Internalizing Intelligent Activity

Three Themes

Directed - Undirected

Self Extension

Free Structure

Directed - Undirected

An algorithm is directed to a purpose - you had to know the purpose before you wrote it

An Artificial Neural Network is directed to a purpose - the flow direction is fixed

It is possible to have an undirected structure - a structure that is not directed to any particular purpose

Undirected Structure

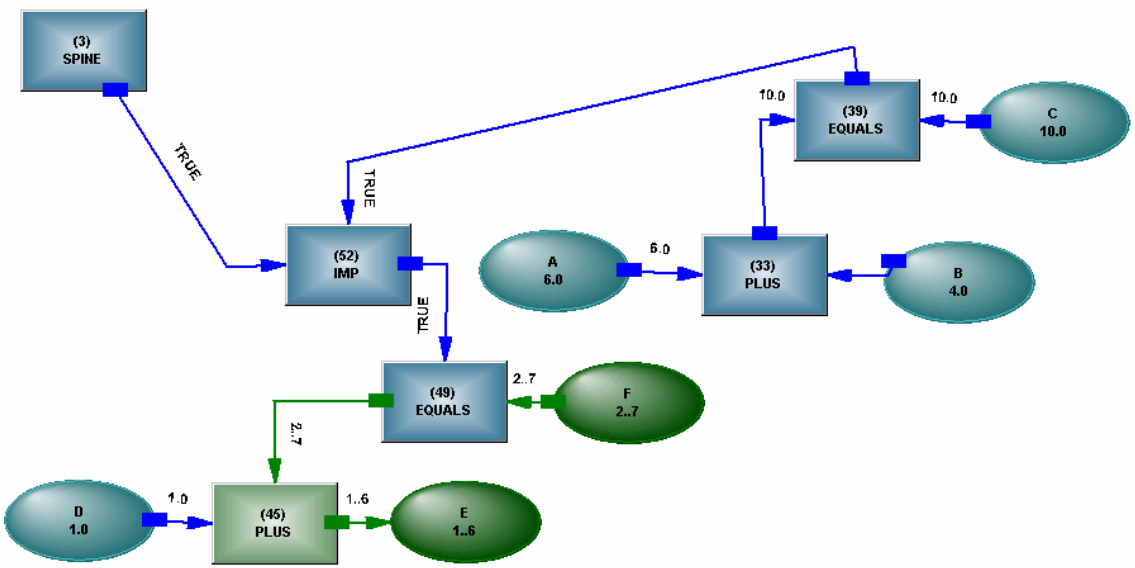

An implication in dynamically directed structure

- It behaves as an IF...THEN...
- It implements *modus tollens*
- It can be used to reason about its own validity

The numerical structures can be used to find any of the variables, used as equations or inequations

- the structure has many uses

Bubbles Rising

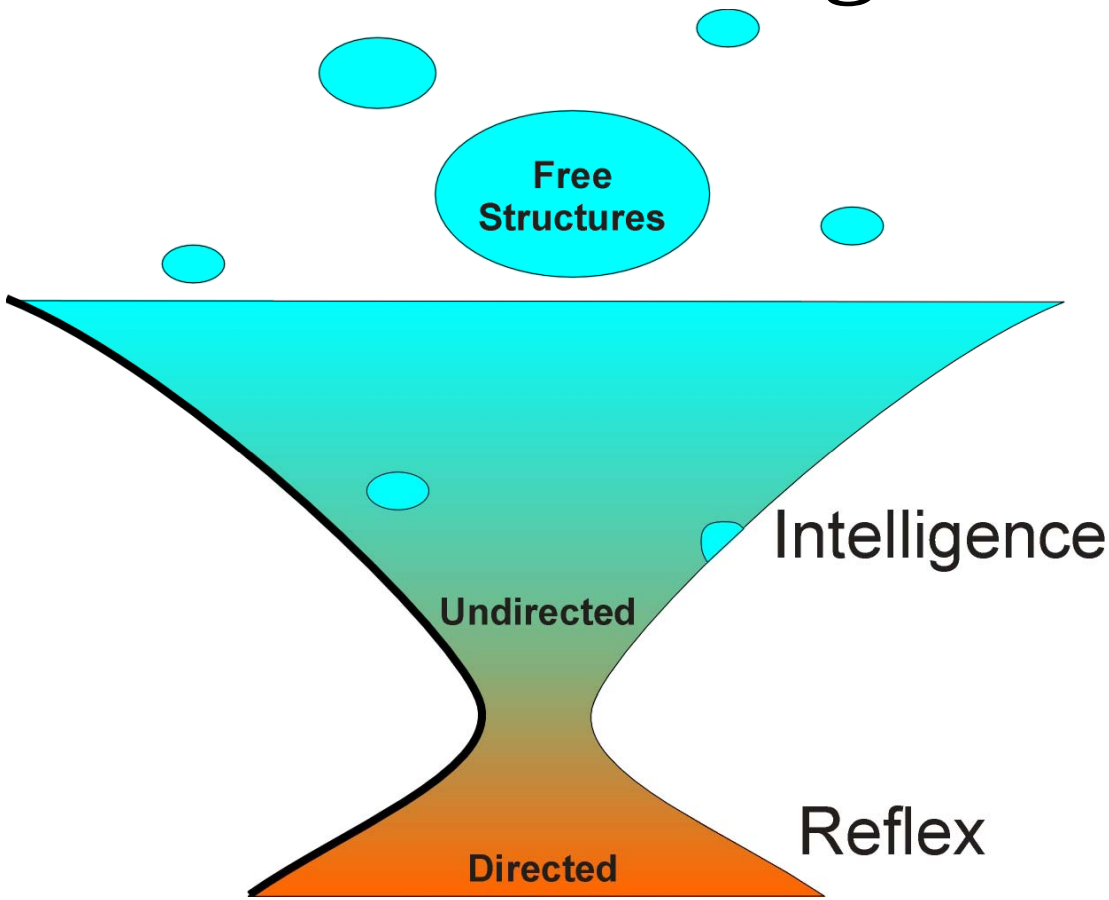

We start with a directed base, layer it so it can be undirected, which reduces the number of elements, but greatly expands the capabilities, and it can even blow bubbles - free bubbles

Explosive Potential

Total neurons	10E11	(100 billion in your head)
Cognitive	10E10	(say 10% of total)
Undirected	10E8	(say it takes 100 neurons to synthesize one undirected one)
Potential use	10E100	(an ordinary sentence)

If you can get to an undirected level, then each element does double work, your cognitive potential hugely increases, and you can generate structure that is free of *the surly bonds*

Self Extension

We start with a structure
It reads something, builds some structure
The new structure immediately becomes active,
becomes part of the original structure,
is used to read more, build more

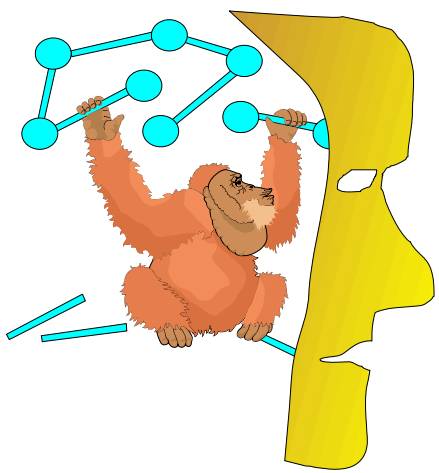

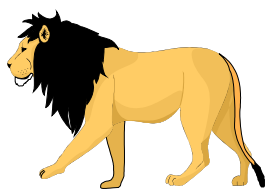

Structure Building

A person reads a book and converts it into internal structure.
What they have read helps them to read further
- the structure builds on itself.

Building Structures

The Jacquard Loom

“a program weaving a structure”

Bright Idea

People looked at the Jacquard loom and thought

“We could get a program to write a program”

Didn’t work

The Concept Was Good

Let’s look at the loom again

“a program weaving a structure”

Why don’t we make it

“a structure weaving a structure”

That might work

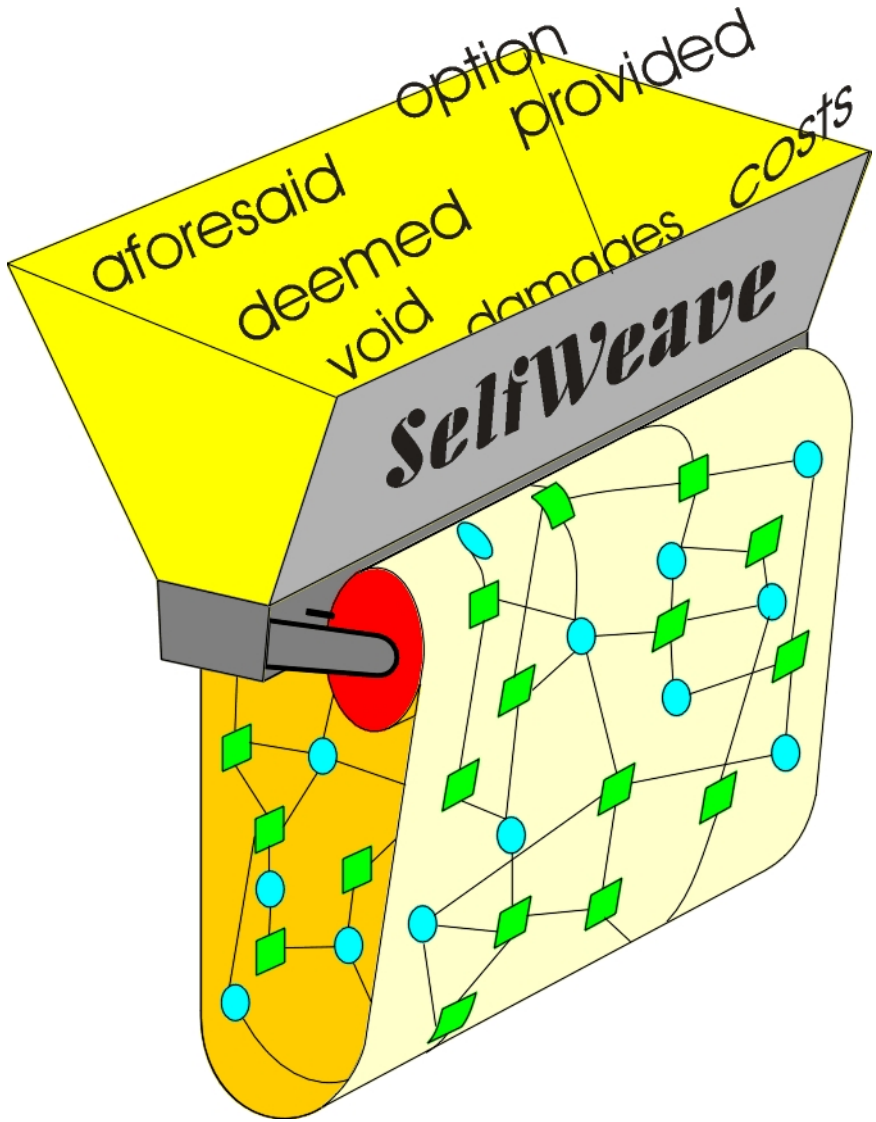

What Stuff?

Self Extension sounds like a useful property - what sort of stuff has that property?

Neurons do, but we would need millions of them, and they are a bit wet and hard to connect manually

We do have lots of computer memory - how about that, at least as something to keep the stuff in

What Granularity?

If we want self-extension, we have to go down to a very low level, a level where we can connect to states flowing in the structure - down to the atomic level of operation - where we can accrete structure to change the behavior of existing structure

We may not be receiving information at this level from outside - we need structures inside the cognitive core capable of operation at the atomic level

All the Same Stuff

A structure reading text is extending itself by building more structure made out of the same stuff and seamlessly connected to the structure doing the building

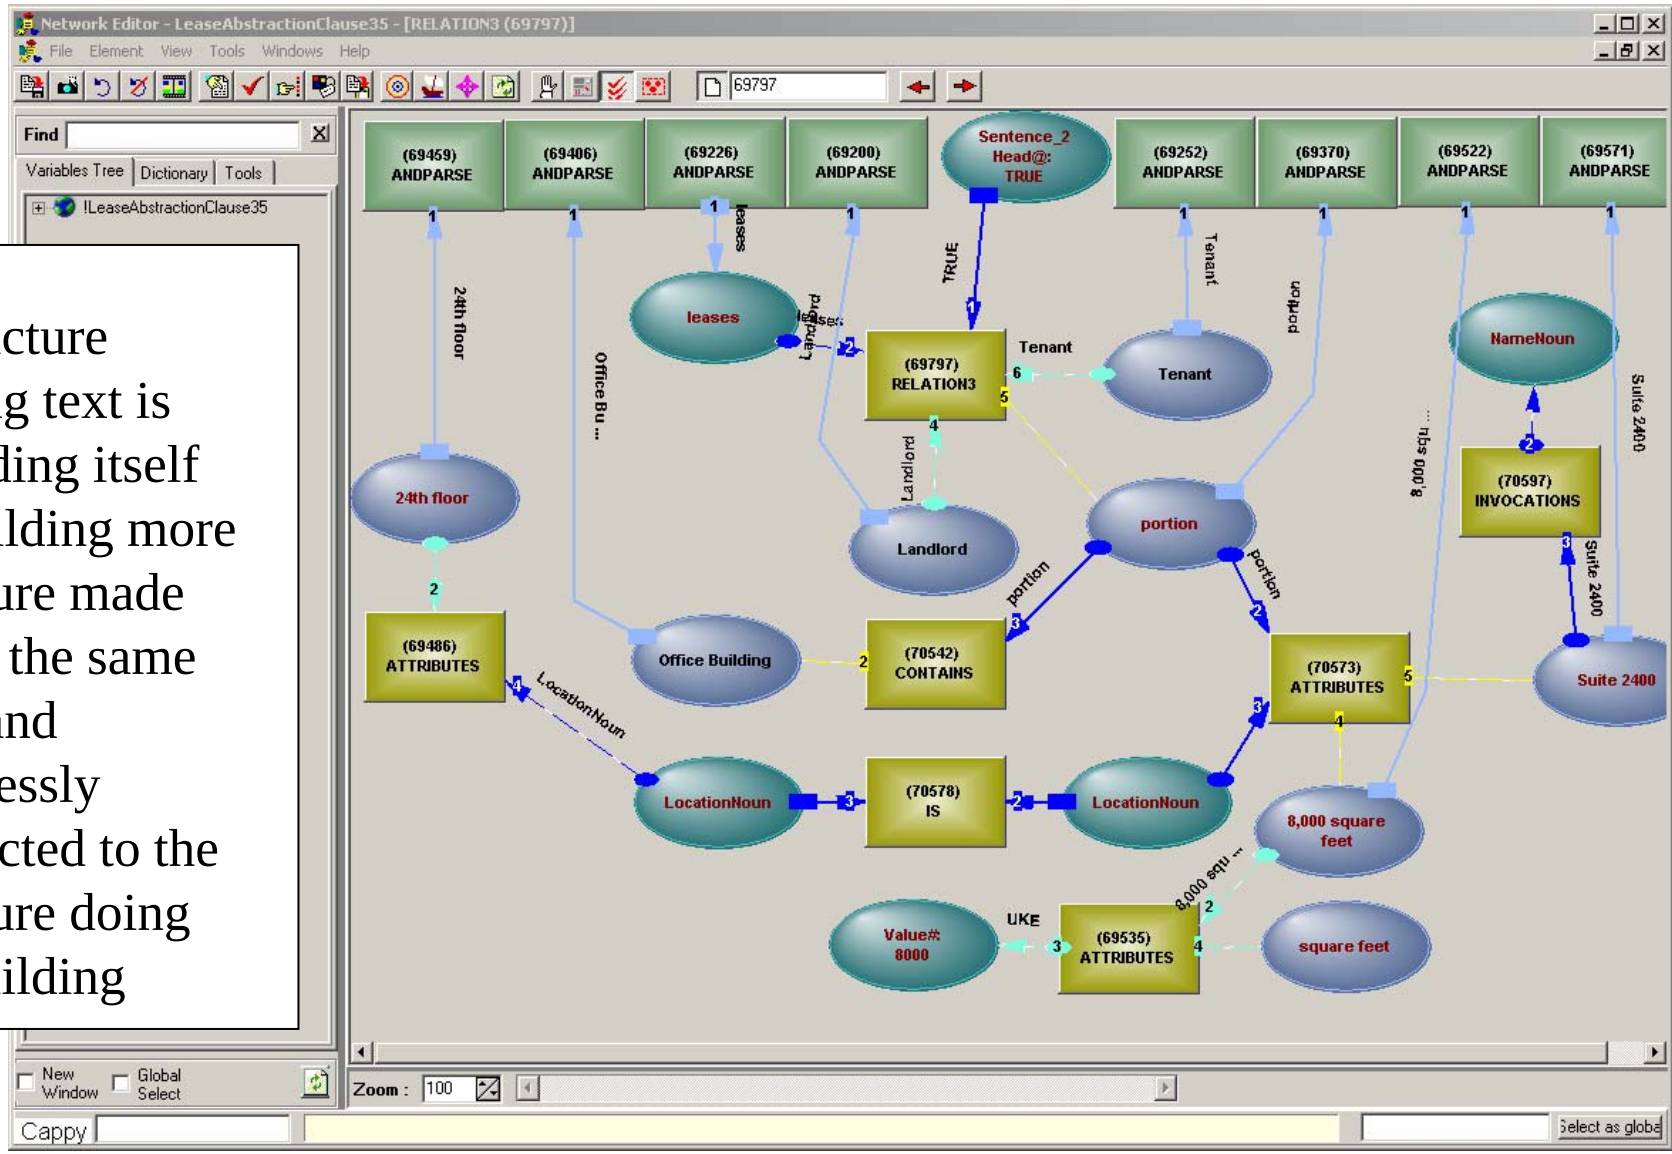

Machine Building

If we are to weave the logical patterns of reality and the universe of thought, we have to expect a structure that is hugely dense and hugely detailed

We will not be able to do it with a few statements in a formal language

We will need a machine and a rich flexible description to drive it - a natural language that meshes logic and relations

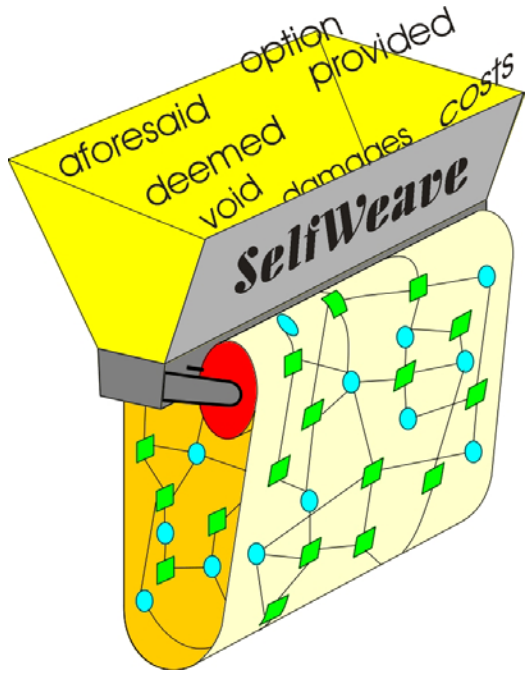

We Can’t Know

When we are reading text, we can’t know how the pieces are being used until the tapestry is finished - we can’t build directed pieces, because we don’t know what they are directed to, and may not know for twenty years, but we do need to build something to help us to understand what we read next - we have no choice but to build undirected structure

Islands in the Sky

If we have a large complex cognitive structure, we won't always be building out from a solid base
Sometimes we will build something in free space, not knowing where it should go until we build it and we can see what it looks like

We are going to need structures that crawl over other structures to repair or join them - structures that are not directly connected and can work alone

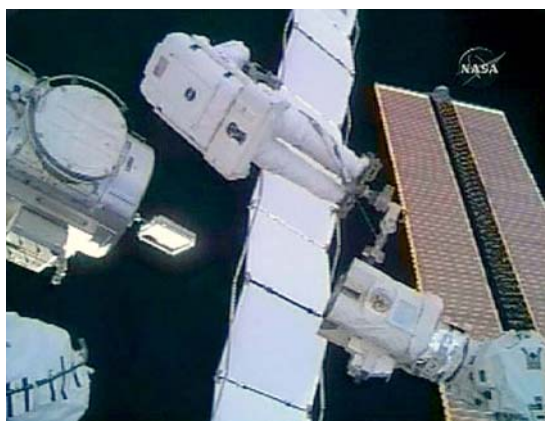

Active Map

An active map is brought to the site of a discontinuity and uses constraint reasoning to check whether it matches the things to be connected. If so, it connects itself, it adds, rearranges or destroys some structure, then disconnects itself and puts itself back on the shelf.

It breaks the connection paradigm, even though there is an implicit connection to some of the things it will connect.

It is like the person attending the loom, fixing anything the static structure of the loom can’t fix, like a broken thread in warp or weft

Simple Example

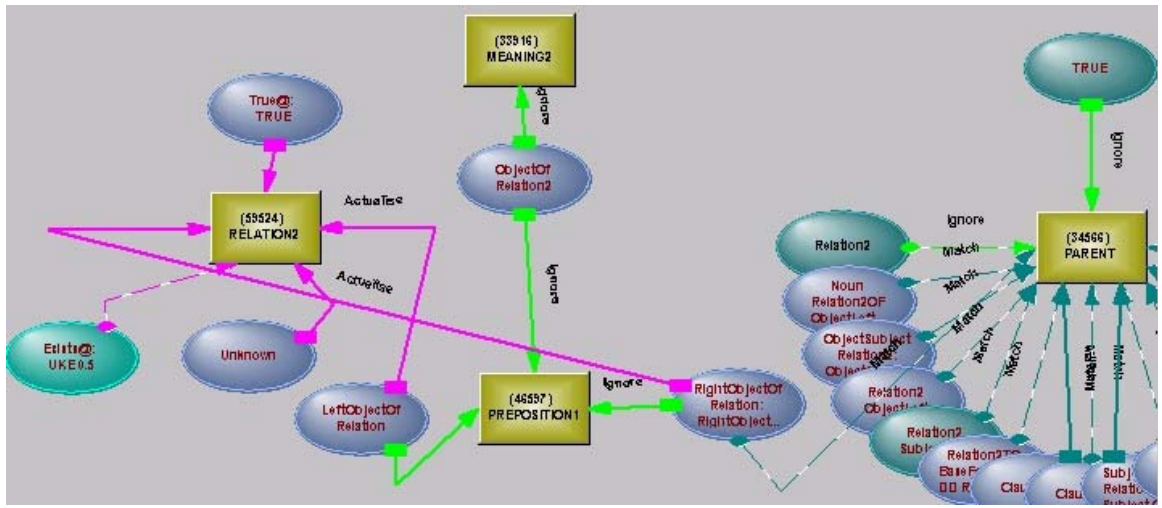

Map takes Object Relation, validates itself, builds the relation

It backtracks on failure, but having built the relation, it can’t backtrack, so cleans itself for next time

Internalizing Activity

If we want undirected structure, then it can’t be directed by something outside it - whatever is outside would get lost

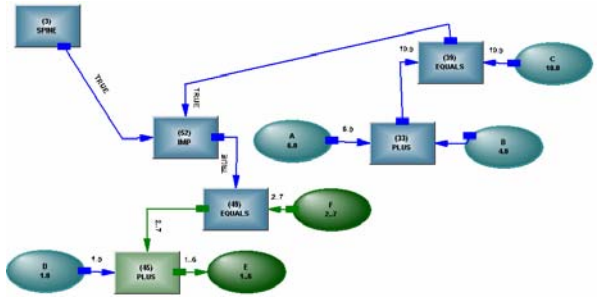

If we want self-extension, then everything has to be in the structure - the semantics, the inferences, everything - connection gives us activity

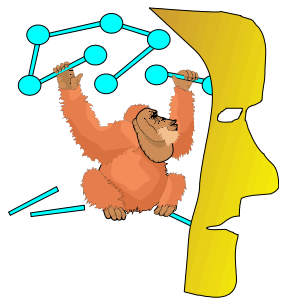

If we want free structures, then everything has to be in the freed structure - it can’t call home to find out what to do

A Single Theme

We started with three themes, but there is a unifying theme

-

an undirected, self-extending structure capable of deploying free structures to aid in its self-extension

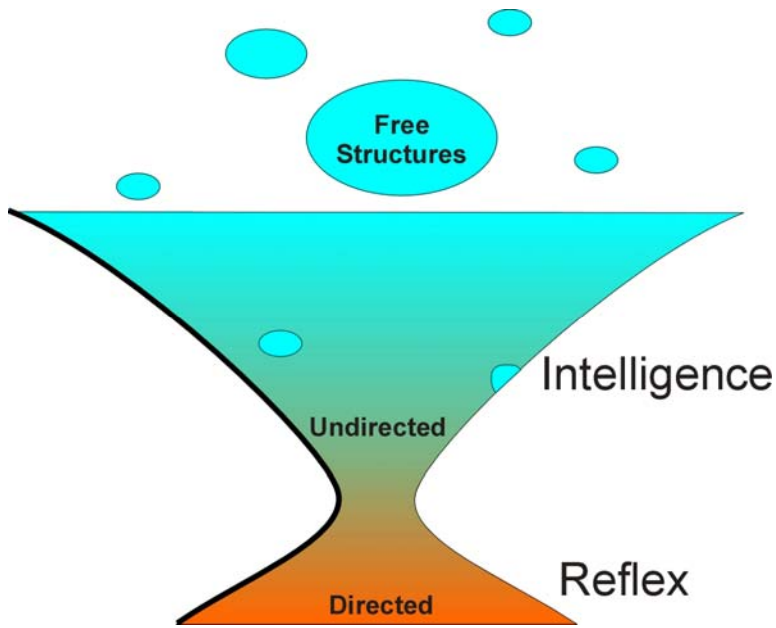

Intelligent Computing

We reach an undirected level, we blow away the limitations of the directed base

If we also make the operations within the structure atomic, we have overcome many of the problems of intelligent computing