Modelling Our World



With origins in antiquity and ancient philosophers speculating on the nature of being, "ontology" and "ontological modelling" have perhaps a reputation of being esoteric, impenetrable and only for the purist of academics.

In the last couple of decades however, there has been great work done to make ontological concepts and models readily accessible and to build the essential toolkits required to apply them at a practical level.

The Ontology Working Group (OWG) provides an introduction to ontology, explains how relevant and important it is and describes the OWG's current objectives in developing new ontological models that will be of benefit to everyone working in domains that involve the management of risk, such as safety and security.

What Is an ontology?

Isn't an ontology some esoteric formalisation, which is nice to have, but not really relevant or accessible to my day-to-day work on real systems?

Ontology is **everything** to do with your day-to-day work! When we build a system, we are modelling our world and this involves developing concepts, terms and their relationships – this is called an ontology, or more formally: "**a specification of a conceptualization**" [1].

Why do we need them?

Even if not being done explicitly, an ontology will always be developed – it's therefore never a case of whether or not an ontology is needed. If they are not explicitly developed, and developed well, you only end up with bad ontologies.

For the system to operate correctly, the model needs to be correct and consistent. Hopefully it is written down, but at worst, it can just be held in the minds of the designers.

Poor ontologies can lead to poor design, misunderstanding, ambiguities and inconsistencies. The system may be difficult to explain or may diverge from the real-world intended behaviour. Many of these issues may actually be difficult to spot or test and may lie latent for years until the right set of circumstances occurs (c.f. the NATS outage [2]). Consider, for example, two people with the same name and date of birth, which can occur in real life, but might not be handled well by a system.

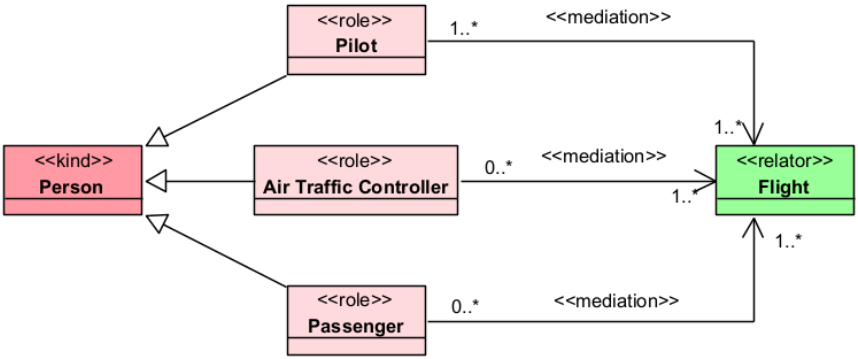
"It's never a case of whether or not an ontology is needed. If they are not explicitly developed, and developed well, you only end up with bad ontologies"

How do we capture models?

The idea of modelling relationships is not new and there are notations such as the **Unified Modelling Language (UML)** [3] to help us to do just that. However, UML can be too general for modelling ontologies and without a framework to impose some constraints, it is very easy to model things that are inconsistent or ambiguous.

This is why the OWG have chosen to use the **Unified Foundation Ontology (UFO)** [4]. This is a specialised framework for modelling in UML tailored for ontological modelling and implemented in **OntoUML** [5] (this is a customisation of UML with associated software support via application plugins).

The following diagram shows an example of a UFO model using OntoUML:



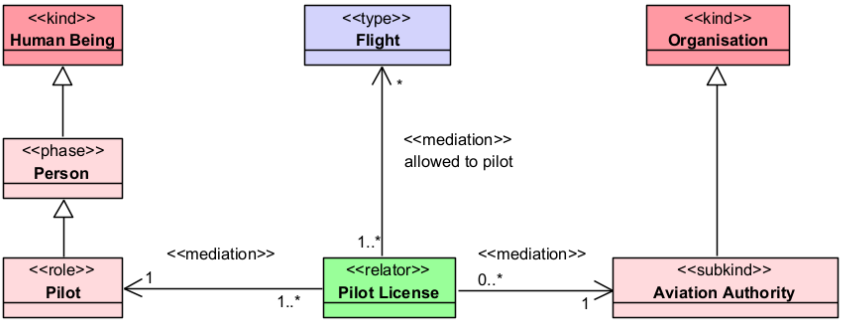
In the diagram, boxes (or more formally *classes)* represent 'things' in our world view. The ontological aspects are expressed as stereotypes (that's the part of the class definition with the Guillemets like this: «role»). Classes of the same stereotype share the same colour in

the diagram. Lines between the boxes show their relationships with other things. Arrows with a filled white ▽ triangle are *specialisations*, or in other words subtypes of the thing that is being pointed to (for example, a **Passenger** is a subset of all **Persons**). An open arrow → relationship shows a looser *association* between things and these also have stereotypes. Expression such as **0..*** and **1..*** show *cardinality* – that is the number of each thing that can be related. These particular ones are read as *zero or more*, and *one or more* respectively so for example, a **Flight** will have at least one **Pilot** but may have more, and those **Pilot**s may be involved in one or more **Flight**s. A particular **Flight** may have one or more **Passenger**s but may have none if the **Pilot** is the only occupant. A **Person** may also be a **Passenger** or more than one **Flight**.

Different Perspectives

This first model shows the concept of a **Flight**, which is hopefully well-known to everyone. In this model, the flight only exists as a relationship between a **Pilot**, and usually the **Air Traffic Controller** and **Passenger**s. In other words, there is no flight without a **Pilot**. **Passenger**s are often also on the flight (but not always!) and the **Air Traffic Controller** usually helps with the navigation (but not always!). Note that this model is likely under-constrained because it does not preclude, for example, **Pilot**s from being **Passenger**s on the same flight. Dealing with this type of issue require additional constraint logic, such as described by the OntoUML anti-pattern catalogue [6].

This perspective might provide a useful way of modelling flights at a small private airfield, where flights are created on an ad hoc basis when a **Pilot** and **Passenger**s are in the plane and want clearance to take off. There might be no prior concept of an actual flight until that moment. In this way of modelling, if the flight never took place, it never existed.

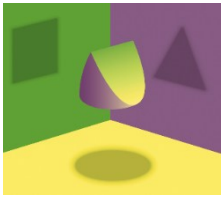


In this second model, **Flight** is considered a type of thing that has its own independent existence, not one that relies on relationships to exist. The **Pilot**'s relationship with the **Flight** is expressed in their ownership of a **Pilot License** to give them authority to pilot the aircraft.

This second perspective could be used to represent flights by a commercial transport airline. In this situation, **Flight**s are planned months or years in advance and might take place, be delayed or even be cancelled. The purely relational approach of the first diagram does not support the permanency of a flight even if it never took place. For example, the cancelled

flight still has relevancy for operational needs (charges and payments to airlines for airport services) or to manage passenger financial compensation.

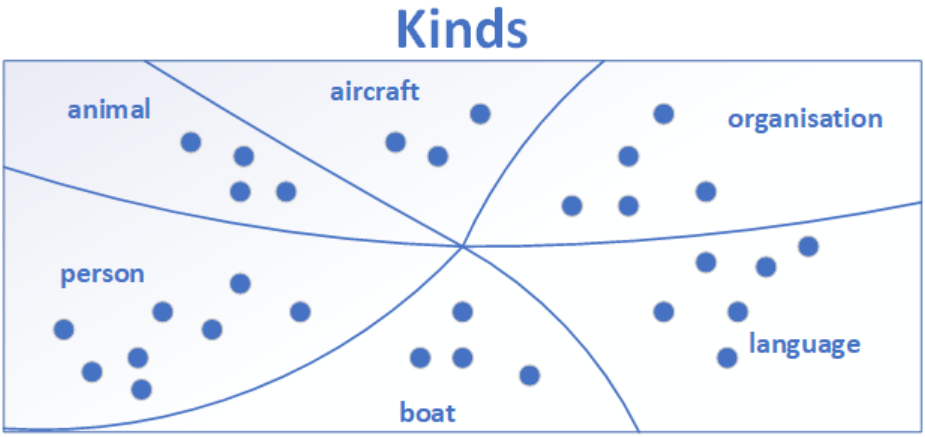
Importantly, **Flight** is modelled completely differently from before, but neither model is incorrect! Also note the differences in how a **Person** is modelled. In the first, a **Person** has identity and is a **Kind**, in the second, the **Person** is seen as a **Phase** of a human being.



These examples illustrate the idea that there can be several perspectives of something in the world, which are different but are not incorrect. Consider the various shadows cast by the object – which is the correct shadow for this object? They all are!

The Unified Foundation Ontology (UFO)

UFO comes in three parts: a top-level model for "things", (UFO-A), a model for "events" (UFO-B) and a model for "social" aspects (UFO-C). UFO-A is where we start normally as it relates to the objects in our world.

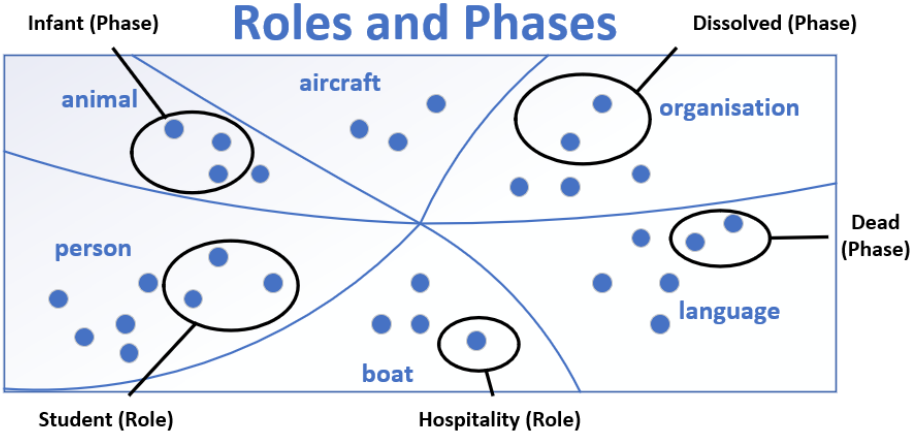


Kinds are a core UFO concept and reflect the objects we find around us like people, aircraft and animals. **Kind**s are objects that endure for their entire lifetime and define essential characteristics that instances must possess (called an identity principle). We can also distinguish between different instances of those **Kind**s: I am a person, you are a person and we are different people as we have different identities based on the properties we usually use to distinguish people (height, age, eye-colour, place of birth etc). Each 'dot' ● in the diagram indicates an individual instance of a particular **Kind**.

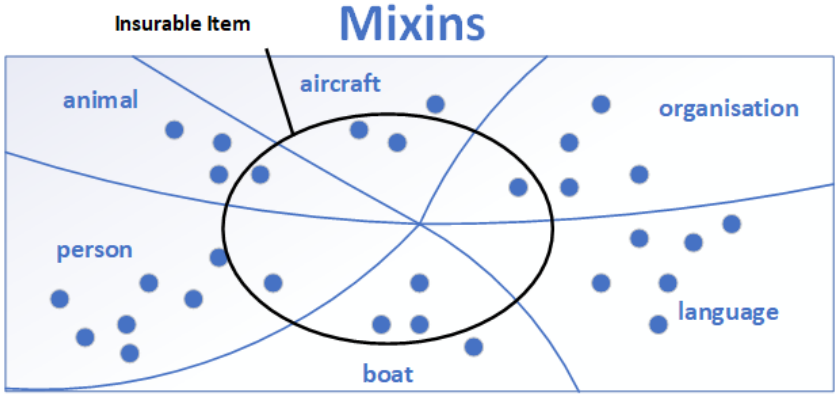
We might also talk about the **Roles** and **Phases** that **Kinds** might have during their lifetime. For example, a person may be a student for some part of their life (a **Role**), but will cease to be a student when they leave education. Similarly, a puppy is a **Phase** a dog will go through, which is transitory and only for the early part of their life. The following diagram shows how particular instances may exist in a particular **Phase** or **Role** as some point in their lifetime.

For example:

- An organisation may be dissolved at some point
- Some languages are no longer spoken as a native language by a community of people so are considered "Dead"
- A boat may at some point in its lifetime be used for hospitality (eg. parties, ceremonies etc.)



There are several other terms that become useful when modelling the world. One useful concept is that of a **Mixin**.



A **Mixin** is something we use to describe shared properties of objects that do not all belong to the same **Kind** and so can span multiple **Kind**s that do not share the same identity principles (a red boat does not have red hair colour, but both are red). The example shows a **Mixin** that some Kinds such as boats, aircraft, animals can be insured.

These terms along with the associated rules give us a powerful and consistent toolkit for describing important concepts and relationships in our world.

The Ontology Working Group (OWG)

The Ontology Working Group (OWG) has been set up to look specifically at developing an ontology of **Risk** and then subsequently **Value**. The OWG 2025 Objective is more formally:

Mission: Develop a clear, systematic terminology, and unambiguous model of risk and value concepts and their relationships.

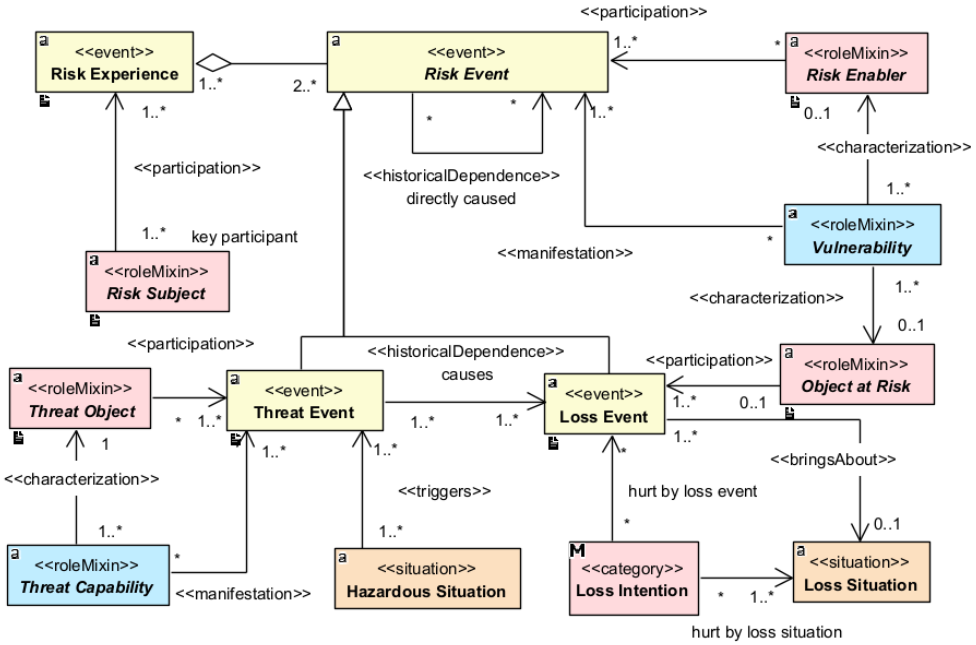
The group will do this by using:

- The Unified Foundation Ontology (UFO)
- Visual modelling tools (Visual Paradigm)
- Ontologically consistent modelling notations (OntoUML)
- Monitoring the work of other working groups (eg. ISO & IEC)

The main reason for pursuing this is to provide a common framework for expressing both **safety and security risk**. Safety and Security Risk management processes are similar but different (or in other words, they are different perspectives of the world), so there is much confusion that can arise when the two groups need to collaborate on projects.

Working Group Developments

The following diagram shows some of the work in progress by the OWG:



There have been many different attempts by other working groups and organisations to model risk, largely because, as we've seen, there are many different perspectives and ways of viewing the problem.

The OWG has assessed many different approaches but believe the most fruitful approach is through models that consider **risk** as an experience for participants and also to consider **value** as a key consideration when thinking about risk.

For example, any discussion around risk needs to really consider how participants *value* those items at risk, and indeed, how those who aim to exploit a vulnerability value *what they will gain* versus how they value *the liberties that may be lost* if they are caught.

It's hard to overstate the importance of this work, not only for the safety/security community where arguably ontologies *have* to be an essential aspect of the work we do. For example, it is difficult to see how a robust understanding of AI decision-making can be developed if not supported by formalised models of the concepts and relationships it operates on.

References

[1] T. R. Gruber, A Translation Approach to Portable Ontology Specifications, in Knowledge Acquisition, 5(2):199-220, Academic Press, 1993, http://www-ksl.stanford.edu/kst/what-is-an-ontology.html

[2] Air traffic control disruption hearing (British Airways), https://committees.parliament.uk/publications/42114/documents/209387/default/

[3] Unified Modelling Language, https://en.wikipedia.org/wiki/Unified_Modeling_Language

[4] Unified Foundational Ontology, https://en.wikipedia.org/wiki/Unified_Foundational_Ontology

[5] OntoUML, https://en.wikipedia.org/wiki/OntoUML

[6] OntoUML Anti-Pattern Catalogue (Relator Mediating Overlapping Types), https://ontouml.readthedocs.io/en/latest/anti-patterns/RelOver/index.html

The Ontology Working Group (OWG)

The OWG was set up in 2020 to help address issues with the lack of formalisation in terms and conceptual relationships used in standards and guidance documents. In particular, the similar but at times different, interpretations of terms and concepts used between safety and security domains is an ongoing cause of confusion and misunderstandings and presents barriers to more formal integration of these disciplines.

The OWG has its origins in the Data Safety Initiative Working Group (DSIWG), which, between 2016 and 2019, explored the development of an ontology model for the concepts described in the Data Safety Guidance (DSG). However, it was realised that the work couldn't be completed without a wider ontological model for risk, and so the OWG was established and is currently working on developing a clear, systematic terminology, and unambiguous model of risk and value concepts and their relationships.