ECAI 2010 Workshop on Spatio-Temporal Dynamics
16th August 2010
Lisbon, Portugal

The Formalities of Affordance

Antony Galton

University of Exeter, UK

Introduction: Affordance and Ecological Perception

J. J. Gibson's "Ecological" Theory of Perception

- The function of vision is not image formation but **information gathering**.
- The retinal image is just a means to this end, and must be understood in the context of a constantly varying succession of retinal images linked to the motion of the eye, the head, and the observer.
- The eye is an instrument for gathering information about the **layout of surfaces** in the environment within which the observer operates.
- It does this by detecting **invariants** underlying the constantly shifting flux of light impinging on the eye.
- **Motion** is an essential element of this: without motion, everything is an invariant.

Affordance

- The term "affordance" was invented by Gibson to refer to a potentiality for action (or inaction) offered to an agent by some feature of the environment.
- Examples: For a human being,
  - *A firm, more or less horizontal surface supported about 50cm above the surrounding ground, if sufficiently wide and deep, affords* **sitting***;*
  - *A sufficiently high and wide aperture in a more or less vertical solid surface affords* **entering***.*
- An affordance is a relation between an agent and its environment. For a given agent, the affordance appears as an intrinsic property of the surface layout of the environment.
- According to Gibson, we perceive surface layouts and their affordances *directly*: they are the primary objects of perception (not "sense data", "inner images", etc.).

Perhaps the composition and layout of surfaces **constitute** *what they afford. If so, to perceive them is to perceive what they afford. This is a radical hypothesis, for it implies that the "values" and "meanings" of things in the environment can be directly perceived. Moreover, it would explain the sense in which values and meanings are external to the observer.*

J. J. Gibson, *The Ecological Approach to Visual Perception* (1979), p.127

The Goals of Affordance Research

- **Ecological questions:**
  - What is the role of affordances in the life of an individual?
  - How can affordances be used to explain behaviour?
  - How can they be exploited for improving the design of environments?
- **Ontological questions:**
  - How are individual affordances defined?
  - What kinds of affordance are there and how can they be classified?
  - How can their properties be formalised?
- **Aetiological questions:**
  - Where do affordances come from, i.e., how does the physical layout of surfaces determine the affordances they provide for any given class of creatures?

An Example: Doors

- Steedman (2002) provides an ontological analysis of the affordances associated with doors, formalised in a *linear dynamic event calculus*:
  - If you push on a closed door, it will open; if you push on an open door, it will close.
  - If a door is open, you can go through it; if it is closed, you cannot.
  - If you are inside, and go through a door, you end up outside; if you are outside, and go through a door, you end up inside.

An Example: Doors

- Steedman (2002) provides an ontological analysis of the affordances associated with doors, formalised in a *linear dynamic event calculus*:
  - If you push on a closed door, it will open; if you push on an open door, it will close.
  - If a door is open, you can go through it; if it is closed, you cannot.
  - If you are inside, and go through a door, you end up outside; if you are outside, and go through a door, you end up inside.
- An ecological analysis would focus on the role of doors in providing passageways and barriers to regulate the movement of people around buildings, etc.

An Example: Doors

- Steedman (2002) provides an ontological analysis of the affordances associated with doors, formalised in a *linear dynamic event calculus*:
  - If you push on a closed door, it will open; if you push on an open door, it will close.
  - If a door is open, you can go through it; if it is closed, you cannot.
  - If you are inside, and go through a door, you end up outside; if you are outside, and go through a door, you end up inside.
- An ecological analysis would focus on the role of doors in providing passageways and barriers to regulate the movement of people around buildings, etc.
- An aetiological analysis would describe the physical characteristics that something must have in order to function as (i.e., possess all the relevant affordances of) a door.

Image Schemas

- **Image schemas** (Talmy, Johnson, Lakoff) are recurring patterns which we use to structure our understanding of the world. They are presumed to play a fundamental role in human cognition and language.
- Important examples are CONTAINER and PATH.
- An image schema may be thought of as a coordinated bundle of affordances:
  - Primary affordances of a container: *putting things in, taking things out*
  - Secondary (optional) affordances of a container: *moving things* (by moving the container they're in), *concealing things, protecting things, storing things*.
- At least in many cases, an image schema may be characterised in terms of the affordances of its instances.

Quantitative and Qualitative Determinants of Affordance

- Warren (1995) showed experimentally that for a set of stairs to be climbable for a given human subject, the ratio between the vertical height of each step and the subject's own leg-length should not exceed 0.88.
- Such numerical measurements are obviously important in determining the affordances of different surface layouts.
- However, the relevant quantitative questions *cannot even be asked* unless suitable qualitative conditions are satisfied first.
- For a flight of stairs there must exist an appropriately configured sequence of alternating horizontal surfaces and vertical displacements — otherwise there is nothing to measure!

QUALITATIVE

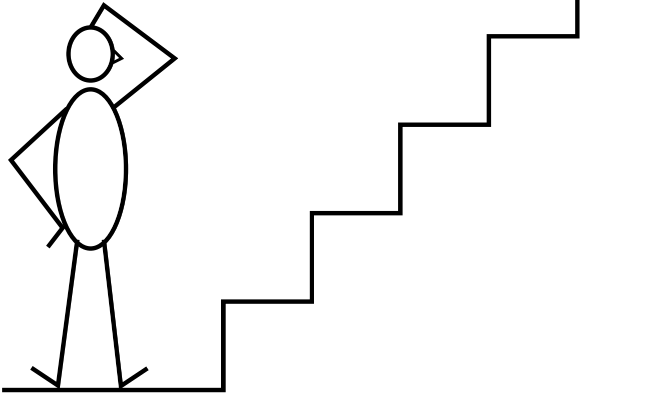

QUANTITATIVE

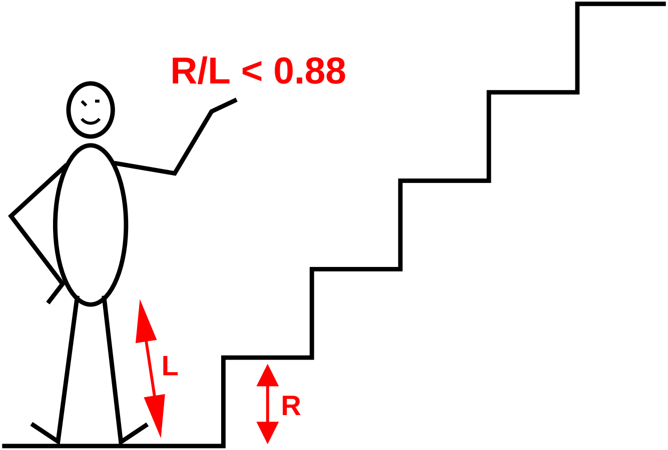

The Goal of this Research

Outline of a research programme:

- To investigate the qualitative conditions that must be satisfied by a surface layout in order for it to have some specified affordance.
- In particular, to determine to what extent the affordance-generating features of surface layouts can be specified in terms on simple qualitative calculi such as the RCC systems.

In the remainder of this paper we will focus on one particular case, the affordance of *containment*.

Formal Preliminaries

Spatial Regions

- We will use standard RCC relations, specifically $P, PP, TP, TPP, EC, DC, O, PO$.

Spatial Regions

- We will use standard RCC relations, specifically $P, PP, TP, TPP, EC, DC, O, PO$.
- Spatial relations between objects are expressed using RCC relations between their positions.

Spatial Regions

- We will use standard RCC relations, specifically $P, PP, TP, TPP, EC, DC, O, PO$.
- Spatial relations between objects are expressed using RCC relations between their positions.
- The position of object o at time t is denoted $pos(o, t)$. This is a spatial region which coincides with the spatial extent of o at t.

Spatial Regions

- We will use standard RCC relations, specifically $P, PP, TP, TPP, EC, DC, O, PO$.
- Spatial relations between objects are expressed using RCC relations between their positions.
- The position of object o at time t is denoted $pos(o, t)$. This is a spatial region which coincides with the spatial extent of o at t.
- Other spatial notions we will need:
  - Boundary: ∂r

Spatial Regions

- We will use standard RCC relations, specifically $P, PP, TP, TPP, EC, DC, O, PO$.
- Spatial relations between objects are expressed using RCC relations between their positions.
- The position of object o at time t is denoted $pos(o, t)$. This is a spatial region which coincides with the spatial extent of o at t.
- Other spatial notions we will need:
  - Boundary: ∂r
  - Convex hull: $cvhull(r)$

Spatial Regions

- We will use standard RCC relations, specifically $P, PP, TP, TPP, EC, DC, O, PO$.
- Spatial relations between objects are expressed using RCC relations between their positions.
- The position of object o at time t is denoted $pos(o, t)$. This is a spatial region which coincides with the spatial extent of o at t.
- Other spatial notions we will need:
  - Boundary: ∂r
  - Convex hull: $cvhull(r)$
  - Congruence: $Congruent(r_1, r_2)$

Spatial Regions

- We will use standard RCC relations, specifically $P, PP, TP, TPP, EC, DC, O, PO$.
- Spatial relations between objects are expressed using RCC relations between their positions.
- The position of object o at time t is denoted $pos(o, t)$. This is a spatial region which coincides with the spatial extent of o at t.
- Other spatial notions we will need:
  - Boundary: ∂r
  - Convex hull: $cvhull(r)$
  - Congruence: $Congruent(r_1, r_2)$
  - Union $(r_1 \cup r_2)$, intersection $(r_1 \cap r_2)$, and difference $(r_1 \setminus r_2)$ (understood set-theoretically or mereologically)

Physical Objects

- Physical objects include
  - Material objects (made of matter)
  - Non-material objects (dependent on material objects, but not themselves material)
- Non-material objects include holes and concavities in material objects, e.g., the space within a container.
- The convex hull of an object (as opposed to that of a region) is also a non-material object.
- These non-material objects associated with material objects are *not* spatial regions: their location, shape, and size depend on those of their hosts, and may change if the latter do. Spatial regions do not change.

Functions and Relations on Objects

Many of the functions and relations used for spatial regions have analogues which apply to objects; we notate these using starred symbols:

Functions and Relations on Objects

Many of the functions and relations used for spatial regions have analogues which apply to objects; we notate these using starred symbols:

- The convex hull $cvhull^*(o)$ of an object o obeys the rule:

$$\forall t(pos(cvhull^*(o), t) = cvhull(pos(o), t)).$$

Functions and Relations on Objects

Many of the functions and relations used for spatial regions have analogues which apply to objects; we notate these using starred symbols:

- The convex hull $cvhull^*(o)$ of an object o obeys the rule:

$$\forall t(pos(cvhull^*(o), t) = cvhull(pos(o), t)).$$

- Objects are congruent if their positions are:

$$\forall t(Congruent^*(o_1, o_2, t) \leftrightarrow Congruent(pos(o_1, t), pos(o_2, t))).$$

Functions and Relations on Objects

Many of the functions and relations used for spatial regions have analogues which apply to objects; we notate these using starred symbols:

- The convex hull $cvhull^*(o)$ of an object o obeys the rule:

$$\forall t(pos(cvhull^*(o), t) = cvhull(pos(o), t)).$$

- Objects are congruent if their positions are:

$$\forall t(Congruent^*(o_1, o_2, t) \leftrightarrow Congruent(pos(o_1, t), pos(o_2, t))).$$

- The boundary $\partial^*(o)$ of an object is a lower-dimensional object satisfying

$$\forall t(P(\partial(pos(o, t)), pos(\partial^*(o), t)))$$

Physical vs Spatial Connection

Modified RCC relations P^*, PP^*, TP^*, ... apply to physical objects, with connection understood to mean physical attachment rather than spatial contiguity (objects are EC^* if actually joined together). Note that these must be relativised to time.

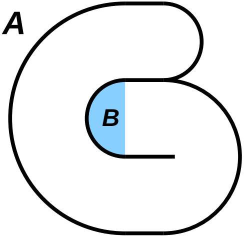

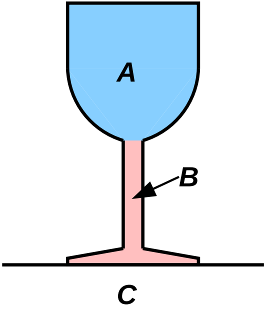

$TPP^*(B, A, t)$
$NTPP(pos(B, t), pos(C, t))$
$PP(\partial(pos(A, t)), pos(\partial^*(A), t))$

$EC^*(A, B, t), DC^*(B, C, t)$
$EC(pos(B, t), pos(C, t))$
$EC(pos(A, t), pos(B, t))$

Non-interpenetrability

The Principle of Non-Interpenetrability for Material Objects:

- If at any time two material objects do not overlap (i.e., have no common part), then their positions at that time cannot overlap either:

$$Material(o_1) \wedge Material(o_2) \wedge \neg O^*(o_1, o_2, t) \rightarrow \neg O(pos(o_1, t), pos(o_2, t)).$$

Non-interpenetrability

The Principle of Non-Interpenetrability for Material Objects:

- If at any time two material objects do not overlap (i.e., have no common part), then their positions at that time cannot overlap either:

$$Material(o_1) \wedge Material(o_2) \wedge \neg O^*(o_1, o_2, t) \rightarrow \neg O(pos(o_1, t), pos(o_2, t)).$$

- If at least one of the objects is non-material, then overlap is possible, e.g., a material object located in a (non-material) cavity in another material object. Thus we can have:

$$Material(o_1) \wedge \neg Material(o_2) \wedge \neg O^*(o_1, o_2, t) \wedge O(pos(o_1, t), pos(o_2, t)).$$

Time

- We use the Method of Temporal Arguments.
- We write

$$S(t)$$

  to mean that state S holds at time t.
- We write

$$E(t_1, t_2)$$

  to mean that an event of type E occurs over the interval $[t_1, t_2]$.

Modality

- We use a notion of modality based on *possible futures*:

  > $\Diamond P$ is true at t if and only if there is some possible future of t such that, if that future is the actual future, then P is true at t.

- This is *not unproblematic*! What is a "possible" future?

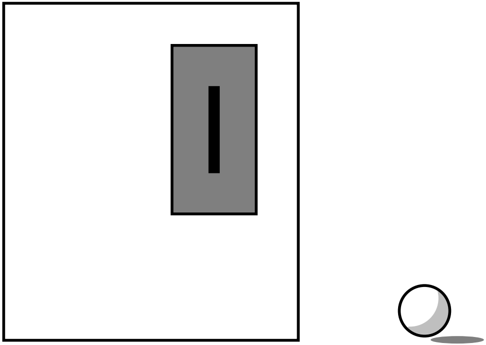

Can the ball fit into the slot?

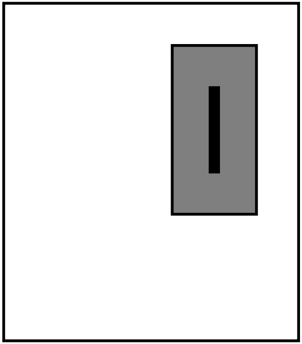

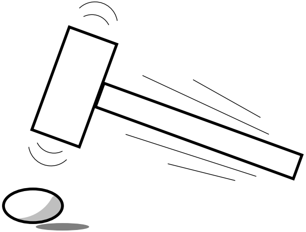

Hammer it into a disc . . .

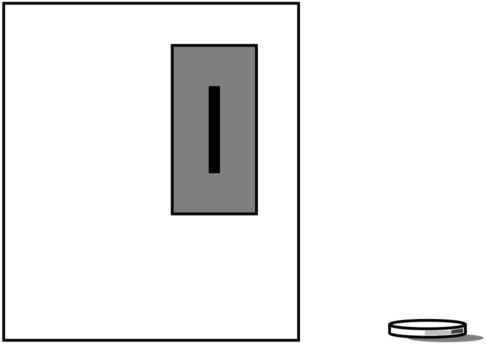

It can now!

Possible futures revisited

- The possible futures we need to found modality on must not be too *affordance-disrupting*.
- We should not (normally) allow hammering a ball into a disc.
- But we should allow, e.g., folding a letter to fit it into an envelope.
- For the purposes of assessing affordances, some objects should be regarded as **rigid**.

Rigidity

An object is rigid if all of its possible positions are congruent:

$$\begin{aligned} \mathit{Rigid}(o) \ =_{\mathrm{df}} \ \forall t \forall t' \forall r \forall r' (\ & \Diamond(\mathit{pos}(o,t) = r) \wedge \\ & \Diamond(\mathit{pos}(o,t') = r') \ \rightarrow \\ & \mathit{Congruent}(r,r') \ ) \end{aligned}$$

Here too, possibility is to be understood as 'non-affordance-disrupting'.

Different degrees of disruption correspond to different degrees of rigidity.

Containers and Containment

What is a Container?

- A container is a material object which can contain other material objects
- But what does it mean for one object to contain another?
  - My pocket contains coins
  - The jug contains water
  - The vase contains flowers (and water)
  - The car contains people
- For simplicity, we shall restrict ourselves to "full containment", in which the contained object is "right inside" the container — as in all the examples above except the flowers in the vase.

The Contained Space of a Container

- An container has a **contained space**. This is a non-material object, dependent on the container, within which an object has to be in order to be contained by the container. (The term is due to Pat Hayes.)
- We shall write $cs(x)$ to denote the contained space of container x.
- The contained space is always
  - joined to x: $\forall t \; EC^*(cs(x), x, t)$;
  - part of the convex hull of x: $\forall t \; P^*(cs(x), cvhull^*(x), t)$.

Open and Closed Containers

- A container is **closed** at time t if the boundary of its contained space is part of the boundary of the container itself (its *inner boundary*):

$$\mathit{Closed}(x, t) =_{\mathrm{df}} \mathit{Container}(x) \wedge P^*(\partial^*(cs(x)), \partial^*(x), t)).$$

  A container can be closed at some times and open (i.e., not closed) at others.

- The **portals** of an open container are the connected components of $\partial^*(cs(x)) \setminus \partial^*(x)$. These are non-material objects, dependent on x, which exist whenever x is not closed.

Containment

- 'At time t, container c contains object o' means that the position of o at t is part of the position of the contained space of c at t:

$$Contains(c, o, t) =_{\mathrm{df}} P(pos(o, t), pos(cs(c), t)).$$

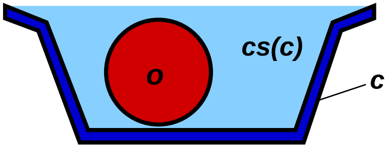

Containment

- 'At time t, container c contains object o' means that the position of o at t is part of the position of the contained space of c at t:

$$Contains(c, o, t) =_{\mathrm{df}} P(pos(o, t), pos(cs(c), t)).$$

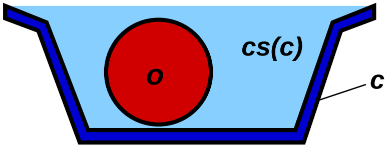

- 'Container c can contain object o' means that it is possible for c to contain o now or in the future:

$$CanContain(c, o, t) =_{\mathrm{df}} \exists t'(t \leq t' \wedge \Diamond Contains(c, o, t')).$$

Containment and Rigidity

- The contained space of a rigid container is also rigid:

$$Container(c) \wedge Rigid(c) \rightarrow Rigid(cs(c))$$

- It is easy to prove that a rigid container can contain a rigid body only if the latter is congruent to part of the contained space of the former:

$$CanContain(c, o, t) \wedge Rigid(c) \wedge Rigid(o) \rightarrow \exists u \forall t (Congruent^*(o, u, t) \wedge P^*(u, cs(c), t)))$$

- But more generally both the container and what it contains may be either rigid or non-rigid.


PURE CEYLON
TEA


PURE CEYLON
TEA

Being in a container is not the same as *entering* it . . .

Being in a container is not the same as *entering* it . . .

Entering a Container

Entering and Leaving

- An important part of the affordance of containment is that, in principle, objects can enter or leave (or be put in or taken out of) containers.
- Suppose that o is outside c at t_0 and inside c at t_1.
- Over the interval $[t_0, t_1]$, both o and c may undergo changes in both position and shape.
- The sequence of positions/shapes assumed by an object over an interval constitutes a **trajectory**.
- A condition for o to come to be inside c is that suitable trajectories for both objects exist, compatible with whatever rules for continuity, rigidity, non-interpenetrability, etc, are in force.

Trajectories

- A **trajectory** *traj* is a continuous sequence of spatial regions, represented by a continuous function

$$\mathit{traj} : [0, 1] \rightarrow \mathcal{R},$$

where $\mathcal{R}$ is the set of all spatial regions.

Trajectories

- A **trajectory** *traj* is a continuous sequence of spatial regions, represented by a continuous function

$$traj : [0, 1] \rightarrow \mathcal{R},$$

  where $\mathcal{R}$ is the set of all spatial regions.
- How is "continuous" defined? Given a metric Δ on the space of spatial regions, continuity of trajectory *traj* is defined in the usual way as

$$\forall t \in [0,1] \forall \epsilon > 0 \exists \delta > 0 \forall t' \in [0,1](\\ |t - t'| < \delta \rightarrow \Delta(traj(t), traj(t')) < \epsilon).$$

Following a trajectory

The following formula says that object o follows trajectory *traj* over the interval $[t_0, t_1]$:

$$Follows(o, traj, t_0, t_1) =_{\mathrm{df}} \forall t \left(t_0 \leq t \leq t_1 \rightarrow pos(o, t) = traj\left(\frac{t - t_0}{t_1 - t_0}\right)\right)$$

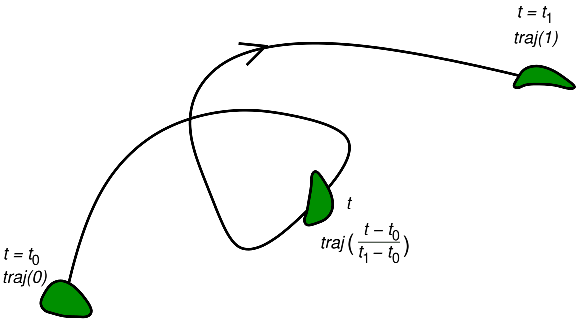

Continuity of Motion

The motion of o over the interval $[t_0, t_1]$ is continuous so long as over that interval it follows a (continuous) trajectory from its position at t_0 to its position at t_1.

We require the motion of any object to be continuous over any interval within its lifetime:

$$\begin{aligned}[t, t'] \subseteq \mathit{lifetime}(o) \rightarrow \quad & \exists \mathit{traj}(\mathit{traj}(0) = \mathit{pos}(o, t) \wedge \\ & \qquad \mathit{traj}(1) = \mathit{pos}(o, t') \wedge \\ & \qquad \mathit{Follows}(o, \mathit{traj}, t, t'))\end{aligned}$$

Entering a Container: Initial position

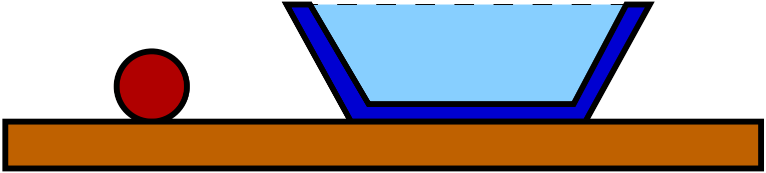

Entering a Container: Just outside

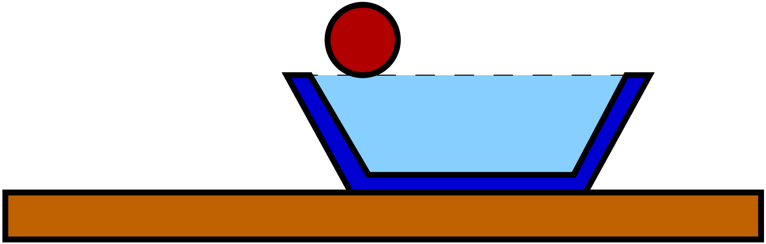

Entering a Container: Entering

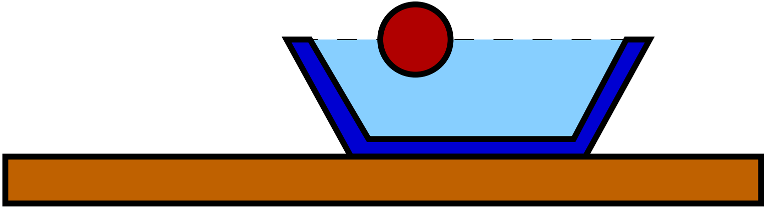

Entering a Container: Just inside

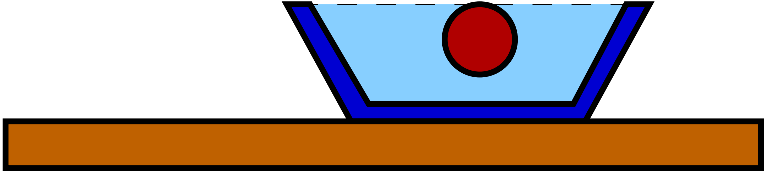

Entering a Container: Final position

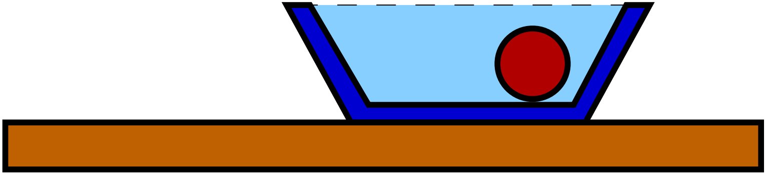

Entering a container

- We will concentrate on the actual entering, i.e., the transition between "just outside" to "just inside".
- For o to enter c over the interval $[t_1, t_2]$, o and c must follow trajectories such that at the start, o is EC to the contained space of c, and at the end, it is TPP. In between, the relation between the two must be neither EC nor TPP.
- The following formula expresses this:

$$
\begin{aligned}
&Enters(o, c, t_0, t_1) \ =_{\mathrm{df}} \\
&\quad \exists traj_o \exists traj_c(\\
&\quad\quad Follows(o, traj_o, t_0, t_1) \wedge Follows(c, traj_c, t_0, t_1)) \wedge \\
&\quad \forall t(t_0 \leq t \leq t_1 \ \rightarrow \\
&\quad\quad EC(pos(o, t), pos(cs(c), t)) \leftrightarrow t = t_0 \ \wedge \\
&\quad\quad TPP(pos(o, t), pos(cs(c), t)) \leftrightarrow t = t_1)
\end{aligned}
$$

Remarks on the definition of $Enters(o, c, t_0, t_1)$

- From non-interpenetrability, we always have $\neg O(pos(o,t), pos(c,t))$, so this does not need to be stated explicitly.

Remarks on the definition of $Enters(o, c, t_0, t_1)$

- From non-interpenetrability, we always have $\neg O(pos(o, t), pos(c, t))$, so this does not need to be stated explicitly.
- From the transition diagram for RCC, for $t_0 < t < t_1$, we have $PO(pos(o, t), pos(cs(c), t))$. If o is a one-piece object, this means that o must intersect a portal of c during this period.

Remarks on the definition of $Enters(o, c, t_0, t_1)$

- From non-interpenetrability, we always have $\neg O(pos(o, t), pos(c, t))$, so this does not need to be stated explicitly.
- From the transition diagram for RCC, for $t_0 < t < t_1$, we have $PO(pos(o, t), pos(cs(c), t))$. If o is a one-piece object, this means that o must intersect a portal of c during this period.
- Both o and c may move, and, if non-rigid, change shape during the entering process — all this is accounted for by the trajectories $traj_c$ and $traj_o$.

Remarks on the definition of $Enters(o, c, t_0, t_1)$

- From non-interpenetrability, we always have $\neg O(pos(o,t), pos(c,t))$, so this does not need to be stated explicitly.
- From the transition diagram for RCC, for $t_0 < t < t_1$, we have $PO(pos(o,t), pos(cs(c),t))$. If o is a one-piece object, this means that o must intersect a portal of c during this period.
- Both o and c may move, and, if non-rigid, change shape during the entering process — all this is accounted for by the trajectories $traj_c$ and $traj_o$.
- It seems obvious that to move from a position outside c to a position inside c, o must enter c. This needs to be proved!

$$\neg O(pos(o,t), pos(cs(c),t)) \wedge Contains(c,o,t') \rightarrow \exists t_0 \exists t_1 (t \leq t_0 < t_1 \leq t' \wedge Enters(o,c,t_0,t_1))$$

The Affordance of Entering

- $\mathit{CanEnter}(o, c, t) =_{\mathrm{df}} \exists t' \Diamond \mathit{Enters}(o, c, t, t')$
- How is the affordance of entry related to the affordance of containment, i.e., how is *CanEnter* related to *CanContain*?
- Conjecture: $\neg O(\mathit{pos}(o, t), \mathit{pos}(\mathit{cs}(c), t)) \rightarrow (\mathit{CanContain}(o, c, t) \leftrightarrow \exists t'(t \leq t' \wedge \mathit{CanEnter}(o, c, t'))).$
- Entry gives a *lower-level view* of the affordance of containment.
- Note: To get the ship into the bottle, we might have to dismantle it first and then reassemble it in the bottle — this raises the ontological question of whether the ship can be referred to even when in the dismantled state . . .

Entry at a Portal

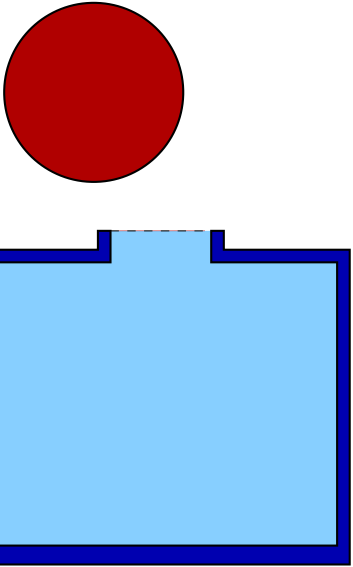

Entry at a Portal

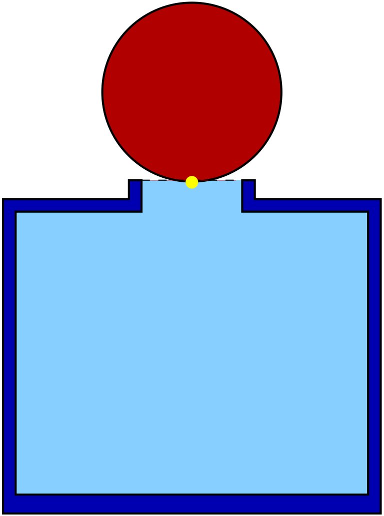

Entry at a Portal

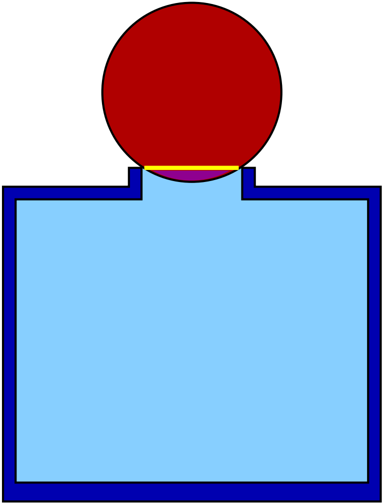

Entry at a Portal

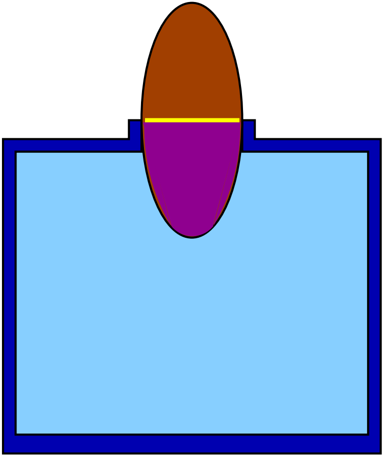

Entry at a Portal

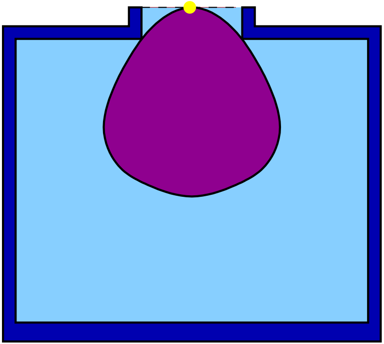

Entry at a Portal

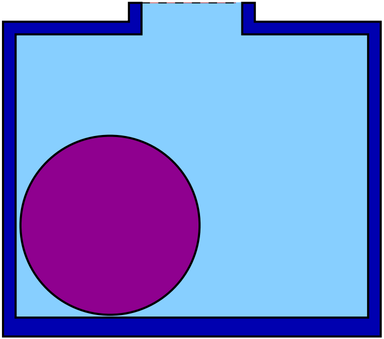

Entry at a Portal

- While o is entering c, we can distinguish the part of o still outside c, and the part already in c. Their common boundary must lie within a portal of c, by non-interpenetrability.
- In fact we need only consider the *positions* of these parts, say

$$\begin{aligned} r_1(t) &= pos(o,t) \cap pos(cs(c),t) \\ r_2(t) &= pos(o,t) \setminus pos(cs(c),t) \end{aligned}$$

- Then we must have

$$Enters(o,c,t_0,t_1) \wedge t_0 < t < t_1 \;\rightarrow\; P(\partial r_1(t) \cap \partial r_2(t),\; pos(\partial^* cs(c) \setminus \partial^* c, t))$$

- $\partial r_1(t) \cap \partial r_2(t)$ is the position of a **cross-section** of o. To enter c, o must be able to fit a continuous series of its cross-sections into a portal of c. This is implicit in the affordance of entering.

Conclusions and Further Work

Specific Summary

Specific Summary

- Initial goal: To characterise formally the conditions under which the affordance of containment exists.

Specific Summary

- Initial goal: To characterise formally the conditions under which the affordance of containment exists.
- High-level characterisation of containment: $Contains(c, o, t)$ and its affordance: $CanContain(c, o, t)$

Specific Summary

- Initial goal: To characterise formally the conditions under which the affordance of containment exists.
- High-level characterisation of containment: $Contains(c, o, t)$ and its affordance: $CanContain(c, o, t)$
- Middle-level characterisation in terms of entering: $Enters(o, c, t_0, t_1)$ and its affordance: $CanEnter(o, c, t)$

Specific Summary

- Initial goal: To characterise formally the conditions under which the affordance of containment exists.
- High-level characterisation of containment: $Contains(c, o, t)$ and its affordance: $CanContain(c, o, t)$
- Middle-level characterisation in terms of entering: $Enters(o, c, t_0, t_1)$ and its affordance: $CanEnter(o, c, t)$
- Low-level characterisation of entering in terms of portals and cross-sections.

Summary of General Approach

Summary of General Approach

- **Initial goal**: To define what it means for an object or collection of objects to afford some action A to an object o.

Summary of General Approach

- **Initial goal**: To define what it means for an object or collection of objects to afford some action A to an object o.
- High-level definition of the affordance is a modalised version of the definition of what it means for o actually to perform A.

Summary of General Approach

- **Initial goal**: To define what it means for an object or collection of objects to afford some action A to an object o.
- High-level definition of the affordance is a modalised version of the definition of what it means for o actually to perform A.
- By invoking general principles such as continuity and non-interpenetrability we tease out successively lower-level conditions for the affordance to exist.

Summary of General Approach

- **Initial goal**: To define what it means for an object or collection of objects to afford some action A to an object o.
- High-level definition of the affordance is a modalised version of the definition of what it means for o actually to perform A.
- By invoking general principles such as continuity and non-interpenetrability we tease out successively lower-level conditions for the affordance to exist.
- We thereby approach by stages the **final goal**, to specify what it is about any particular physical layout that results in its having the affordances it does.

Summary of General Approach

- **Initial goal**: To define what it means for an object or collection of objects to afford some action A to an object o.
- High-level definition of the affordance is a modalised version of the definition of what it means for o actually to perform A.
- By invoking general principles such as continuity and non-interpenetrability we tease out successively lower-level conditions for the affordance to exist.
- We thereby approach by stages the **final goal**, to specify what it is about any particular physical layout that results in its having the affordances it does.
- This will then enable us to explain how, in Gibson's words, we are able to perceive affordances directly: by perceiving these lower-level properties of the physical layout.

A Challenge for AI and Knowledge Representation

- We have only looked at containment here, and even in this case there is still a lot of work to do in order to derive the desired lower-level characterisations from the high-level definitions using an appropriate set of general principles.

A Challenge for AI and Knowledge Representation

- We have only looked at containment here, and even in this case there is still a lot of work to do in order to derive the desired lower-level characterisations from the high-level definitions using an appropriate set of general principles.
- That leaves everything else to investigate, such as affordances for
  - shifting
  - lifting
  - hiding
  - opening
  - closing
  - climbing
  - grasping
  - ...

THE END