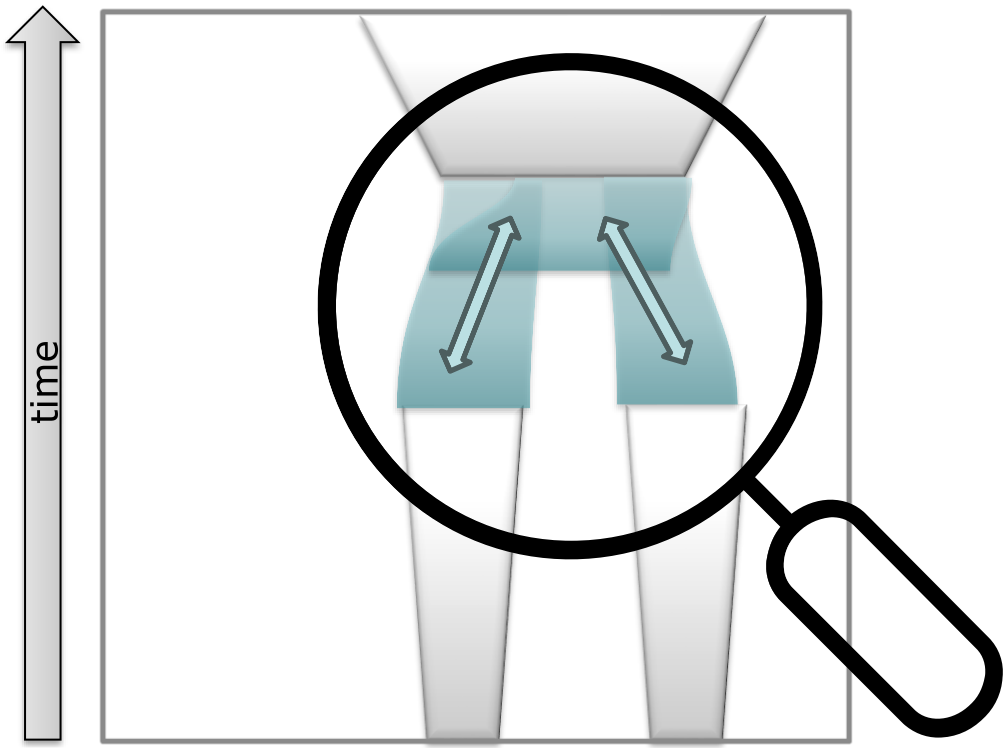

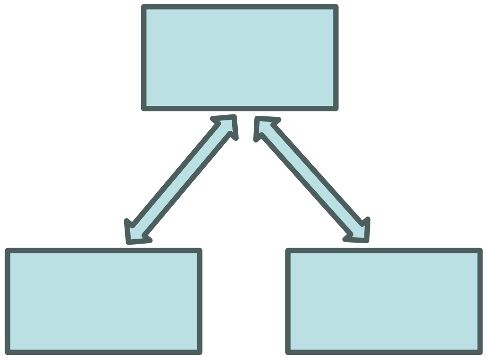

Schematically, there are end nodes.
The data exchange migrates and transforms data between these nodes.
The transformation changes the data from the format of one system into the format of another

Data exchange transformation

- The data exchange involves the transformation of data in the format of one system into the format of another
- Review two aspects of the transformation
  - transformation as translation and
  - translation reflecting reference

Translation - interlingua

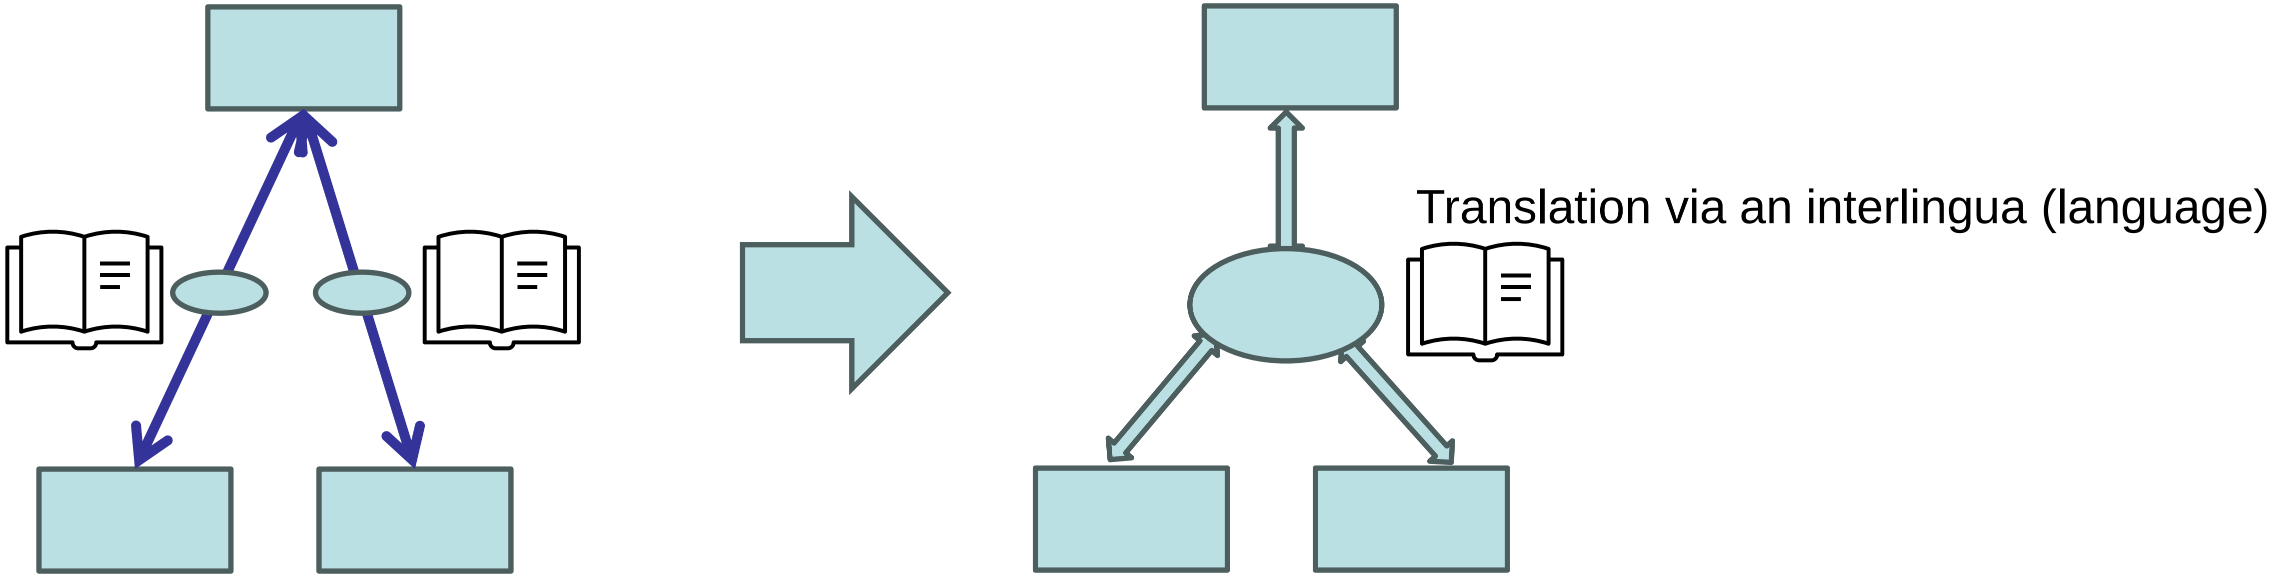

Use a common interlingua as the hub
for each translation spoke.

Reference - triangulate

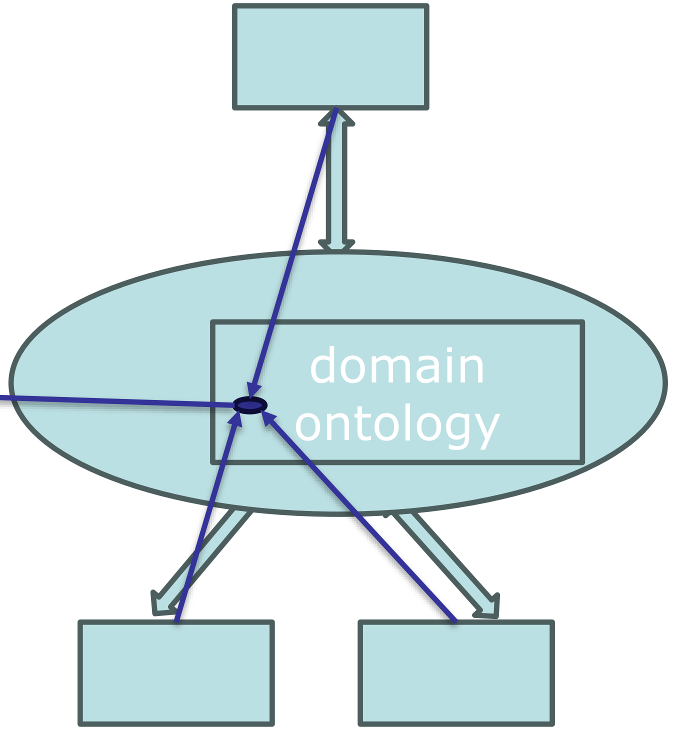

Triangulate the mapping using 'reference' to objects in the domain – that is, a domain ontology

This assumes the systems represent (refer to) the same things in the world.

Why bother with ontologies?

- Because improving interoperability is key to improving computer ecosystems, and
  - ontologies enhance interoperability
    - by helping to enable the use of reference
      - to triangulate the interlingua translation
        - in the data exchange process

Foundational ontologies – purpose

Why bother with foundational ontologies?

Why bother with foundational ontologies?

- We have a domain ontology
- Why do we need a foundational ontology?
- Show that:
  - the domain ontology, on its own, is undermined by ontological relativity’s
    - three theses of indeterminacy
  - this indeterminacy is resolved by the foundational ontology

Business Object Reference Ontology

Quine’s three theses of indeterminacy

- Ontological relativity is underpinned by three theses of indeterminacy
- “Three theses of indeterminacy have figured conspicuously in my writings:
  - indeterminacy of **translation**,
  - inscrutability of **reference**, and
  - underdetermination of scientific theory.”
    - *Quine, W. V. O. (2008). Three Indeterminacies.*
  - *NOTE: translation and reference are key to successful interoperability*
- Regarded as an important topic: for example:
  - “W. V. O. Quine’s contention that translation is indeterminate has been among the most widely discussed and controversial theses in modern analytical philosophy. It is a standard bearer for one of the late twentieth century’s most characteristic philosophical preoccupations …”
    - *Crispin Wright (1999). "Chapter 16: The indeterminacy of translation*
- Leave aside the controversy
  - we can use this to explain the role of foundational ontology
    - as a tool to manage metaphysical indeterminacy

Business Object Reference Ontology

Three consequences of indeterminacy

- There may be **no** unique way to:
  - translate between two languages
    - so, no single candidate interlingua
      - weakens the translation interlingua approach
  - determine the reference of terms in the languages
    - so, no unique set of real-world objects to refer to
      - weakens the real-world approach
  - build a scientific account of our empirical data
    - so, we cannot empirically test our theories

- Roughly:
  - the metaphysics deals with the ‘stuff’ science cannot

For more background on Quine’s ontological relativity

- WVO Quine’s essay, “Ontological Relativity” (1969) turned 50 in 2019
- Recurring topic in his work, see (for example):
  - 1960, Word and Object
  - 1970, On the Reasons for the Indeterminacy of Translation
  - 1974, The Roots of Reference
  - 1975, On Empirically Equivalent Systems of the World
  - 2008, Three indeterminacies
  - …
- See also: https://plato.stanford.edu/entries/quine/#UndeTheoEvidIndeTran

Why bother with foundational ontologies?

- The domain ontology, on its own, is undermined by ontological relativity's
  - three theses of indeterminacy
- There are metaphysical choices that aren't empirical

- Foundational ontologies are a way to regain determinacy
  - the foundational ontologies address these choices in a systematic metaphysical architecture
  - this mitigates the indeterminacy
  - but, it can be resolved in multiple ways,
    - by foundational ontologies with different metaphysical architectures
      - in other words, there are multiple possible foundational ontologies

Visualising the revised architecture

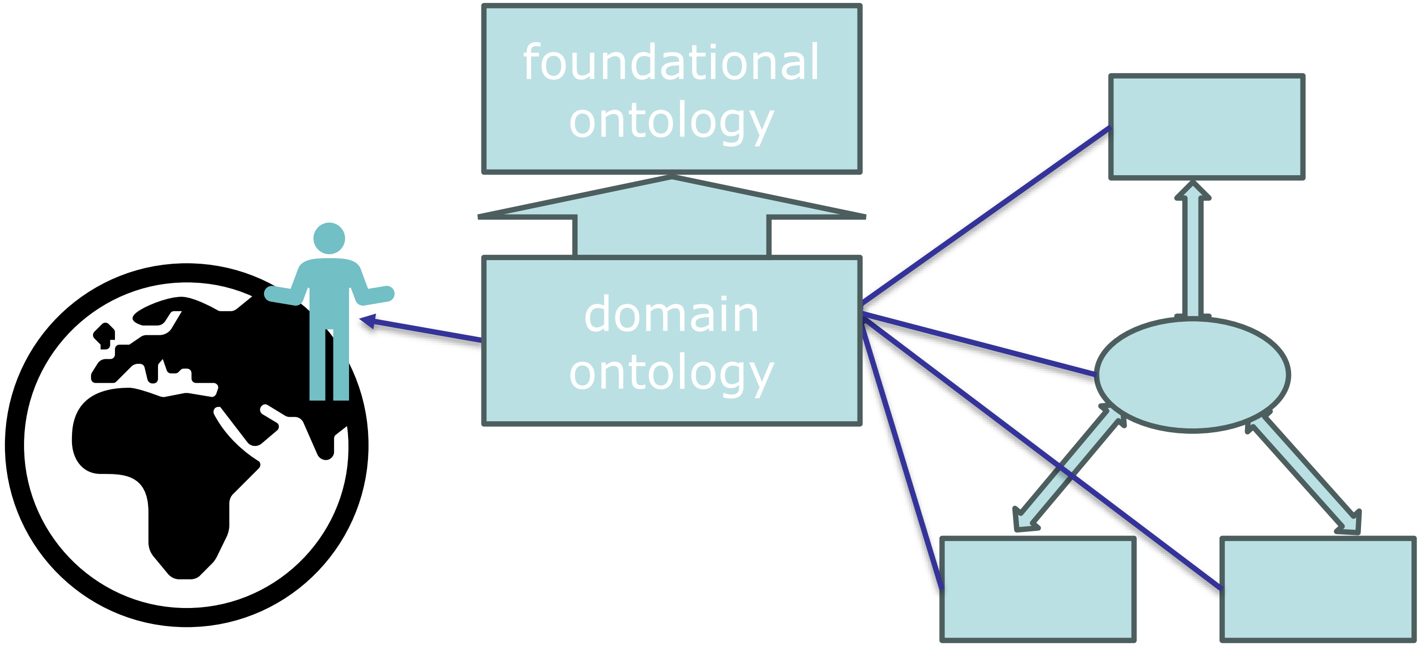

Whither truth?

- The existence of multiple foundational ontologies,
  - where the choice between them cannot be determined by empirical tests
    - is evidence of the indeterminacy
- But what happens to ‘truth’ in this situation?
  - “This is an issue on which Quine has not merely changed his mind but vacillated, going back and forth between what he calls the “**sectarian**” and the “**ecumenical**” responses.
    - the sectarian response is to say that we should not let the existence of the alternative in any way affect our attitude towards our own theory: we should continue to take it seriously, as uniquely telling us the truth about the world. (We are assuming that the two theories possess all theoretical virtues to equal degree; clearly Quine would say that if one theory were superior in some way then we would have reason to adopt it.)
    - the ecumenical response, by contrast, counts both theories as true.”
      - *https://plato.stanford.edu/entries/quine/#Unde*

Foundational ontologies – nature

How do foundational ontologies resolve this indeterminacy?

Explaining the nature of foundational ontologies

- One way of characterising the indeterminacy is through metaphysical choices
- A foundational ontology makes a range of metaphysical choices in a coordinated way
  - this results in a metaphysical architecture
  - the range of the metaphysical choice is what generates indeterminacy
  - making the metaphysical choice resolves the indeterminacy
- So, one way of explaining the nature of foundational ontologies is through the metaphysical choices that underpin their metaphysical architecture

Example category of architectural choice: whether to stratify or unify

One of the categories of architectural choice: horizontal aspects

“4.2.2 Horizontal aspects: stratification versus unification
There is a group of fundamental choices that impact the ontological architecture which involves whether or not to make a distinction. If one chooses not to make the distinction, one only introduces a single type. If one chooses to make the distinction, one introduces two types; one for each alternative. The choice boils down to whether to horizontally stratify or unify. One can describe choosing to make the distinction as ‘separating one potentially unified type into two’, creating a horizontal stratification in the hierarchy – and not making the distinction, ‘**unifying the potentially separated two types into one**’.”

web-based: https://digitaltwinhub.co.uk/a-survey-of-top-level-ontologies/#a_survey_of_TLOs_contents

Example: Stratification as a journey

Table 8 – Summary of the horizontal stratification choices introduced

Label	Unified type	Separate types	Stratifying relation
spacetime	spatio-temporal objects	spatial objects, temporal objects	spaces are multiply located at times (though this is often a derived relation – from an occupying object's links to both space and time)
locations	supersubstantival objects	(physical) objects, locations	objects are (exactly) located at their locations
properties	objects	substances, properties	substances are bearers of properties
endurants	perdurants	continuants, occurrents	occurrent is dependent upon continuant
immaterial	(physical) objects	material objects, immaterial objects	immaterial objects are part of material objects

4.2.2.7 Stratification journey

There is limited inter-dependence between the choices meaning that a range of permutations are possible. For a top-level ontology, one can visualise the architectural stratification choices being adopted in a sequence, starting with no stratifications and introducing the choices one or two at a time – as illustrated in the figures below. This sequence or journey is a rational reconstruction – the original development of the top-level ontology is most likely ad hoc and bottom up. However, this reconstruction gives us a good picture of the underlying architecture.

Example metaphysical choice: locations

Label	Unified type	Separate types	Stratifying relation
locations	supersubstantival objects	(physical) objects, locations	objects are (exactly) located at their locations

web-based: https://digitaltwinhub.co.uk/a-survey-of-top-level-ontologies/#a_survey_of_TLOs_contents

4.2.2.2 Locations

People often talk of physical objects and their locations, where physical objects occupy their locations, suggesting two related types; objects and locations (let's leave the decision whether the location is spatial, temporal or spatiotemporal to the previous choice). For example, "today your car is parked in the same place as mine was yesterday" could be regarded as a location which was occupied by my car (a physical object) yesterday and your car today. There is a debate going back to Newton and Leibnitz in the 17th century as to whether location is absolute or relative. If it is relative, then location is clearly fundamentally different from physical objects – which aren't. However, if it is absolute, a kind of substance, then this opens the possibility that one could unify objects and their locations as fundamentally the same, technically known as supersubstantivalism. If one does not have cases of interpenetration (see 4.3.3) then this resolves the oddity where physical objects exactly occupy a single location throughout their life – unifying eliminates this double counting. If there is interpenetration, then two objects may collapse to the same location – which may have unintended consequences. After the unification, the physical object and its location, two kinds of substance, are replaced by a single supersubstantival object. One has a broad choice between separating or unifying physical objects and locations.

Two (of many) levels of unifying-stratifying

matter and the space-time it is located in are unified as supersubstantival objects – matter is then a way space-time can be

locations in space and locations in time are unified as locations in space-time

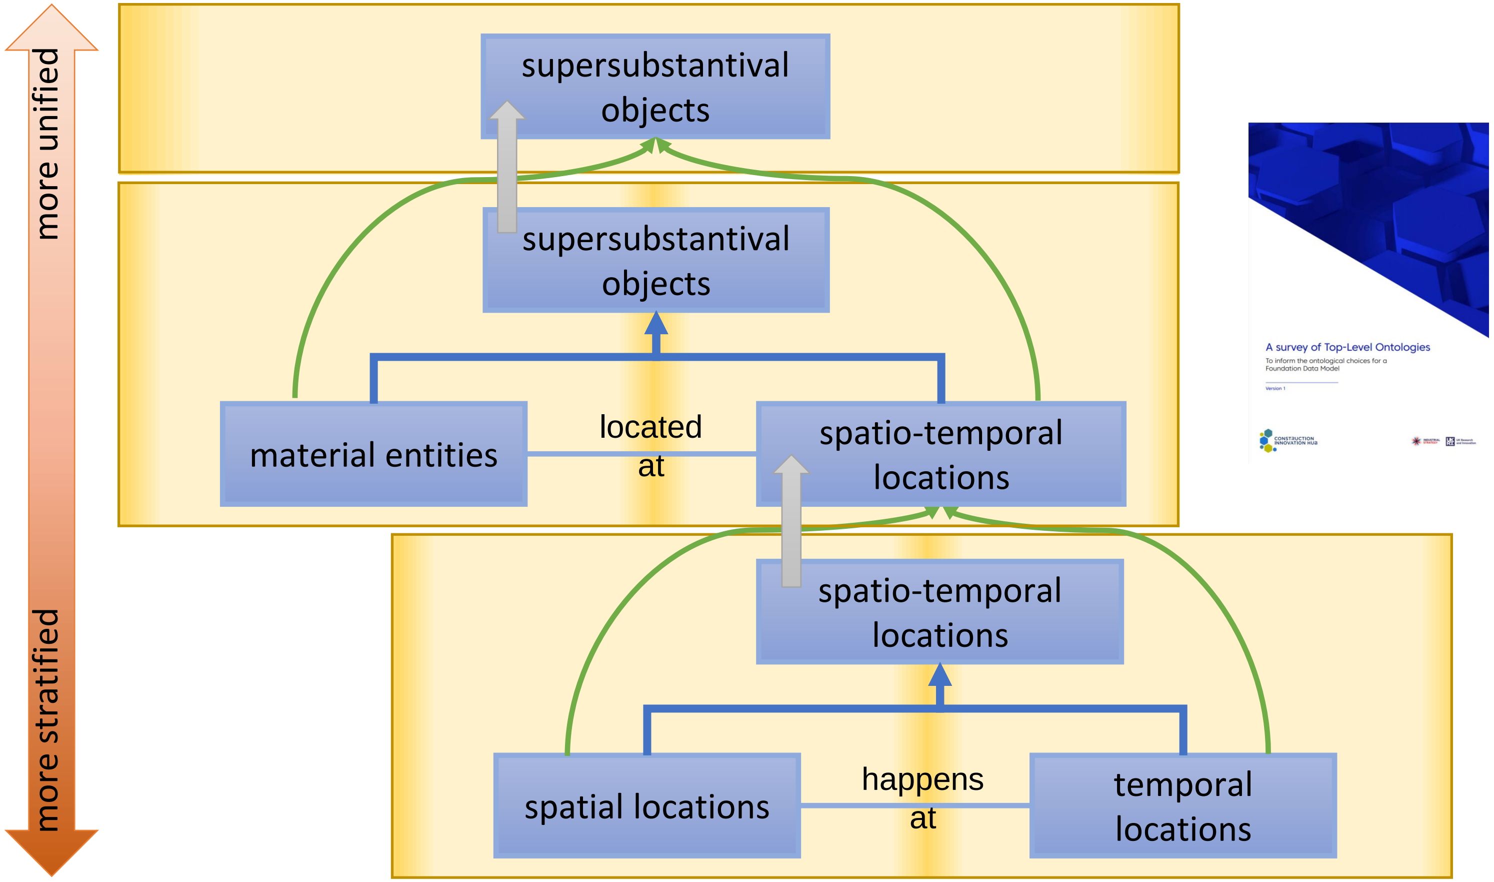

web-based: https://digitaltwinhub.co.uk/a-survey-of-top-level-ontologies/#a_survey_of_TLOs_contents

Locations example choice

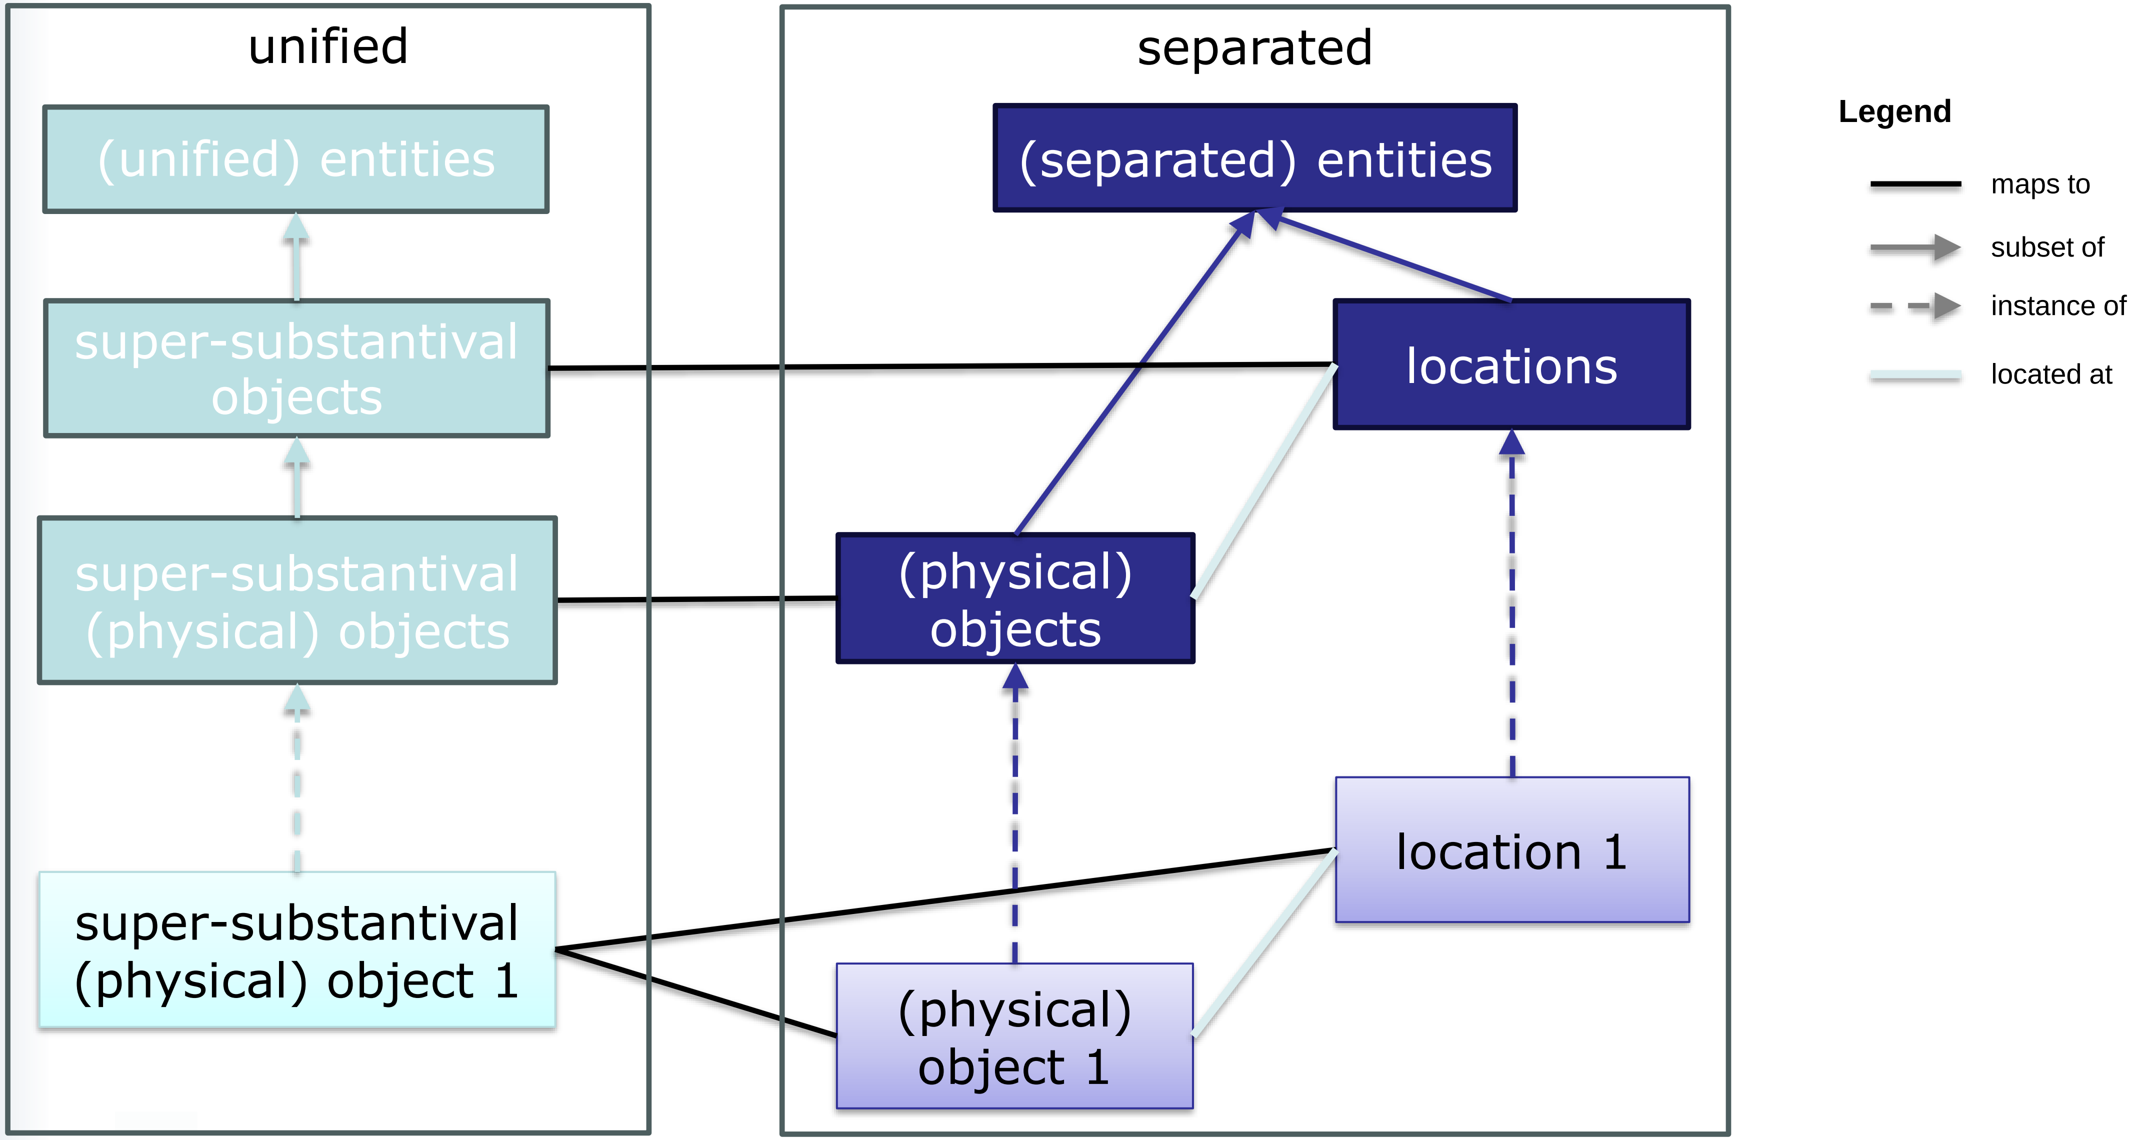

The HMELR test

- We developed the HMELR test
  - to expose the implicit ‘horizontal’ metaphysical choices
- Acronym:
  - **H**ow-**M**any-things-**E**xactly-**L**ocated-in-this-**R**egion?
- Basing the test on a region is empirically conservative
  - it does not introduce empirical differences
- The test trades on
  - a difference
    - stratification: leads to a multiplication of exactly coinciding objects of the stratified types
      - and a linking relation between the types
    - unification: leads to a collapse into a single object of the unified type
      - with identity as the linking relation (if you need one)
  - this difference leads to a difference in numerical identity
    - the count reveals the underlying metaphysical choice
- The ‘numerical identity’ difference
  - also illustrates **how data behaves under the choice**
    - multiplying objects leads to multiplying their representations (rows, etc.)

Locations example - HMELR test

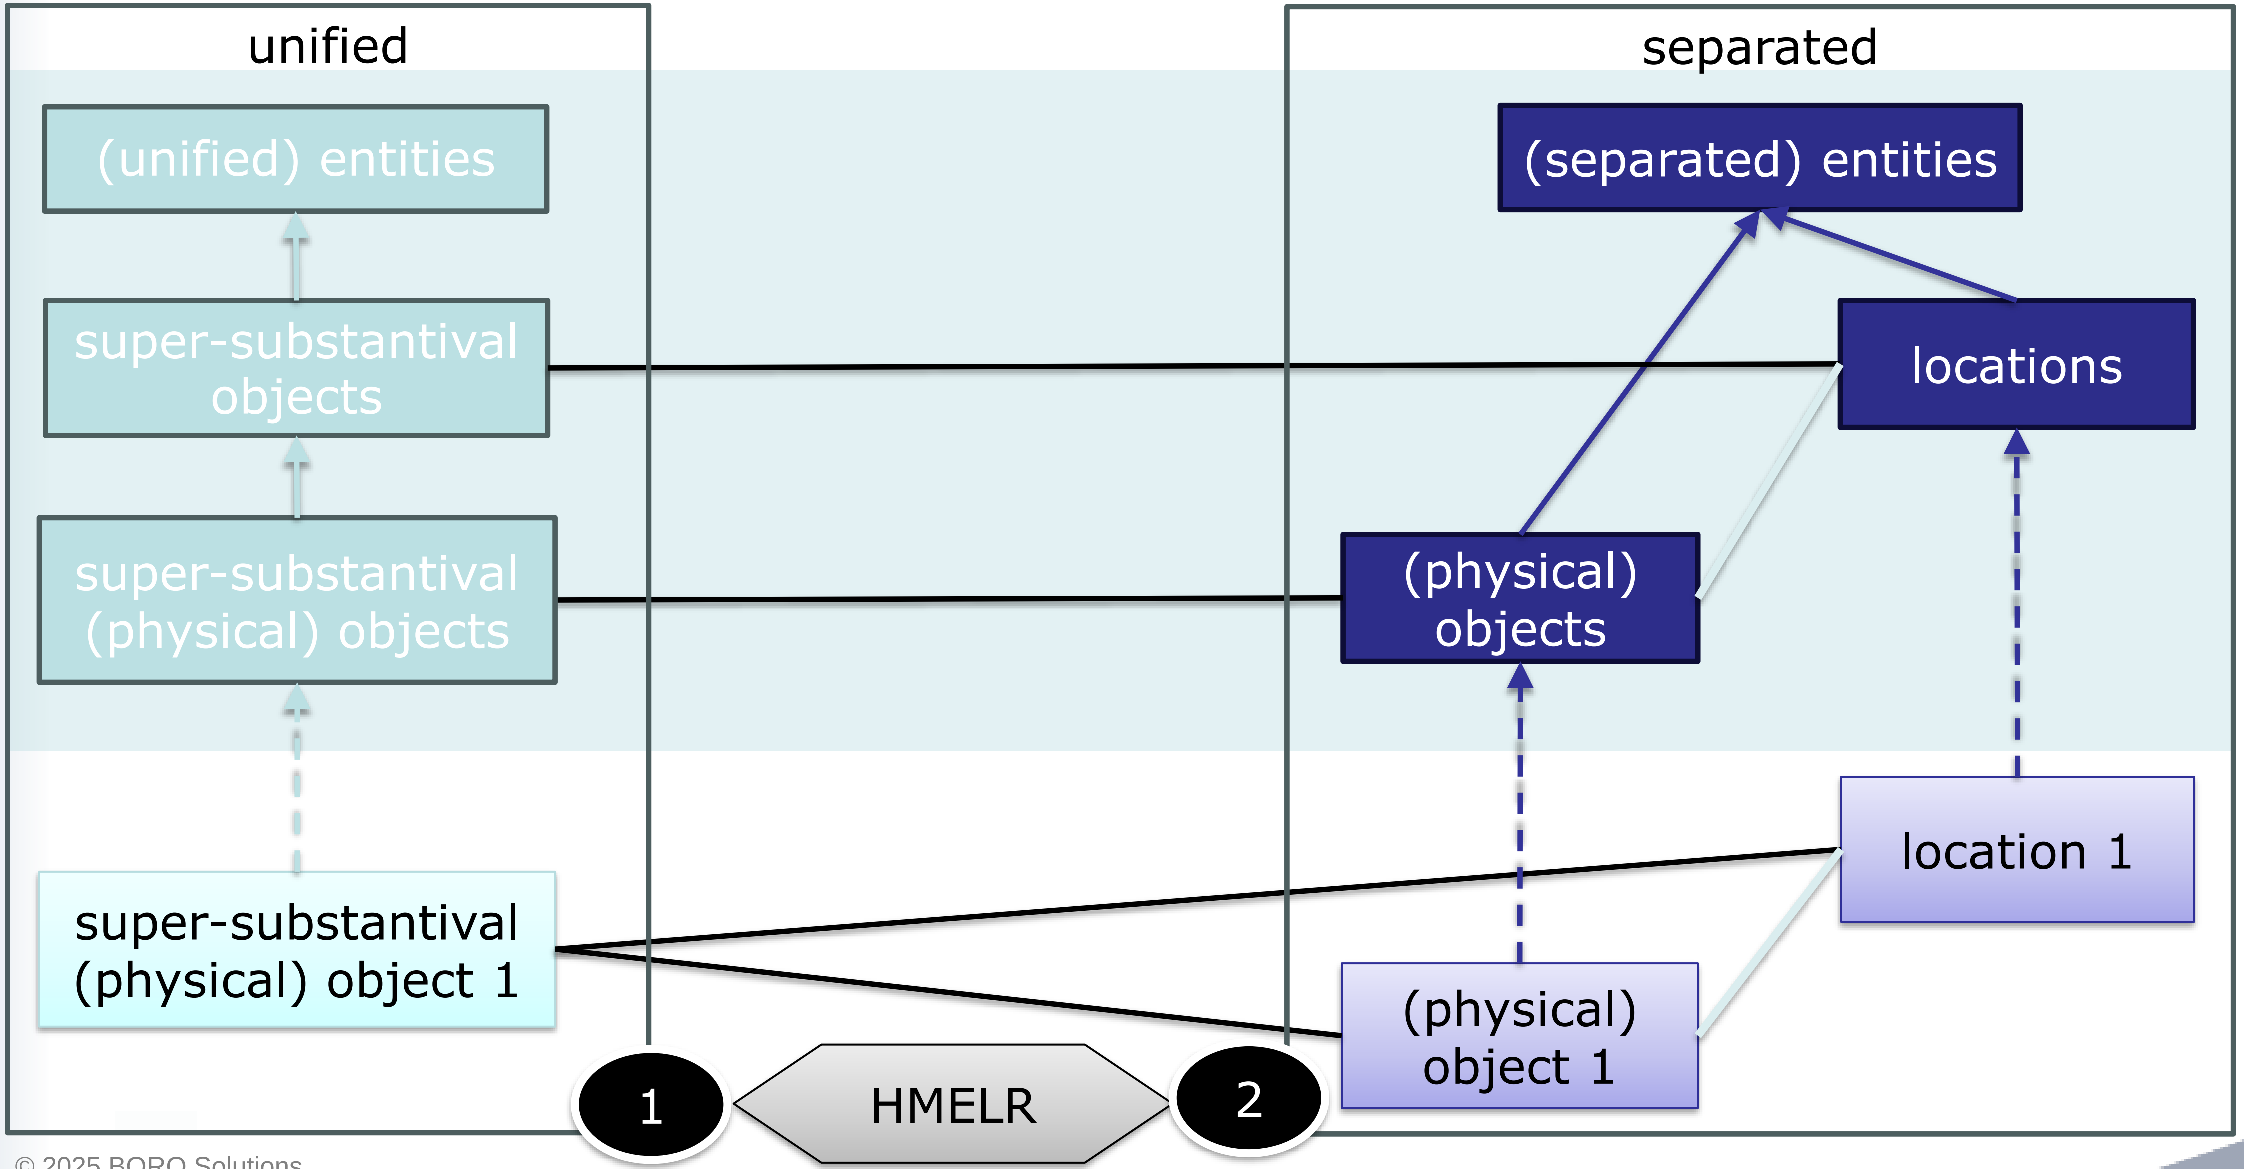

Corresponding database indeterminacy

There is a corresponding ‘numerical identity’ difference for data behaves under the choice.
Multiplying objects leads to multiplying their representations (rows, etc.)

super-substantival database

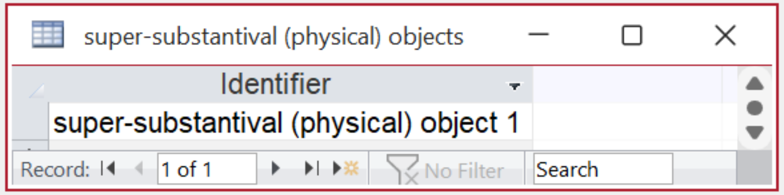

substantival database

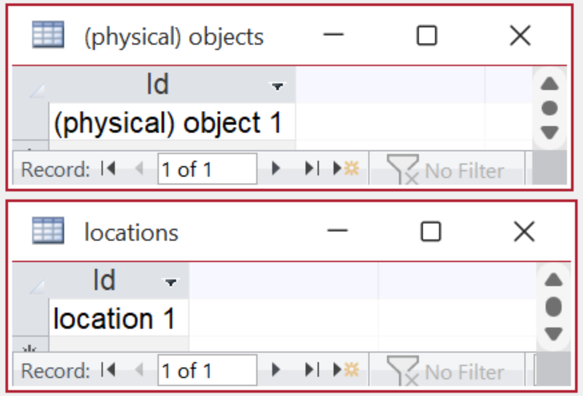

Business Object Reference Ontology

A framework for: mapping the metaphysical architecture landscape

A map to navigate the different architectural possibilities

31 ontological choices

37 top ontologies shortlisted and assessed

Appendix E: Summary of Framework Assessment Matrix Results

category	vertical aspect								
type	parent-arity		transitivity	boundedness			stratification	formal generation	
relation	type-instance	super-sub-type	super-sub-type	type-instance			type-instance	whole-part	
characteristic				downwards	fixed finite levels	number of fixed levels		fusion	compler
choice	single or unconstrained	single or unconstrained	yes or no	bounded or unbounded	fixed or not-fixed	[a number]	stratified or unstratified	yes or no	yes or n
BFO	unconstrained	single	yes	bounded	fixed	2	stratified	no	no
BORO	unconstrained	unconstrained	yes	bounded	not-fixed	not applicable	unstratified	yes	yes
YAMATO	not assessed	single	yes	bounded	fixed	2	stratified	yes	yes
HQDM	unconstrained	unconstrained	yes	bounded	not-fixed	not applicable	unstratified	yes	yes
IDEAS	unconstrained	unconstrained	yes	bounded	not-fixed	not applicable	unstratified	yes	yes
ISO 15926-2	unconstrained	unconstrained	yes	bounded	not-fixed	not applicable	unstratified	yes	yes
UFO	unconstrained	unconstrained	yes	bounded	not-fixed	not applicable	stratified	no	no
GFO	unconstrained	single	yes	bounded	not-fixed	not applicable	unstratified	yes	yes
KR Ontology	not yet assessed	unconstrained	yes	not yet	not	not assessed	not assessed	not	not asse
DOLCE	unconstrained	single	yes						
ConML+CHARM	unconstrained	single	yes						
CIDOC (ISO 21127:2014)	unconstrained ?	unconstrained	yes						

The ontological choices shape the architecture of the ontology

https://www.repository.cam.ac.uk/handle/1810/313452

web-based: https://digitaltwinhub.co.uk/a-survey-of-top-level-ontologies/#a_survey_of_TLOs_contents

Mapping the metaphysical architecture landscape

Business Object Reference Ontology

Horizontal stratification choices

Table 8 – Summary of the horizontal stratification choices introduced

Label	Unified type	Separate types	Stratifying relation
spacetime	spatio-temporal objects	spatial objects, temporal objects	spaces are multiply located at times (though this is often a derived relation – from an occupying object's links to both space and time)
locations	supersubstantival objects	(physical) objects, locations	objects are (exactly) located at their locations
properties	objects	substances, properties	substances are bearers of properties
endurants	perdurants	continuants, occurrents	occurrent is dependent upon continuant
immaterial	(physical) objects	material objects, immaterial objects	immaterial objects are par material objects

Mapping the metaphysical architecture landscape

Horizontal Aspect: Stratification Clustering

Stratification: (in the survey) unifying and separating clusters at each end of the spectrum – especially for the more ontologically commited.

Dependencies seem lead to attraction one way or the other.

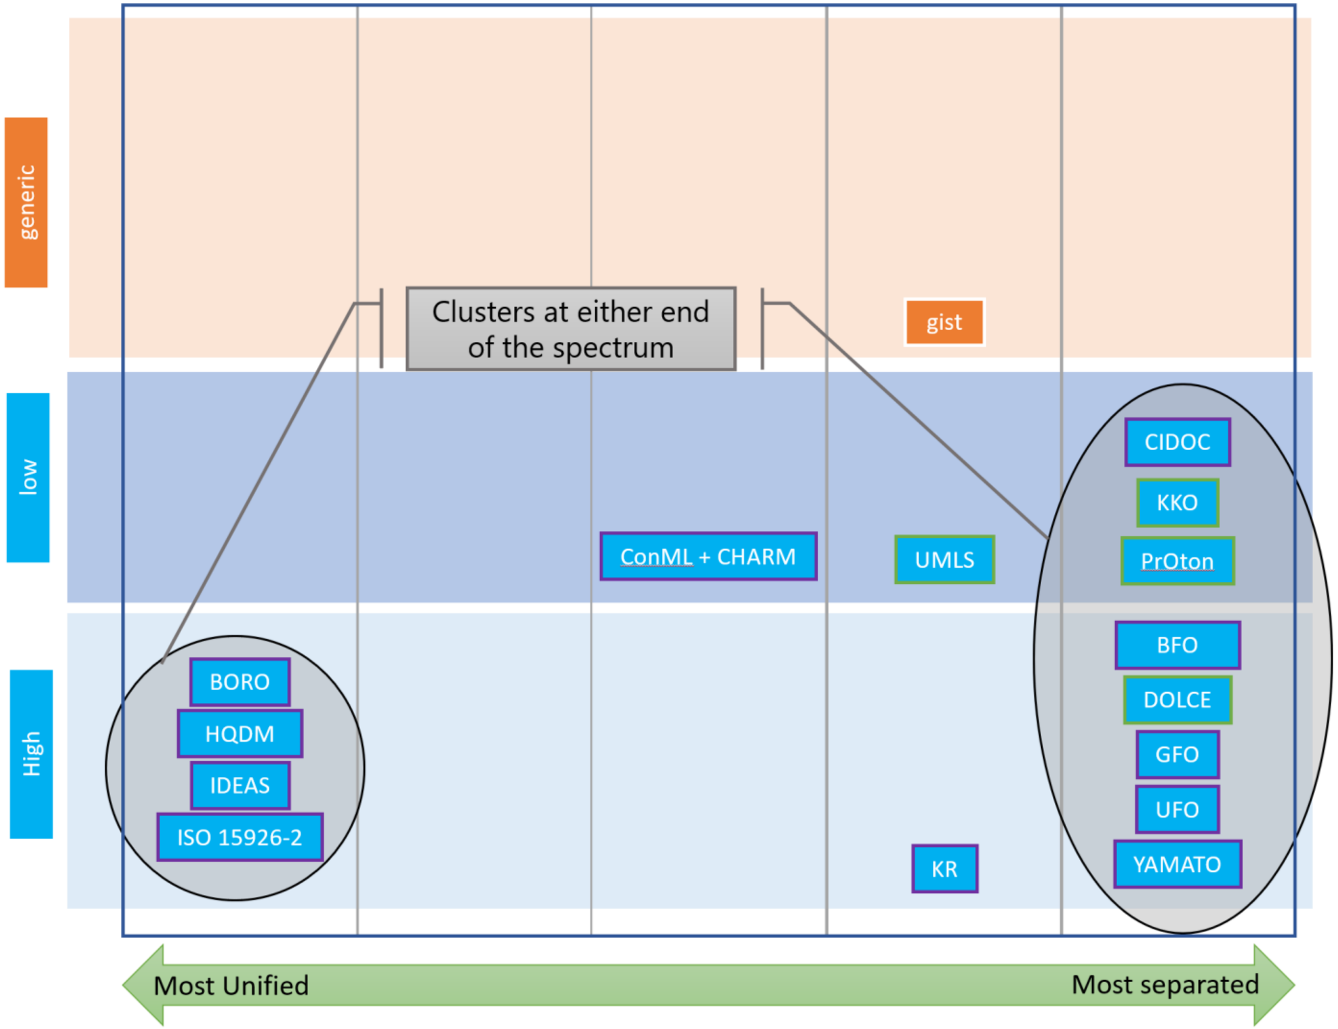

web-based: https://digitaltwinhub.co.uk/a-survey-of-top-level-ontologies/#a_survey_of_TLOs_contents

What kind of thing are foundational ontologies?

- Foundational ontologies
  - provide a metaphysical architecture
    - different top ontologies provide different metaphysical architectures
  - the architecture is underpinned by metaphysical choices
    - addressed in a coordinated way
    - the existence of the choice introduces indeterminacy
    - making the choice resolves the indeterminacy
  - resolving the metaphysical indeterminacy
    - makes translation and reference (more) determinate
    - enabling better interoperability

Summary

Summary

- Improving interoperability is key to improving computer ecosystems, and
  - ontologies enhance interoperability
    - by helping to enable the use of reference
      - to triangulate the interlingua translation
        - in the data exchange process
    - but these are undermined by ontological relativity’s indeterminacy
- Ontological relativity’s indeterminacy can be resolved by a foundational ontology
  - it can be resolved in multiple ways,
    - by foundational ontologies with different metaphysical architectures
      - in other words, there are multiple possible foundational ontologies
- Foundational ontologies
  - provide a metaphysical architecture
  - the architecture is underpinned by metaphysical choices
  - resolving these resolves metaphysical indeterminacy
    - makes translation and reference (more) determinate
    - enabling better interoperability

Questions

Business Object Reference Ontology

THE END.