

THE DOMAIN OF POLYFLAPS AND OTHER DOMAINS FOR ACTING AND LEARNING

(This file is now <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps/>)

Also available in PDF format as

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps/polyflaps.pdf>

Investigation of scenarios to drive ontology-formation, perceptual development, development of action, and growth of an information processing architecture in an 'altricial' robot in a 3-D environment. An artificial 3-D domain with a number of interesting and challenging features is proposed as a domain for robotics research in perception, learning, action, planning, ontology formation, and perhaps later on natural language communication (after the robot has something to communicate about).

Written by [Aaron Sloman](#), on the basis of discussions with Jackie Chappell (Biosciences Birmingham), Jeremy Wyatt, Marek Kopicki, Nick Hawes, Somboon Hongeng, in [the Birmingham CoSy Project\[1\]](#), Tony Cohn (Leeds), David Hogg (Leeds), and other members of the CoSy Project (www.cognitivesystems.org).

The germ of the idea was born in a conversation between Marek Kopicki and Aaron Sloman.

INTRODUCTION TO POLYFLAPS

The domain of polyflap configurations illustrates the notion of a domain of related configurations of objects that can be seen, explored, manipulated, and progressively learnt about, with or without a teacher. It is just one of many domains, including many that occur in nature and many that are products of a culture (including industrial products).

Domains that occur naturally include configurations of water, sand, mud, stones, trees, grass, other plants, and many kinds of animals. Some of the objects naturally move autonomously while others move only when manipulated by something else or blown by the wind, etc.

Non-natural domains include things like [string](#), [string plus other materials](#), paper, cloth, cotton-wool, plasticene, toy blocks, dolls and their clothing, toy vehicles, cutlery, crockery, and many kinds of construction kits, e.g. Lego, [Meccano](#), Tinker toys, Fischer technik, etc. (See Michael Adler's Meccano web site: <http://www.meccanotec.com/> and the International society of Meccanomen: <http://www.internationalmeccanomen.org.uk/> -- why only men???)

Combining domains provides new richer domains with novel capabilities that can be explored and used. E.g. combining string and paper produces possibilities like parcels that can be pushed, thrown, etc. without falling apart. Combining cutlery, crockery and various kinds of solid and liquid edible substances produces many new possible ways of consuming food and liquid.

There are also abstract domains, such as arithmetic, many kinds of games, words, phrases, sentences, poems, limericks, and many branches of mathematics. In the last few decades new domains have been developed that young children learn to interact with using computing devices even though their ancestors never encountered anything similar. Some primates can also learn to interact with them.

Our domain of polyflaps was invented partly to produce shapes and processes that might be easier for a robot to see than some of the more common shapes (e.g. avoiding curved surfaces and curved edges), while being cheap to make, and lightweight, and indefinitely extendable.

The [pictures shown below](#) are examples of 'polyflaps'. Each polyflap is made from a single sheet of paper cut in a polygonal shape (convex or concave, regular or irregular) and then folded once, but not folded flat, so as to form a 3-D structure. (Future examples may be made of cardboard, or some other material).

Any collection of polyflaps can clearly be extended indefinitely, forming more and more complex domains of perception and action, e.g.

- by folding an already used polygon shape in a new way
- by adding a new polygon shape with the same number of sides as used previously (e.g. varying relative line-lengths or angles)
- by adding a new polygon shape with more sides
- by providing larger numbers of polyflaps making more complex configurations possible, including new 3-D structures built from polyflaps
- by combining polyflaps with other physical objects of varying complexity
- by having two or more sources of light producing complex shadows
- by allowing more varieties of actions like pushing, grasping, lifting, rotating, placing on something else, inserting into a space, etc.

The set of possible polyflaps is very large (a) because there is a very large set of possible polygons, taking into account different numbers of sides, different angles between sides, and different lengths of sides, and (b) because for each such polygon a large number of more or less different polyflaps can be produced from one fold of between 0 and 180 degrees, to produce a 3-D object.

Different polyflaps will have different multi-stable possibilities depending on the shape of the original polygon and how it is folded.

So there is a large set of possible configurations of a single polyflap. Each configuration can be viewed from multiple viewpoints, adding to the complexity.

ADDED 22 Nov 2017: Note on triangle deformation

I now realise that videos of polyflaps, either taken with a moving camera while the polyflaps are stationary, or taken with a stationary camera while polyflaps are moving, or taken with moving camera while polyflaps are moving, will include rich examples of processes with mathematical structures that enable causal relations within the processes to be understood, even on a first viewing. I suspect this will turn out to be closely related to mechanisms required to explain the mathematical discovery process illustrated in this document on triangle deformation:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html>

Also relevant is this discussion of the role of perception of partial orderings and topological changes, in reducing or eliminating the need for numerical estimation in perception of affordances and how they change:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html>

A short paper on polyflaps as a domain for AI was included in the AAAI Fellows Symposium in 2006

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0605>

<http://www.aaai.org/Fellows/Papers/Fellows16.pdf>

Polyflaps as a domain for perceiving, acting and learning in a 3-D world (PDF)

CONTENTS

- [INTRODUCTION TO POLYFLAPS](#)
 - **ADDED 22 Nov 2017:** [Note on triangle deformation](#)
 - [BACKGROUND](#)
 - [POLYFLAP CONFIGURATIONS](#)
 - [THE PICTURES](#)
 - [POLYFLAP COMPETENCE REQUIRES 'GENERATIVE' COMPETENCE](#)
 - [WHAT IS THE POINT?](#)
 - [RELATION TO CURRENT AI VISION RESEARCH](#)
 - [PREREQUISITES FOR ACQUIRING POLYFLAPPERY COMPETENCE](#)
 - [DISCLAIMERS](#)
 - NB 1: Not necessarily the best or only domain
 - NB 2: Little or no role for object recognition initially
 - [PROGRESSIVE SCENARIOS](#)
 - [Static fragments](#)
 - [Larger fragments and combinations of fragments](#)
 - [Seeing fragments of processes](#)
 - [SURPRISE CAN DRIVE LEARNING](#)
 - [THINGS TO DO](#)
 - [BROADER RESEARCH ISSUES](#)
 - [INFLUENCES](#)
 - [SEE ALSO](#)
 - [NOTES AND NEWS](#)
 - [Acknowledgements](#)
-

BACKGROUND

We have been using the label 'altricial' somewhat confusingly to connote both

- the biologists' notion of a species whose members are born or hatched underdeveloped and incompetent
and also to refer to

- a pattern of cognitive development found in many such species that leads to unusual breadth and depth of cognitive competence most of which is not genetically determined but a result of learning about the environment.

In contrast, species described as 'precocial' by biologists are born or hatched relatively advanced physiologically, and behaviourally and cognitively competent (e.g. deer that run with the herd soon after birth, and chicks that peck for food soon after hatching). Their cognitive competence must be largely genetically determined, since there is no time to learn about the environment between birth and demonstrating the competence.

In fact, all species involve a mixture of both kinds of features as explained in [2]. While continuing to use the labels 'precocial' and 'altricial' to refer to types of species, we now propose to use different labels for the two main kinds of capabilities found in those species.

- **Preconfigured** capabilities (which may involve mechanisms, forms of representation, architectures, and behavioural competences based on them) are mainly determined genetically, though they need not be manifested at birth and there may be some calibration in use or re-shaping through reinforcement learning, usually a slow process.
- **Non-preconfigured** capabilities are discovered, created, constructed, developed by the individual during exploration and play, sometimes as a result of observing what others do sometimes as a side-effect of performing some old task in a new context where a surprise is encountered, or a feature of the task is noticed for the first time. Often, though not always, this is very fast learning, leading immediately to a new, lasting competence involving a structural change in behaviour, and possibly also a new understanding of structures and processes in the environment.

Note that the non-preconfigured capabilities depend on the operation of "higher level" pre-configured (i.e. genetically determined) bootstrapping capabilities. Adult members of altricial species seem to have a higher proportion of non-preconfigured competences than adult members of precocial species (born or hatched physically developed and behaviourally competent).

We are interested in explaining kinds of competence in the second category, competences that are not genetically determined but arise from play, exploration, observation and imitation using powerful genetically determined 'meta-level' learning mechanisms. This document is about some possible domains in which such mechanisms might operate in robots and animals.

DISCLAIMERS

-- NB 1: Not necessarily the best or only domain

No claim is made that the polyflap domain is the only or the best domain for such robotic experiments. It is hoped that in other projects other domains can be defined and explored and the results compared. E.g. we may find that the lack of curved surfaces and edges in the polyflap domain is a very serious restriction because a robot with an articulated arm will naturally produced curved movements, and will have to learn about them, which may be made more difficult in an environment where everything else is flat or straight.

-- NB 2: Little or no role for object recognition initially

Object recognition plays very little role in the early stages of the proposed experiments. The task is NOT to recognise the polyflaps from various angles and viewpoints, but rather to *perceive* them, which is something different and far more fundamental, since we can perceive and act on totally unfamiliar and unrecognisable objects and we certainly do not need to be able to name or describe spatial structures in order to perceive them. Finding out what that pre-linguistic competence, shared by young human children and possibly other non-linguistic species amounts, to is part of the goal of this research.

However, as the scenes perceived become more complex, cluttered or indistinct, the segmentation and recognition of familiar objects and the perception of spatial structure and affordances can use a mixed top-down and bottom-up perceptual processing mechanism, e.g. as frequently proposed by researchers of many sorts, e.g. [here](#)

POLYFLAP CONFIGURATIONS

There are of course even more configurations of two or more polyflaps and for each configuration infinitely many viewpoints. Some configurations include one polyflap partially or wholly supported by one or more other polyflaps (as illustrated in several of the pictures -- though some of the pictures may not be clear enough to be easily perceivable).

When the lighting produces shadows or highlights or merely variations of intensity the complexities increase. Likewise structure in the background.

In addition to the variation in the possible scenes and the possible images that can be produced in those scenes, there are multiple possible actions that can be performed by a robot with a two-fingered gripper on each polyflap in each of its stable configurations. Once the polyflap is grasped and lifted there are large numbers of additional configurations many not included in the stable configurations when the polyflap is resting.

Moreover from each configuration there is a large space of possible types of movement (combinations of translation and rotation) that can lead to many other configurations.

THE PICTURES

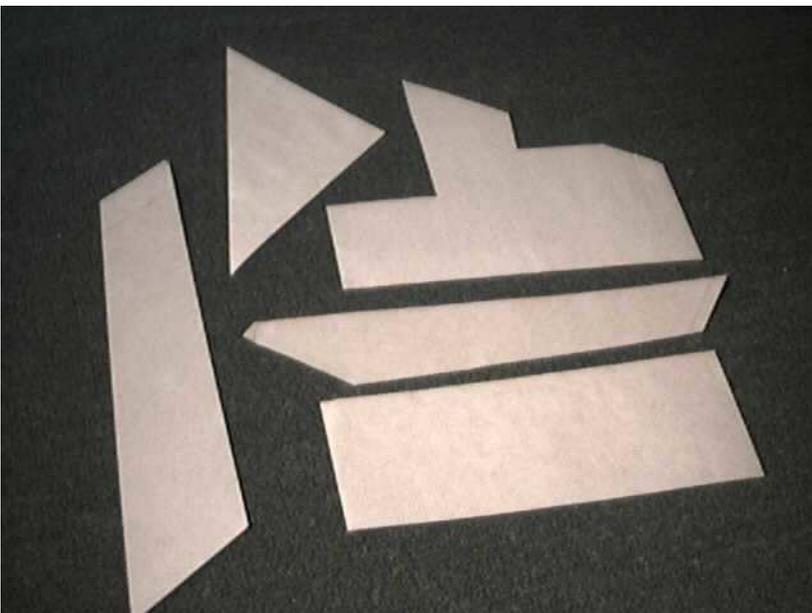
The following initial set of pictures is provided to indicate some of the properties of polyflaps. The polyflaps were made in a hurry from some old envelopes that happened to be available, which included some extraneous markings, and the pictures were taken using a simple webcam so they are not of very high quality. By making lighting more or less diffuse it would be possible to remove or sharpen shadows.

Two views of some unfolded polyflaps lying on the carpet

(which the robot will not necessarily encounter -- e.g. because they are too hard to pick up with a typical gripper, though at a later stage the robot might learn how to use the corner of a polyflap to lift an unfolded polyflap):



[Fig 1: flat1.jpg](#)

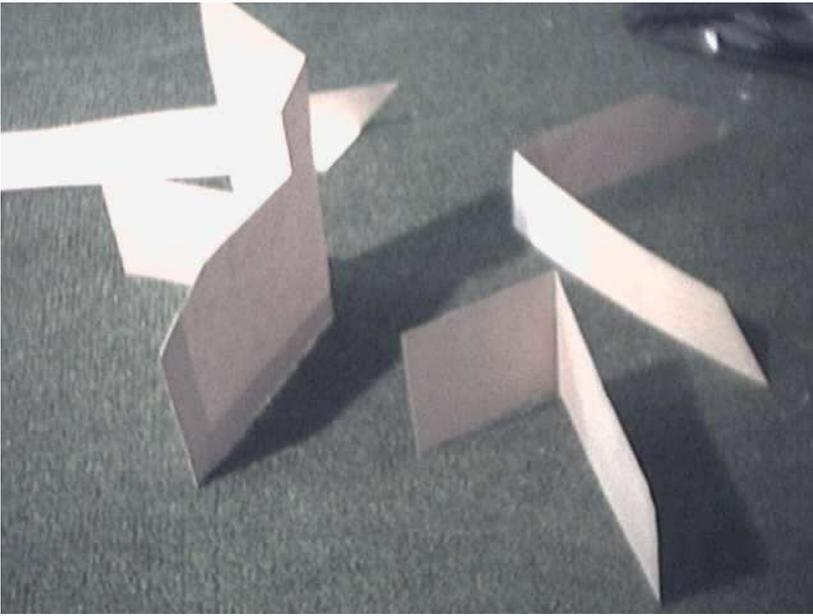


[Fig 2: flat2.jpg](#)

Each polyflap is made by folding one of those polygonal shapes once and leaving the fold 'open' to produce a 3-D shape. Several polyflaps can be placed close together, in some cases with one resting on two or more others, to produce 3-D polyflap configurations. Some examples are in the next lot of pictures.

3-D Polyflap configurations derived from the above polygons:

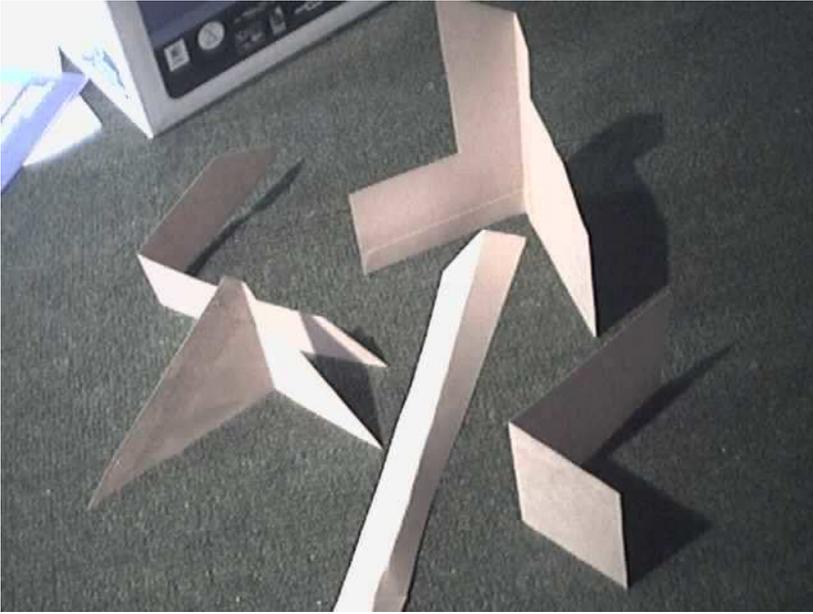
Can you tell which polyflap comes from which polygon?



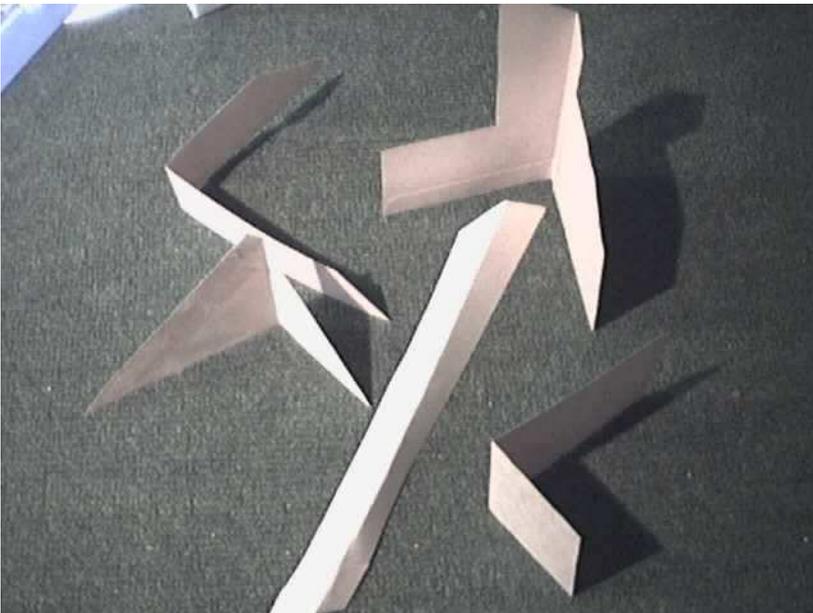
[Fig 3: pol01.jpg](#)



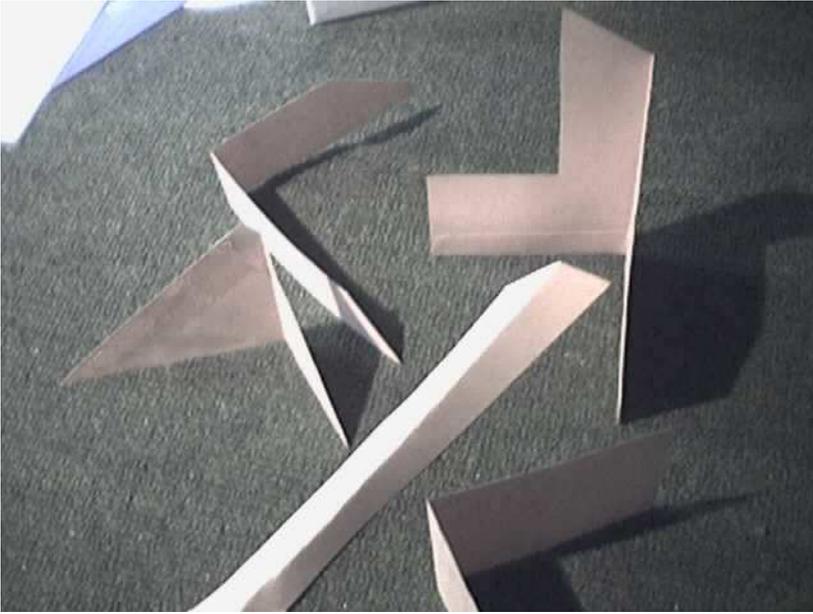
[Fig 4: pol02.jpg](#)



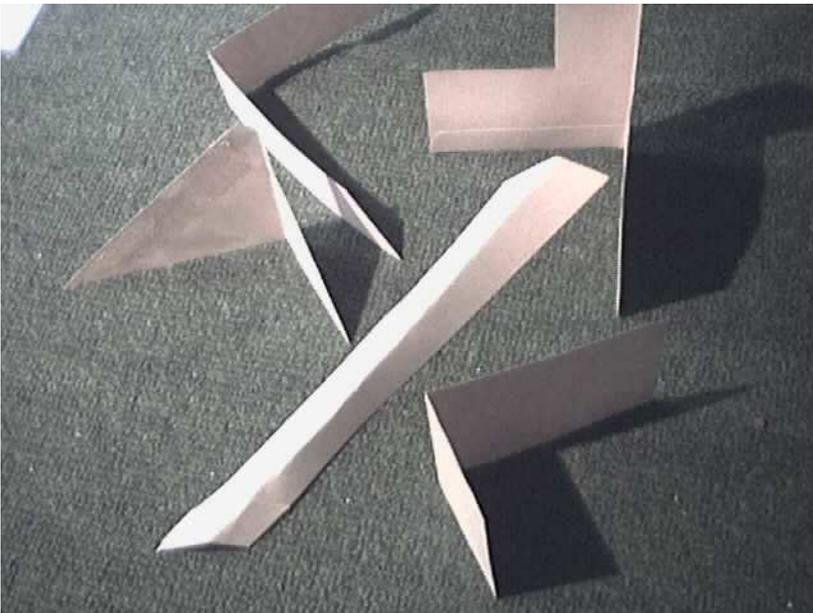
[Fig 5: pol03.jpg](#)



[Fig 6: pol04.jpg](#)



[Fig 7: pol05.jpg](#)



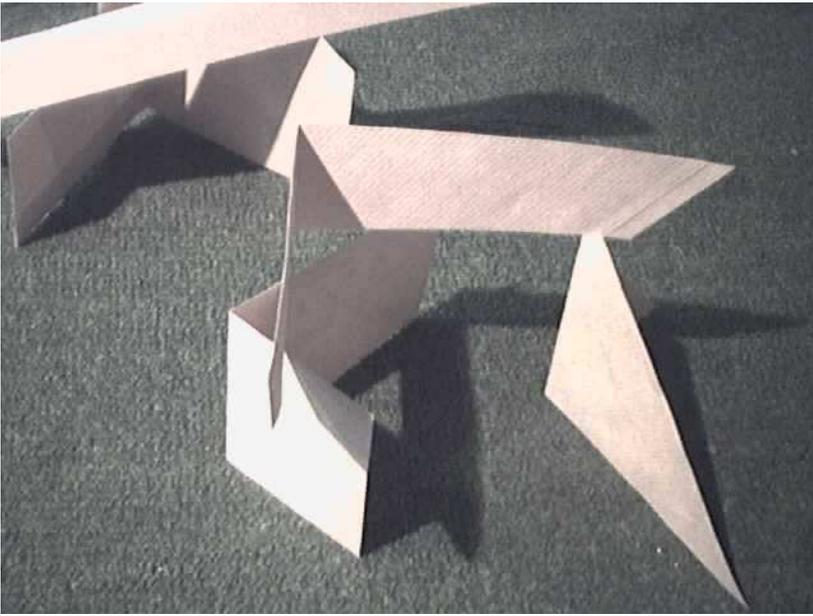
[Fig 8: pol06.jpg](#)



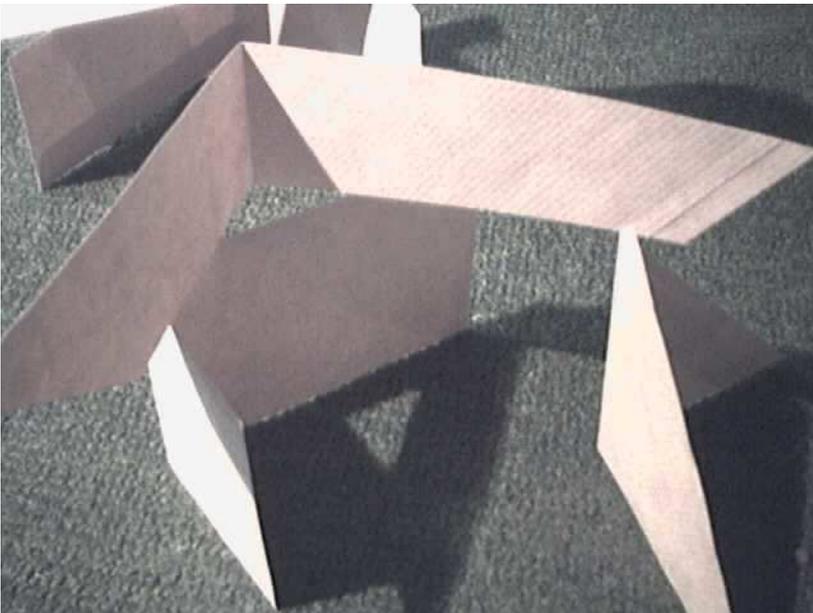
[Fig 9: pol07.jpg](#)



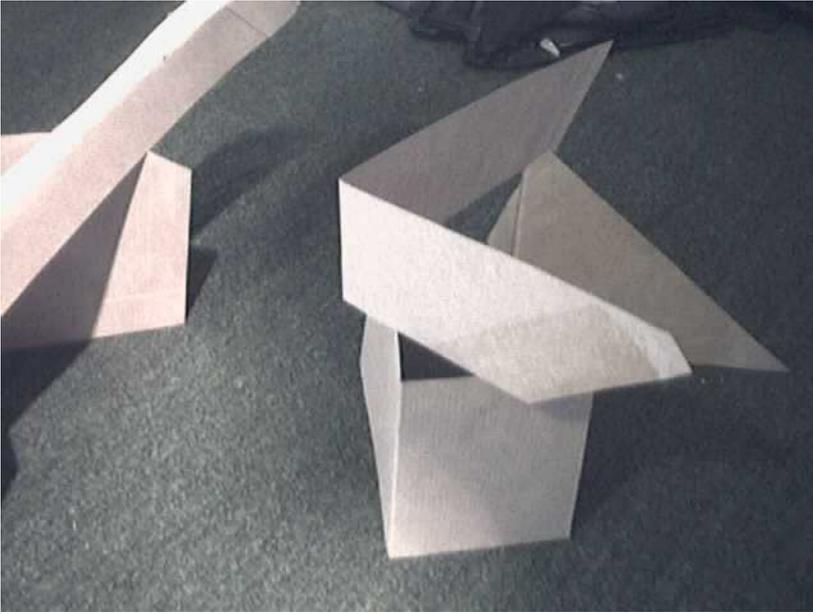
[Fig 10: pol08.jpg](#)



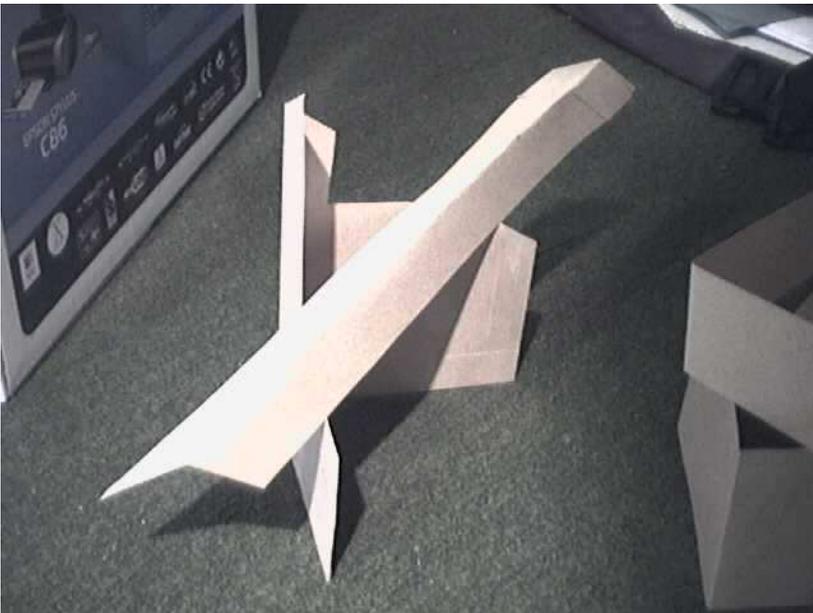
[Fig 11: pol09.jpg](#)



[Fig 12: pol10.jpg](#)



[Fig 13: pol11.jpg](#)



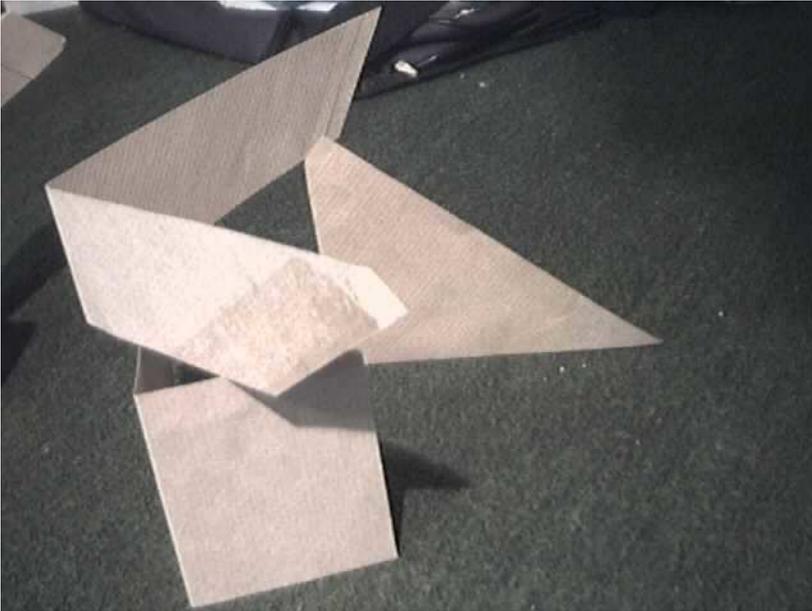
[Fig 14: pol12.jpg](#)



[Fig 15: pol13.jpg](#)



[Fig 16: pol14.jpg](#)



[Fig 17: pol15.jpg](#)



[Fig 18: pol16.jpg](#)

(Some of those are quite hard to see! Shadows can sometimes help. They also sometimes introduce confusing ambiguities.)

The pictures shown here all include some hastily made polyflaps, and are merely illustrative. The polyflaps were made by cutting an old brown paper envelope into polygons, then folding each one once.

For the purposes of the research (on requirements for an altricial robot learning how to perceive, understand and use spatial affordances) I would expect to move to more carefully made polyflaps which are not as flexible as paper. Plastic polyflaps may be used by Jackie Chappell in experiments with parakeets.

No flat polygons are being used because robot hands likely to be available to us would not easily be able to lift a completely flat object. (Likewise a very young child?)

We would start with simpler configurations than those shown, e.g. providing one polyflap at a time at first.

Sub-projects

An interesting sub-project would aim to produce a robot that when faced with a previously constructed configuration of polyflaps can dismantle it one piece at a time without knocking anything over in the process.

Another interesting sub-project would be to get a robot to develop a **mathematical theory of polyflaps**, e.g. specifying conditions under which they are stable, and relationships between initial polygonal shapes and the 3-D structures that can be obtained from them.

Note added 25 Sep 2017

An extension of this would include understanding proto-affordances, e.g. ways of reasoning about possible variations in polyflap configurations, and considering questions such as which are stable, which would be useful intermediary stages in constructing new more complex structures (given a particular collection of primitive polyflaps).

These would require development of topological and geometric forms of reasoning, not learning statistical regularities (which seems to have been the only use of polyflaps in working AI systems since this project was first proposed). E.g.

Marek Kopicki, Sebastian Zurek, Rustam Stolkin, Thomas Moerwald, Jeremy Wyatt, 'Learning modular and transferable forward models of the motions of push manipulated objects'

Autonomous Robots, June, 2017, 01, vol 41, issue 5, pp. 1061--1082, Springer, US.

<https://doi.org/10.1007/s10514-016-9571-3>

POLYFLAP COMPETENCE REQUIRES 'GENERATIVE' COMPETENCE

The point I am making is that polyflap competence requires understanding that is not expressible as a fixed collection of stored cases, and also not expressible as a grasp of some vector space, since the space of possible polyflap configurations and actions and processes involving polyflaps is a structured space with regions of varying complexity partly related to nested structures (larger configurations made of smaller ones, larger perceivable processes made of smaller ones, and larger actions made of smaller ones).

So polyflap competence requires structural-generative competence, probably using forms of representation with compositional semantics. Thus we have a domain that seems to have some of the key features often attributed to linguistic competence but not involving linguistic competence.

The suggestion is that this is closely related to the kinds of domains in which altricial animals gain competence without using language, and that the abilities that evolved to support development of such competence were important precursors to evolution of linguistic capabilities.

Note added: 22 Nov 2017

I have just realised that although my interest in these problems stemmed originally from my interest in philosophy of mathematics, and my attempt

WHAT IS THE POINT?

The point is that the polyflaps are a sort of microcosm of the environment that a typical human child, nest-building bird, or animal with manipulative skills used in hunting, foraging, playing, fighting or feeding learns to perceive, act in, think about, and, in the case of humans, eventually talk about.

Like Galileo's and Newton's thought experiments, the polyflap world abstracts away from a huge amount of complex and messy detail of the real world (including [some of the detail](#)) involved in the ability to produce fast fluent behaviours.

We conjecture that what remains includes a substantial challenge that can lead to a major advance in our understanding of a variety of issues regarding perception of affordances, learning, reasoning, forms of representation, and architectures that build themselves while acting in the world, as seems to occur in the so-called 'altricial' species (unlike 'precocial' species whose individuals start with a much richer collection of genetically determined capabilities whose subsequent development seems to be far more circumscribed).

It is also conjectured that despite the complexity of the polyflap domain there is an interesting subset of it that can be fruitfully explored in the near future in research on cognitive robotics -- both in CoSy and in other related projects by our collaborators.

Moreover, this is only one of a collection of domains that can be investigated in similar fashion within a larger interdisciplinary collaborative research programme including cognitive robotics, linguistics, developmental psychology, neuroscience, biology, and philosophy (among others). In parallel with the polyflap domain we are also investigating simpler domains, e.g. objects being pushed around on 2-D surfaces.

Our physical world allows a huge variety of more or less distinct domains to be learnt about by children, animals, and robots, and the polyflap domain is merely one of those selected because it seems to suit some of the constraints of what can be done with a rather limited one armed robot with a gripper. We also expect to investigate what happens when birds with manipulative capabilities meet polyflaps. Perhaps parallel research could introduce children to polyflaps.

The polyflap domain is just one among many other domains associated with classes of toys and types of materials, e.g. blocks, sand, stones, water, mud, string, paper, pencils, crayons, paint, plasticine, lego, meccano, tinker-toys, dolls of many kinds and their accessories, electronic-circuit building toys, drawing and painting kits, and of course other animals and/or other humans.

We hope that collaborators in other projects will do parallel investigations, so that we can share our tools and our discoveries in a process of joint exploration.

Of course, one of the important questions is how learning about different domains supports creative new combinations, e.g. combining string with other things, in actions like wrapping up, pulling, swinging, tying down, or assembling structures from things that don't intrinsically stick together, e.g. making things out of sticks.

One conjecture underlying this work is that the notion of learnt *sensorimotor contingencies* will not prove rich enough to capture what a child or robot has to learn in dealing with a domain like this. The issue is discussed further in <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0603> (COSY-DP-0603: Sensorimotor vs objective contingencies).

RELATION TO CURRENT AI VISION RESEARCH

One of our presumptions is that a great deal of ongoing research on recognition of objects and production of verbal descriptions of objects and their relationships can to some extent be done using techniques that by-pass most of the characteristic features of human and animal vision involved in both visual and tactile perception and understanding of 3-D shape and spatial structure and the use of that understanding in manipulating things. This is especially true of research that goes from 2-D static or changing visual data to labelling of objects and their relationships, or tracking them as they move around, without producing perception of the 3-D structure of the surfaces of the objects perceived nor any understanding of the affordances for action inherent in the 3-D structure.

So we have deliberately chosen tasks where the recognition of 'whole' objects is not an intermediate goal (though it may be a side-effect), and where most of the visual perception that is produced does not depend on recognition of objects, though it may depend on learning to recognise various 3-D structure fragments and their relationships (described below) and also fragments of motion of 3-D structure fragments.

So our aim is to find ways of perceiving and understanding what is seen that does not presuppose recognition of previously encountered objects nor any linguistic competence.

If we can explain the nature of that spatial understanding we may hope to shed light on a number of central features of human intelligence:

- (a) generative learning capability which involves the *unending* ability to extend the ontology used for thinking about, perceiving and acting in the world
 - (b) the fact that thinkers of many kinds have attested to the power of human-spatial forms of representation and reasoning, and the use of spatio-temporal metaphors in thinking about and dealing with problems in many domains that are not spatial, including such things as family relationships, search strategies in abstract problem solving, and many projects of mathematics.
-

PREREQUISITES FOR ACQUIRING POLYFLAPPERY COMPETENCE

Part of the research involves trying to identify a starting point for a learning robot that includes

a: a set of primitive manipulative competences

b: a set of primitive perceptual competences, for perceiving simple spatial structures, relations, and processes (e.g. two parts of an edge moving closer, two edge-fragments changing their relative orientation, a part of an object sliding along a part of an edge, a fragment of an object coming into contact with a part of a surface of another, or sliding along that surface)

c: mechanism and forms of representation for encoding those information fragments

d: a set of exploration-generating mechanisms that can produce actions that in turn can be perceived

e: a set of meta-level mechanisms determining classes of phenomena that are 'interesting' and should be labelled, the information stored, and later produced in varying ways

f: a set of meta-level 'syntactic' mechanisms for combining previously learnt perceptual or action fragments, or ontological components into larger, re-usable, structures, possibly parametrised in a way that copes with continuous variation. (This can lead to exponential growth in speed of problem solving, but in limited application areas: as seems to happen in humans.)

g: an architecture in which all these resources can be combined so as to grow the competences and extend the architecture in a manner that, at least to some minimal degree of approximation reflects hitherto unexplained aspects of learning and development in altricial animals.

PROGRESSIVE SCENARIOS

The research will be broken into a sequence of increasingly complex stages. Initially the robot will perceive and manipulate only relatively simple polyflaps, learning about various kinds of fragments of both static scenes and scenes where one or more objects move.

-- Static fragments

Static fragments include edges of objects, corners where two edges meet, corners where two edges and a concave or convex fold meet, flat surface regions, portions of shadows, places where an edge or a corner rests on the table or on part of another polyflap, places where two edges or an edge and a fold, or two folds cross each other.

Exactly which 3-D configuration fragments are learnt will be determined by the robot, as a result of applying very general criteria of 'interestingness' (still to be investigated) to processes that occur within its visual system when the robot performs various actions (e.g. moving its head or cameras), or when it passively observes things change as a result of its actions or the actions of others.

E.g. some 'interesting' fragments would be invariants during motions of various cases. Other interesting fragments might be fragments that can be created or destroyed by motion.

-- Larger fragments and combinations of fragments

The robot would also need to learn how different sorts of 3-D fragments can be combined to form larger structures, which themselves may be components of even larger structures. This process will aid perception of complex scenes as a result of learning a sort of 'grammar' for 3-D configurations in the polyflap domain, along with newly developed extensions to the perceptual architecture to enable parsing of increasingly complex polyflap configurations. Exactly what primitive kinds of competence will support such learning, and what architectural developments need to occur during the process is part of the research.

-- Seeing fragments of processes

In parallel with the process of learning to see 3-D fragments and structures made of them, the robot will also learn to see fragments of processes involving those fragments.

processes and causation

affordances

mapping causal powers (enabling, preventing, constraining, etc.) back to the static fragments

after the development of some basic competence in seeing and acting on polyflaps start introduction language learning processes.

(It may be that the mechanisms will be there from the start but unable to do anything. Or it may be that in humans the physical development of language learning mechanisms is 'deliberately' postponed until after infants have developed sufficient non verbal competence to provide a rich semantics to drive the linguistic learning: it seems that an important feature of altricial species is *staggered* development of cognitive mechanisms controlled by genetic mechanisms, of which sexual cognitive development is one of the most obvious examples.)

SURPRISE CAN DRIVE LEARNING

(Compare Sussman: Bugs drive learning. Also John Holt on learning)

Clearly there are various kinds of surprise (Matthias Scheutz and I have an incomplete partial conceptual analysis of surprise in a larger incomplete paper. We should make this available.)

- Passive surprise: something novel

There's noticing something that has not been seen before (lump of yogurt on the carpet in the video below[4]) which can trigger various new actions, e.g. pick it up, put it in spoon, put it in mouth. But then having done that, a new goal can be triggered: put another lump of yogurt on carpet.

- Active surprise 1: failed prediction
- surprise 2: failed action
- Other surprises: unexpected (side-)effects of actions,
 - interesting ones
 - good ones
 - bad ones

How do surprises of various kinds help to drive increased self-understanding and other-understanding. (Note that some of the representational requirements for mechanisms relating to self-description and other-description will be similar, for both involve meta-semantic competence. This is important in relation both to evolution and to individual development -- a recurring theme in the Cognition and Affect project.[5])

THINGS TO DO

- sketch some other domains
- identify initial architecture and patterns of growth of architecture
- identify initial ontology and patterns of growth of ontology (both in the robot and also in *us* as we study these phenomena!)
- identify initial forms of representation and patterns of growth of forms of representation
- identify initial chunking mechanisms for isolating fragments that should be stored, 'labelled' and made re-usable in various ways (what ways?)
- identify initial forms of 'syntactic' or 'combinatorial' competence for recombining chunks in perception and action. (For a very simple example in a very simple world with a moving finger and a counter, originally invented by Oliver Selfridge see <http://www.cs.bham.ac.uk/research/poplog/teach/finger>)
- identify forms of motivation and affective responses that will help to drive the robot's exploration and learning.
- some of this may include mathematical properties of patterns of sensory input or patterns of motor output that trigger some 'store this as a chunk' mechanism. (Compare Gardenfors, among others.)
- once the robot has started learning about chunks it can start learning generalisations associating those chunks.
- Having learnt to construct representations of what is perceived, or done, it could also (with an architectural extension -- towards deliberative mechanisms) start 'playfully' constructing representations of possible percepts, possible actions, not currently perceived or done.
- With an appropriate architectural change some of those representations could begin to function as goals. (How would they be chosen? more on motivation.)
- Using the deliberative capability (representing what does not exist), stored associations will trigger predictions of various sorts (forward chaining) but can also trigger action selection processes, then later on planning processes (backward chaining), and still later explanation processes.
- Once predictions are made, predictions can be checked (another architectural development), and the difference between successful and unsuccessful predictions learnt. (Or does that distinction have to be innate -- like Kant's presumption of an 'objective' external reality underlying what is perceived??).
- and then various ways of dealing with failed predictions can be explored and learnt, including changing associations, debugging prediction and planning strategies, etc. (Sussman's HACKER, Push Singh's PhD thesis and other refs.)

More on affect: perceived internal and external states and processes, and various kinds of actions need to be evaluated. It's clear from watching very young children that a lot of the motivation and evaluation has nothing to do with satisfying biological physical needs (getting food, reducing various kinds of physical discomfort, generating 'physical pleasures') or even social needs. There are deeper cognitive needs.

(See also this paper on "Architecture-based motivation" contrasted with "Reward-based motivation":

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html>

BROADER RESEARCH ISSUES

There are many potential links to neuroscience, psychology, biology. (Much of this arises out of discussions with Bioscientist Jackie Chappell who has been working on animal cognition, especially crows and parrots.)

One of the interesting questions to be addressed in the long run on the basis of research like this could be: what are the differences in learning capabilities of human children born with or without arms, e.g. those affected by the thalidomide tragedy in the 1960s. E.g. See

<http://cerhr.niehs.nih.gov/genpub/topics/thalidomide2-ccae.html>

<http://www.fda.gov/cder/news/thalidomide.htm>

Other investigations of a related type could investigate learning capabilities of blind children, using robots with only arms, and no vision.

If we can identify different sorts of 'initial' architectures etc. we may be able to come up with a new way of classifying types of animal species, going far beyond the altricial/precocial distinction.

This can also lead to new kinds of research into evolutionary trajectories and the kinds of environments that trigger, facilitate or inhibit them.

Various Polyflap-based domains could provide useful environments in which to test the ideas about pre-configured and meta-configured competences in [Ref \[2.b\]](#), especially the ideas about cascaded interactions between genome and environment summarised in this diagram:

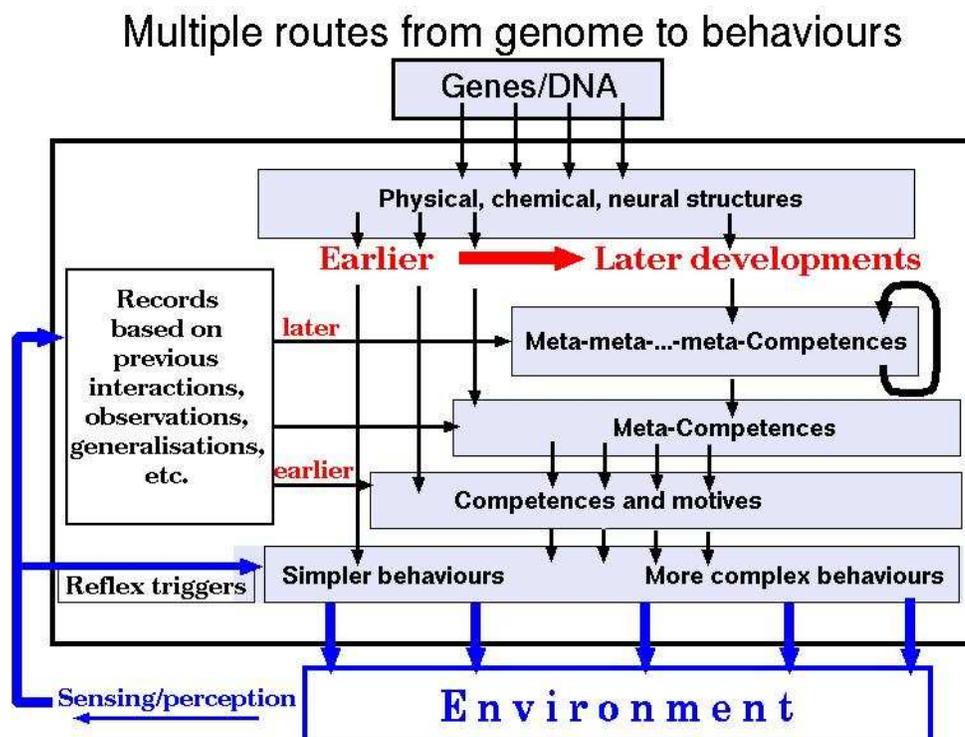


Diagram based on ideas in [Ref 2b](#).

Three multidisciplinary research proposals to investigate these ideas, by relating experiments on animal cognition and work on robot cognition were turned down.

The possibilities for expanding the robot work in different kinds of directions and also the interdisciplinary collaborations are endless.

Note added 4 May 2011:

I have just realised there is a strong connection between the motivation for this proposed domain and the ideas about Representational Redescription in Annette Karmiloff-Smith's 1992 book

Beyond Modularity: A Developmental Perspective on Cognitive Science

See my draft discussion of that book in

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html>

Note added 26 Sep 2013:

The Polyflap domain could be used to test the ideas in this presentation on vision and the need to extend Gibson's ideas about affordances.

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#gibson>

What's vision for, and how does it work?

From Marr (and earlier) to Gibson and Beyond

(With some potted, rearranged, history)

Since this paper was written, individual polyflaps have been used in various robot projects (e.g. search for "pushing polyflap"), but so far only in much simpler configurations than proposed here as a challenge for research in robotics (including vision).

INFLUENCES

The impact of these ideas, since they were first presented (2005-6) can be gleaned by searching for papers that use the word "polyflap" or "polyflaps" along with one or more of "robot", "robotics", "vision", "learning", "ontology", and similar words. So far (10 Jan 2016) it seems that only researchers in Europe have used the idea.

Note added 10 Jan 2016: Compare the Author's submission to the Royal Society on machine learning in 2016.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/royalsoc-deep-learning.html>

Note added 22 Nov 2017

So far nobody seems to have noticed the deep opportunities for mathematical discovery -- learning about structures, relationships, effects of various changes, including effects of changes of viewpoint on how things appear - necessary connections between structures and processes, including impossibilities:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html>

(Compare the notion of an "aspect graph".)

SEE ALSO

To be extended.

For more on the context see

[1] <http://www.cs.bham.ac.uk/research/projects/cosy/PlayMate-start.html>

Introduction to the PlayMate scenario of the CoSy research project

[2.a] <http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>

The altricial precocial spectrum for robots

Paper co-authored with Jackie Chappell, for IJCAI'05.

The next paper was a sequel, expanding these ideas.

[2.b] <http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717>

Jackie Chappell and Aaron Sloman, (2007)

Natural and artificial meta-configured altricial information-processing systems, invited paper for *International Journal of Unconventional Computing*, 3, 3, pp. 211--239,

[3] <http://www.cs.bham.ac.uk/research/cogaff/challenge.pdf>

Challenge to AI vision researchers and others, regarding perception of shape-based affordances.

[4] Baby and yogurt video

[3.6MB Mpeg](#)

[19MB Mpeg](#)

[5] <http://www.cs.bham.ac.uk/research/cogaff/>

<http://www.cs.bham.ac.uk/research/cogaff/talks/>

Papers and talks in the Cognition and Affect project directory.

[6] The UKCRC 'Architecture of Brain and Mind' research grand challenge.

<http://www.cs.bham.ac.uk/research/cogaff/gc/>

<http://www.ukcrc.org.uk/>

[7] On scenario driven research:

[Overview of methodology](#)

[Draft Scenario template](#)

[8] Provisional work on ontologies for a pre-linguistic agent

[ontology.pdf](#)

[ontology-for-a-manipulator.txt](#)

[evolving-ontologies-for-minds.txt](#) (newsgroup posting 1999)

NOTES AND NEWS

This file is <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps/>

Also available as

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps/polyflaps.pdf>

This web site is still under construction.

Updates:

24 Mar 2008; 4 May 2011; 7 Jun 2011; 26 Sep 2013 (Added PDF version); 22 Nov 2017

- **REORGANISED 6 Jun 2006**

- **NOTE ADDED: 16 Jan 2006:**

My thoughts have moved on a lot since I wrote these notes between May and October 2005. In particular

- o See the discussion of [Orthogonal recombinable competences](#).

The ideas presented there suggest that a project of the sort envisaged should also explore learning about different kinds of physical stuff.

The ideas are expanded in

- o A presentation on [A \(possibly\) new theory of vision as concurrent simulation of processes at different levels of abstraction in partial registration with the optic array](#),

- o A partly overlapping presentation on [Two views of child as scientist: Humean and Kantian](#)

Part of the background to this paper is a claim that AI *vision* research (as opposed to pattern recognition, tracking, and learning research has made very little progress in the last 30 years, for reasons indicated in <http://www.cs.bham.ac.uk/research/cogaff/challenge.pdf>

- 10 Aug 2005

Introduction expanded, terminology clarified (following IJCAI presentation with Jackie Chappell).

A new, related web page concerned with the hypothesis of [Primacy of dynamical systems \(PDS\)](#) provides some of the context for the polyflap discussion.

- Tue Jul 26 11:45:08 BST 2005

(Converted from plain text to html)

- Wed May 25 12:10:18 BST 2005

(Extended descriptions of what the robot will learn, and elaborated on the context)

(New pictures from stiffer, brown paper)

Acknowledgements

Much of this work was inspired by things I learnt many years ago from Marvin Minsky (see his 1962 paper 'Steps towards artificial intelligence' and the online draft book [on his web site](#) 'The Emotion Machine')

More recently there has been collaboration with Push Singh[*]. See his PhD thesis: ([PDF](#), [HTML](#))

[*]Alas Push died in February 2006.

Maintained by [Aaron Sloman](#)