COSC451: Artificial Intelligence

Lecture 18: Object perception and the syntax of noun phrases

Alistair Knott

Dept. of Computer Science, University of Otago

A general question

Take some concrete sentences.

There is a dog in the garden
Many dogs were brown

There are sm dogs in the garden
The dog was a dachshund

How do these linguistic expressions interface with perceptual processes?

(Clearly there *is* an interface, because the sentences report information which can be obtained perceptually.)

Two hypotheses

Language and perceptual processing are modules

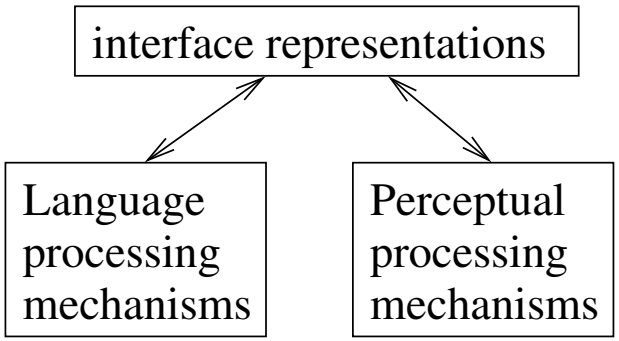

Interface representations abstract away from details of SM processing

Language and perceptual processing share mechanisms

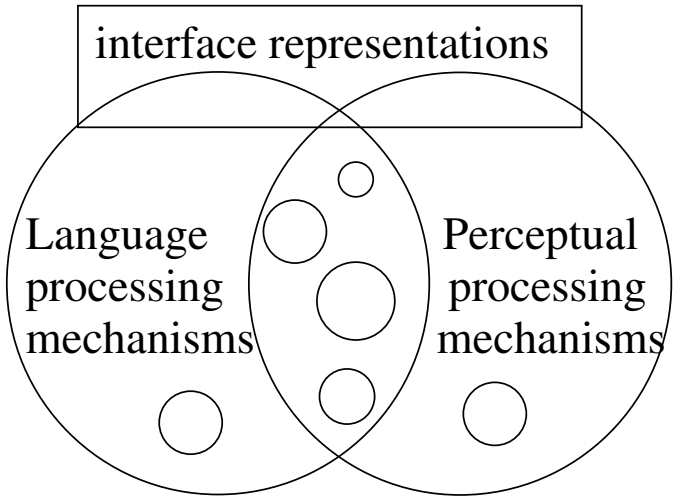

Interface representations retain details of SM processing

The shared mechanisms hypothesis

I'm investigating the second hypothesis.

A strong version of this hypothesis:

> *The syntactic representation of a concrete sentence conveying proposition P can be understood as a* description *of the perceptual processes involved in acquiring P.*

A method for investigating the hypothesis:

- Choose a concrete sentence conveying *P*.
- Give a model the perceptual processes involved in acquiring *P*.
- Give a model of the syntactic structure of the sentence.
- Look for formal similarities between the two models.

Noun phrases and their perceptual correlates

In this lecture, I'll be thinking about the perceptual correlates of noun phrases.

There is a dog in the garden
Many dogs were brown

There are sm dogs in the garden
The dog was a dachshund

The task in this case:

- Model the syntactic/semantic contribution of NPs in the chosen sentence.
- Are there any elements in perceptual model which have a similar contribution?

What can NPs do?

A simple proposal: 'NPs model the processes involved in perceiving objects'.

This certainly isn't the whole story.

1. NPs don't always denote objects.

   - NPs can refer to *groups* of objects. (*The dogs barked*)
   - NPs can denote *substances*. (*John ate dog for supper*)
   - NPs can denote *properties*. (*Fido is a dog*)

What can NPs do?

2. It's impossible to study NPs without considering their relationship to the clauses they're embedded in.

- NPs can introduce *quantifiers*. (*Every dog barked*)
- NPs can refer to *bound variables*. (*Every dog loves its master*)
- NPs can be *expletive*. (*There is a cat outside*)
- NPs have *agreement features*. (*The dogs are barking*)

So minimally, we need to study how processes of object perception are incorporated into broader processes of 'establishing propositions'.

Talk outline

Outline of today's lecture

Choosing a syntactic framework

How can syntactic structures be descriptions of perceptual mechanisms?

- Everyone has the same perceptual mechanisms.
- But NPs look very different in different languages.

To maintain the hypothesis, we must adopt a syntactic theory which distinguishes between 'surface' syntactic representations (which differ between languages) and 'underlying' syntactic representations (which don't).

- I'll use a Minimalist framework (Chomsky, 1995).
- (Some of the arguments carry over to other frameworks.)

X-bar theory

Minimalism assumes that syntactic structures are built from uniform building blocks called maximal projections, or XPs.

Each XP looks like this:

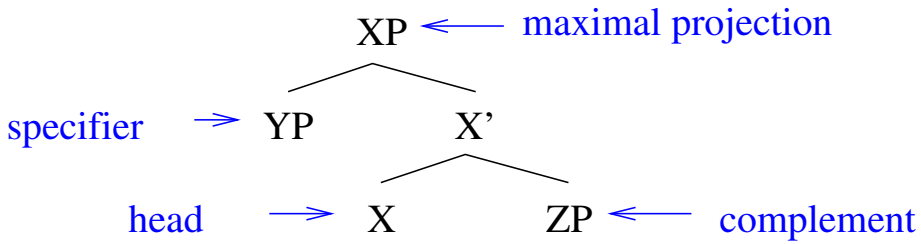

- Each word sits at the head of its own XP.
- In Minimalism, there are also XPs with functional heads, which contribute morphology or semantic elements.

The DP hypothesis

The simplest syntactic model of NPs is as follows:

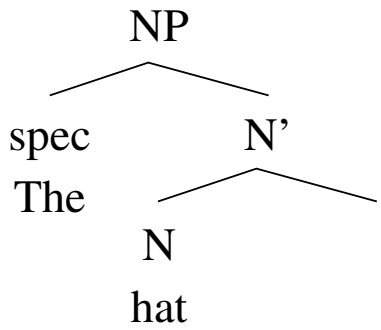

But a more complex structure is now preferred:

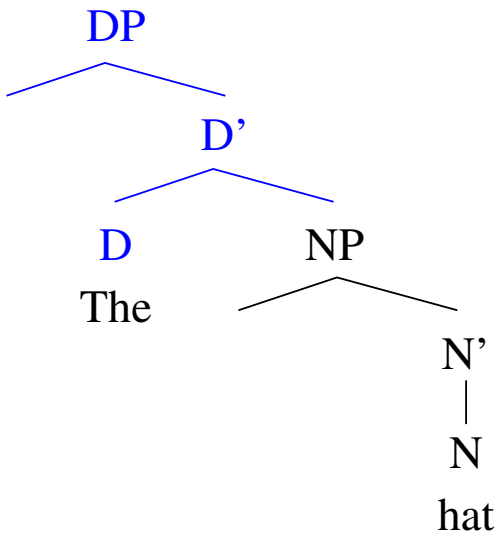

The DP hypothesis

Abney (1987): 'NPs' are headed by determiners, not nouns.

1. 'NPs' have a lot in common with clauses. For instance, they can have correlates of subjects and objects:

   Peter has a hat *Peter's hat*

2. In some languages, the 'subjects' in NPs agree with their 'objects':

   My hat-1sg *His hat-3sg* *Our hat-1pl*

The DP hypothesis

Abney (1987): ‘NPs’ are headed by determiners, not nouns.

1. ‘NPs’ have a lot in common with clauses. For instance, they can have correlates of subjects and objects:

   Peter has a hat *Peter’s hat*

2. In some languages, the ‘subjects’ in NPs agree with their ‘objects’:

   My hat-1sg *His hat-3sg* *Our hat-1pl*

3. In clauses, agreement morphology heads its own XP Infl (or I):

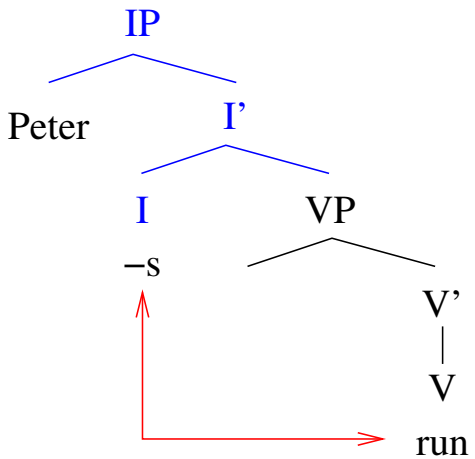

The DP hypothesis

Abney (1987): ‘NPs’ are headed by determiners, not nouns.

1. ‘NPs’ have a lot in common with clauses. For instance, they can have correlates of subjects and objects:

   Peter has a hat *Peter’s hat*

2. In some languages, the ‘subjects’ in NPs agree with their ‘objects’:

   My hat-1sg *His hat-3sg* *Our hat-1pl*

4. In a ‘NP’, we can envisage a similar mechanism.

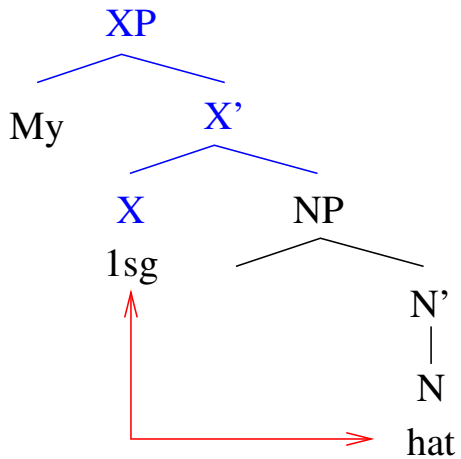

The DP hypothesis

Abney (1987): ‘NPs’ are headed by determiners, not nouns.

1. ‘NPs’ have a lot in common with clauses. For instance, they can have correlates of subjects and objects:

   Peter has a hat *Peter’s hat*

2. In some languages, the ‘subjects’ in NPs agree with their ‘objects’:

   My hat-1sg *His hat-3sg* *Our hat-1pl*

5. The functional projection can be headed by D.

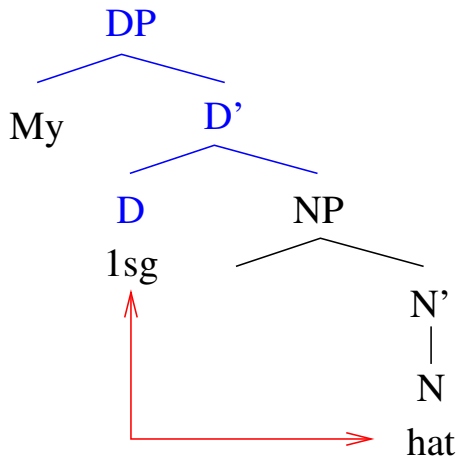

The DP hypothesis

Abney (1987): 'NPs' are headed by determiners, not nouns.

1. 'NPs' have a lot in common with clauses. For instance, they can have correlates of subjects and objects:

   Peter has a hat *Peter's hat*

2. In some languages, the 'subjects' in NPs agree with their 'objects':

   My hat-1sg *His hat-3sg* *Our hat-1pl*

5. The functional projection can be headed by D.

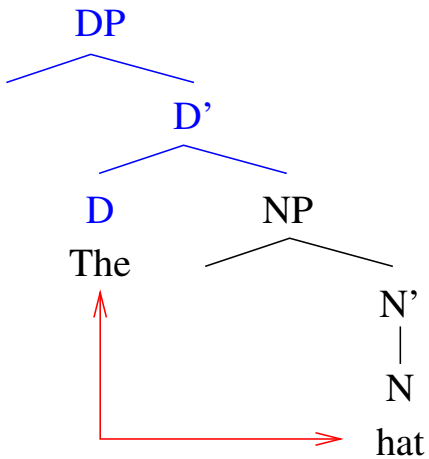

The DP hypothesis

Abney (1987): ‘NPs’ are headed by determiners, not nouns.

1. ‘NPs’ have a lot in common with clauses. For instance, they can have correlates of subjects and objects:

   Peter has a hat *Peter’s hat*

2. In some languages, the ‘subjects’ in NPs agree with their ‘objects’:

   My hat-1sg *His hat-3sg* *Our hat-1pl*

5. The functional projection can be headed by D.

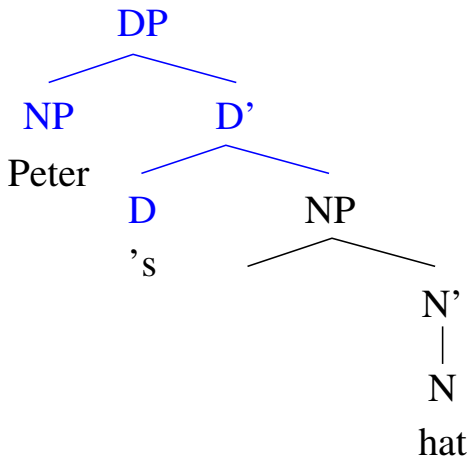

The DP hypothesis

Semantic argument for the DP hypothesis:

1. Quantified sentences have a standard form:

   many *dogs* *bark*

The DP hypothesis

Semantic argument for the DP hypothesis:

1. Quantified sentences have a standard form:

$$many(x,\ dog(x),\ bark(x))$$

The DP hypothesis

Semantic argument for the DP hypothesis:

1. Quantified sentences have a standard form:

$$most\ (x,\ teacher(x), strict(x))$$

The DP hypothesis

Semantic argument for the DP hypothesis:

1. Quantified sentences have a standard form:

$$Q \quad (x, \; P1(x), \quad P2(x) \;)$$

The DP hypothesis

Semantic argument for the DP hypothesis:

1. Quantified sentences have a standard form:

$$Q \quad (x, \; P1(x), \quad P2(x) \;\;)$$

2. The 'semantic template' of a quantified sentence is contributed by the determiner of one of its 'NP's.

$$[[most]] = \lambda P1 \lambda P2[most(x, P1(x), P2(x))]$$

The DP hypothesis

Semantic argument for the DP hypothesis:

1. Quantified sentences have a standard form:

$$Q \quad (x, \; P1(x), \quad P2(x) \;)$$

2. The 'semantic template' of a quantified sentence is contributed by the determiner of one of its 'NP's.

$$[[most]] = \lambda P1 \lambda P2[most(x, P1(x), P2(x))]$$

3. In compositional semantics, heads contribute templates, and their specifiers/complements fill in slots in these templates.

The DP hypothesis

Semantic argument for the DP hypothesis:

1. Quantified sentences have a standard form:

$$Q \quad (x, \; P1(x), \quad P2(x) \;)$$

2. The 'semantic template' of a quantified sentence is contributed by the determiner of one of its 'NP's.

$$[[most]] = \lambda P1 \lambda P2[most(x, P1(x), P2(x))]$$

3. In compositional semantics, heads contribute templates, and their specifiers/complements fill in slots in these templates.

4. So determiners should be heads of 'noun phrases'.

Reference and quantification in discourse

Current semantic accounts of reference and quantification originate from work by Kamp and Heim.

- An indefinite DP introduces a referent into the discourse.
  A definite DP presupposes a referent with certain properties.

A dog came in $\longrightarrow$ 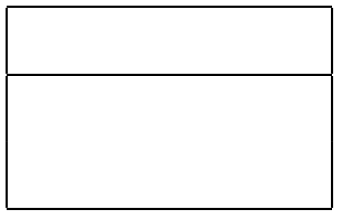

Reference and quantification in discourse

Current semantic accounts of reference and quantification originate from work by Kamp and Heim.

- An indefinite DP introduces a referent into the discourse.
  A definite DP presupposes a referent with certain properties.

A dog came in $\rightarrow$

x
dog(x) came-in(x)

Reference and quantification in discourse

Current semantic accounts of reference and quantification originate from work by Kamp and Heim.

- An indefinite DP introduces a referent into the discourse.
  A definite DP presupposes a referent with certain properties.

A dog came in $\rightarrow$

x
dog(x)
came-in(x)

The dog sat down $\rightarrow$

sat-down(y)
?y
?dog(y)

Reference and quantification in discourse

Current semantic accounts of reference and quantification originate from work by Kamp and Heim.

- An indefinite DP introduces a referent into the discourse.
  A definite DP presupposes a referent with certain properties.

A dog came in →

x
dog(x)
came-in(x)

The dog sat down →

sat-down(y)
?y
?dog(y)

- Quantifying determiners introduce referents too.
  These can only be referred back to in specific syntactic contexts.
  (And in these contexts, they contribute to a quantified expression.)

Many dogs bark →

x		
dog(x)	MANY	bark(x)

The referent-introducing function of D

Semantically, a determiner contributes a ‘referent’ (bound by a quantifier or free in the discourse), while a noun contributes a ‘property’ of that referent (either asserted or presupposed).

- The referent is ‘an anonymous individual’: we know it’s something, but we know nothing else about it.
- Quantified constructions tell us about *sets* of individuals with certain properties.

Summary: the contributions of DP and NP

Here's the idea so far:

- D introduces a referent.
- N associates a category with that referent.

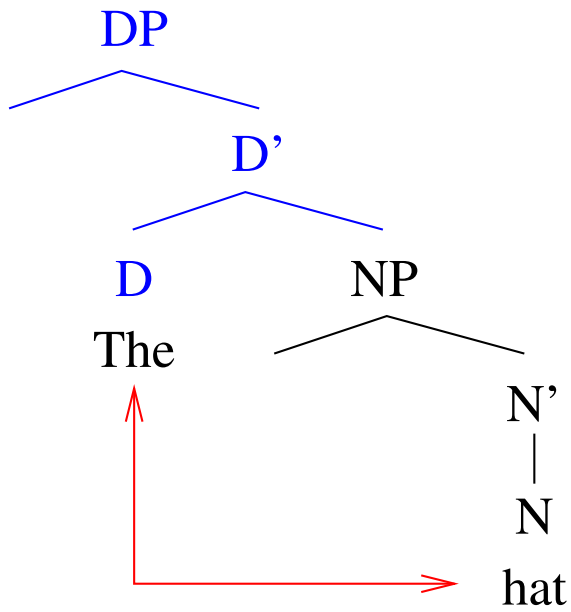

Number

Compare: *There is a dog outside / There are sm dogs outside*.

- In each case, the determiner contributes a single referent, which can be referred to presuppositionally. (By *it* or *they*.)
- But in one case, the referent is a single individual; in the other, it is a group of individuals (all of which are dogs).

Where in a DP does information about number originate?

- Noun stems don't carry it.
- 'Introducing a referent' is neutral as regards the number of the referent.
- Maybe number information originates at D, or at N?
- Or maybe number is contributed by a separate XP?

Predicate nominals and NumP

Some ‘DPs’, called predicate nominals, don’t contribute a referent.

This animal is a dog. *These animals are sm dogs.*

But note:

- They do carry information about number.
- They do have determiners.

Zamparelli (1989): there are two ‘determiner’ positions:

- The determiner in predicate nominals just carries number information. (We could call it NumP.)
- ‘Regular DPs’ have two layers of determiners: first a DP (contributing a referent), and then a NumP (contributing number).

Existential sentences and NumP

A similar argument can be framed for DPs in existential sentences.

There is a dog outside. *There are sm dogs outside.*

- 'There' contributes a 'bare referent'. (Perhaps a point in space.)
- The sentence predicates a property of this referent.
- The predicate nominal doesn't introduce a referent, yet it does have a determiner, and it does convey number information.
- The predicate nominal is a NumP, stripped of its DP.

Summary: a simple syntactic account of DPs

‘Referential’ and ‘quantifying’ nominals look like this:

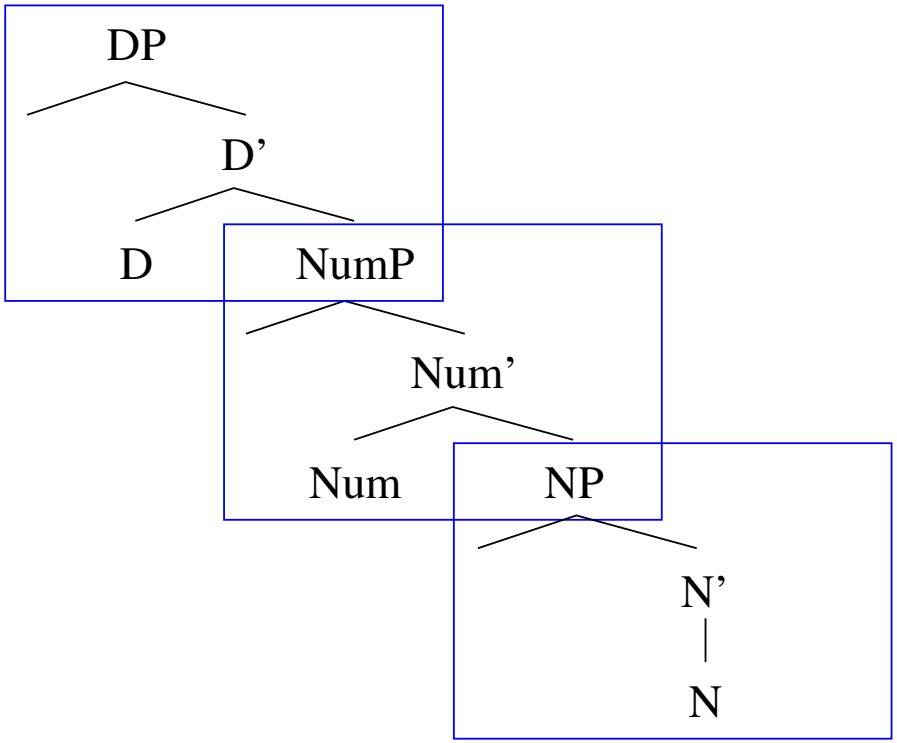

Predicate nominals (and those in existential sentences) look like this:

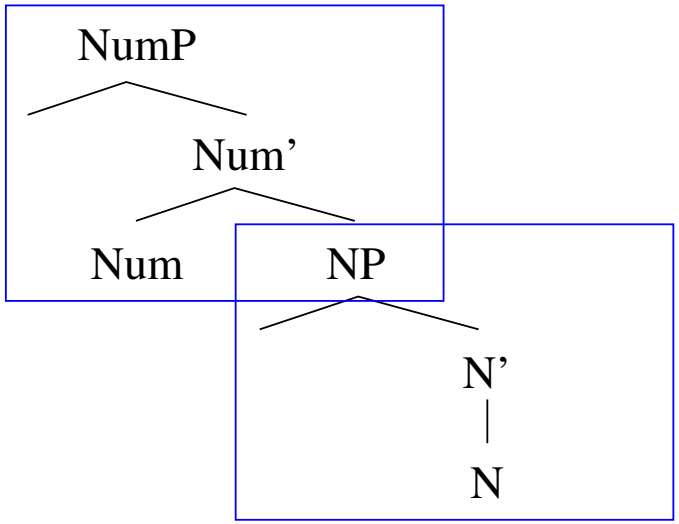

Summary: a simple syntactic account of DPs

‘Referential’ and ‘quantifying’ nominals look like this:

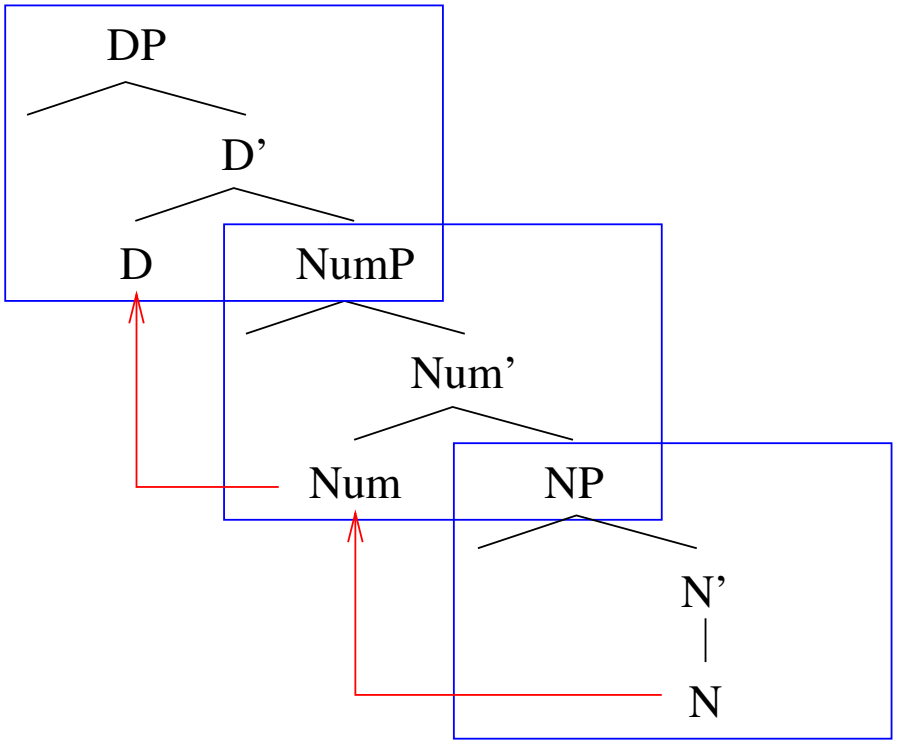

Predicate nominals (and those in existential sentences) look like this:

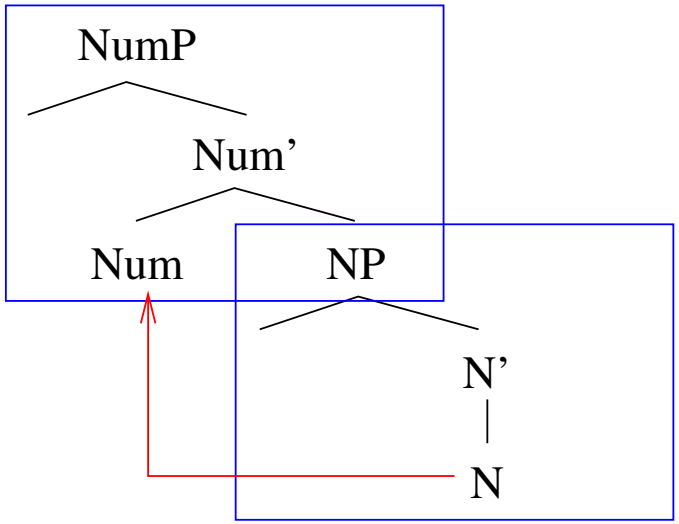

In each case, transmission of agreement features happens through head movement.

What's up next

- Outline a basic model of object perception
- Extend it to deal with groups, predication, existentials
- See if there's anything in the perceptual model which might correspond to DP, NumP and NP.

Outline of today's lecture

Neural pathways involved in object perception

Here's the model I presented way back in Lecture 2.

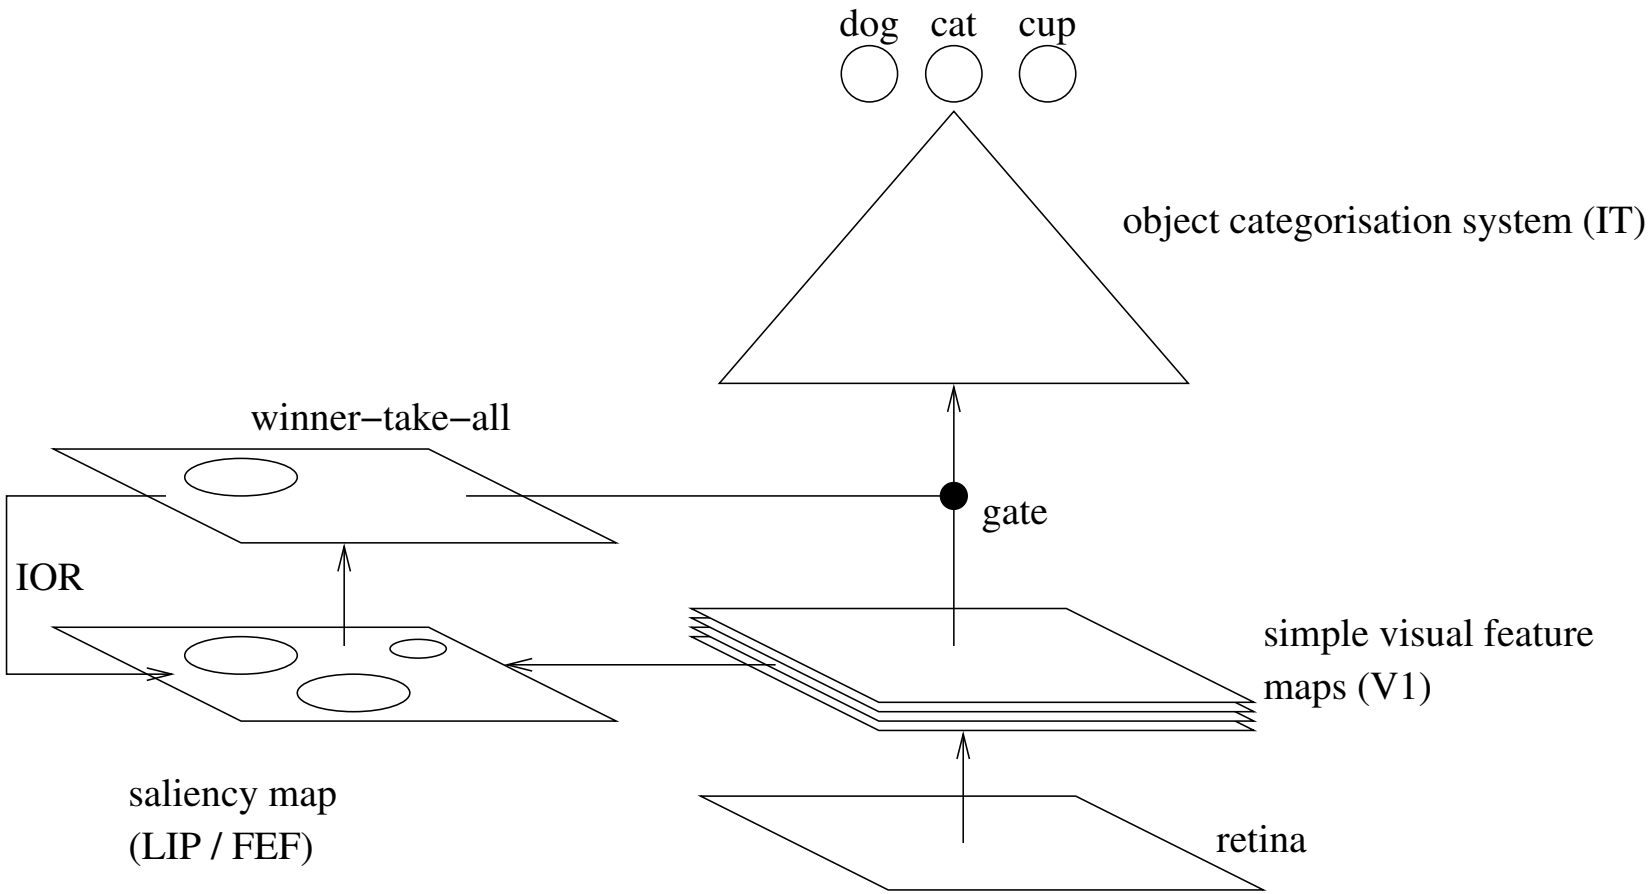

Evidence for the model: recap

- Visual feature maps: Hubel and Wiesel (1968)
- Saliency map: Itti *et al.* (1998); Gottlieb *et al.* (1998); Bichot *et al.* (2001); Schall (2004)
- Inhibition-of-return: Posner *et al.* (1984)
- Object categorisation: Logothetis *et al.* (1995); Tanaka (1996); Kourtzi and Kanwisher (2001)
- Gating of V1 activity by FEF: Moore & Armstrong (2003); Grosbras & Paus (2003)

A model of the visual categorisation system

The categorisation system is modelled as a convolutional NN. Le Cun & Bengio (1995); Mozer & Sitton (1996); Riesenhuber & Poggio (1999)

- Its input is a map of simple visual features.
- Each layer takes a map of features and returns a map of combined features.
- To avoid a combinatorial explosion, each layer also abstracts over space.

This is a reasonably good model of cells in the IT pathway.

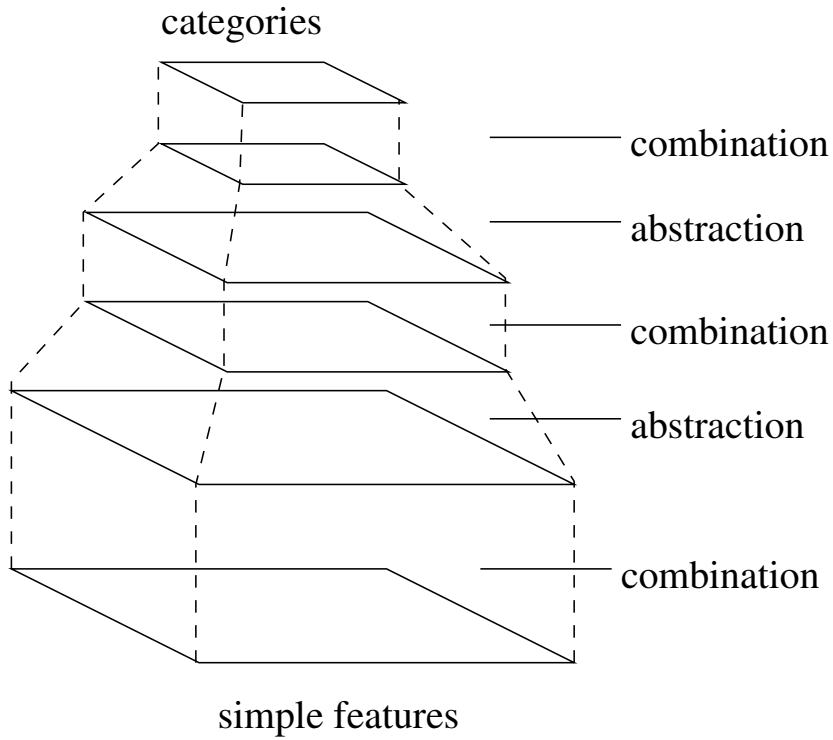

A model of the visual categorisation system

Combination layers look like this:

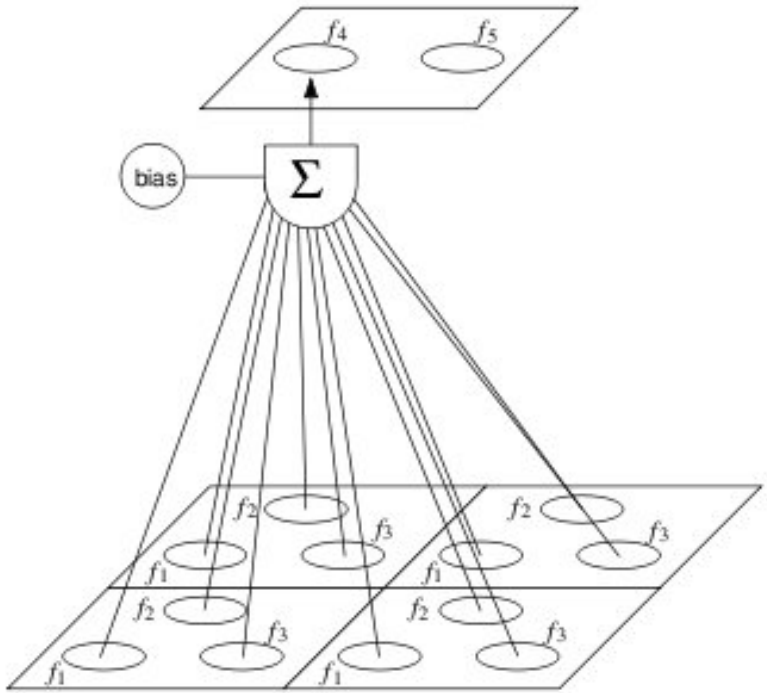

A model of the visual categorisation system

Abstraction layers look like this:

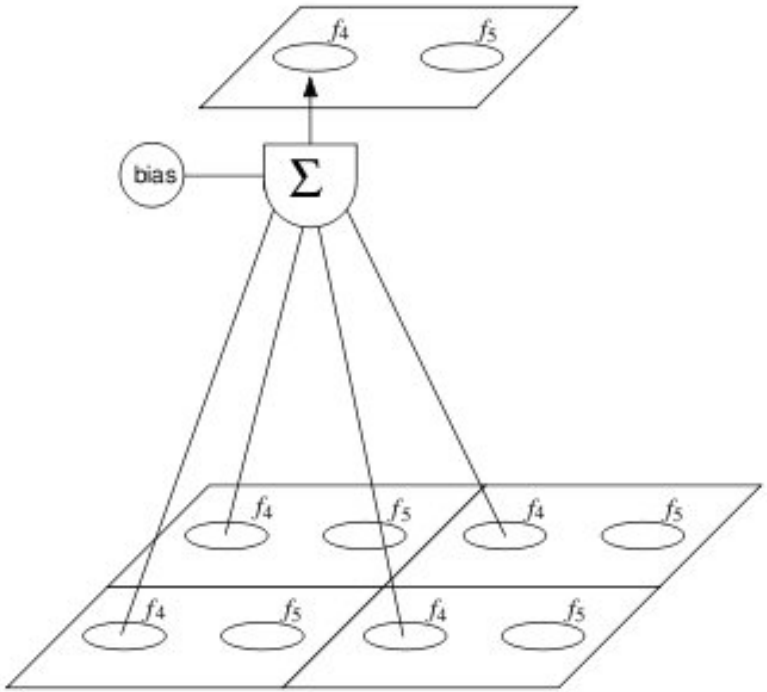

Translation and scale invariance of a convolutional NN

The abstraction operations allow an object to be categorised anywhere on the retina.

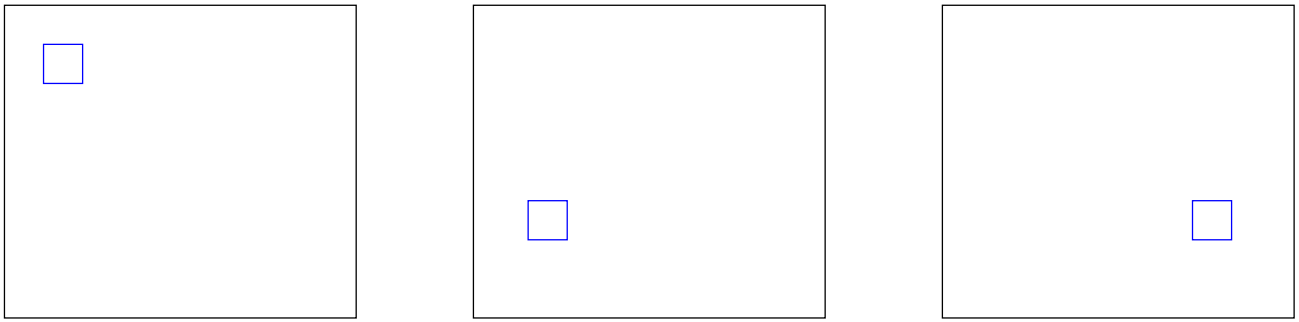

Input feature maps of different spatial frequencies allow an object to be categorised at a range of sizes.

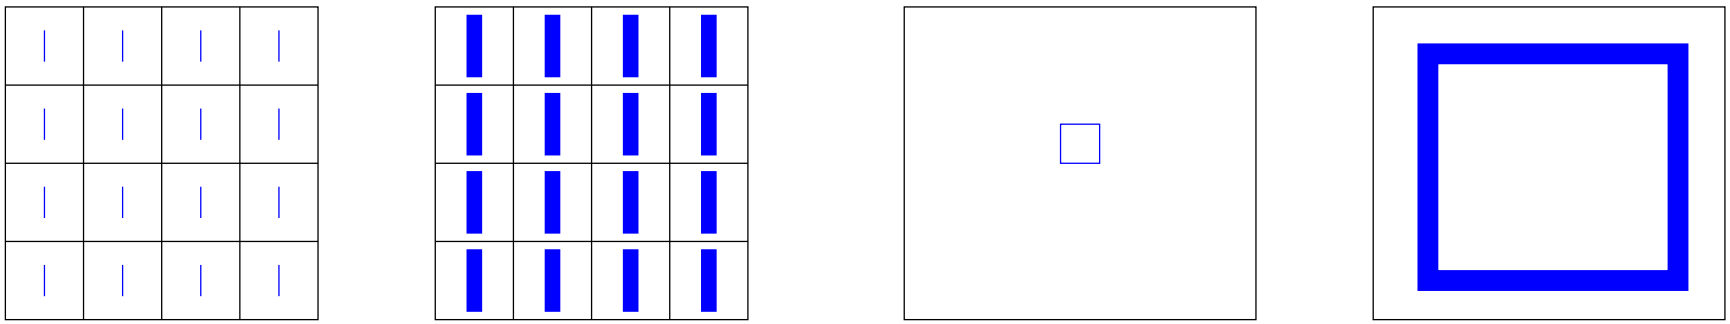

The sequential structure of object establishment

When you are attending to an object's location and you have evoked a category in the categorisation system, you have established it.

Note that perceptual establishment is necessarily sequential:

- You must establish the location of an object before you can determine its other properties.

Perceptual correlates of DP syntax: an initial idea

One suggestion:

- A *right-branching structure* of XPs denotes a *sequence* of perceptual operations.

For DPs, we might then speculate:

- DP signals the selection of a winning saliency map region;
- NP signals the activation of an assembly in the categorisation system.

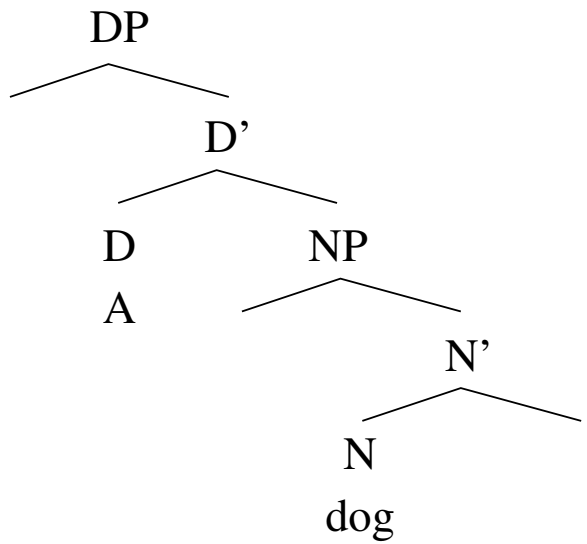

What about NumP?

We now need a perceptual interpretation for NumP.

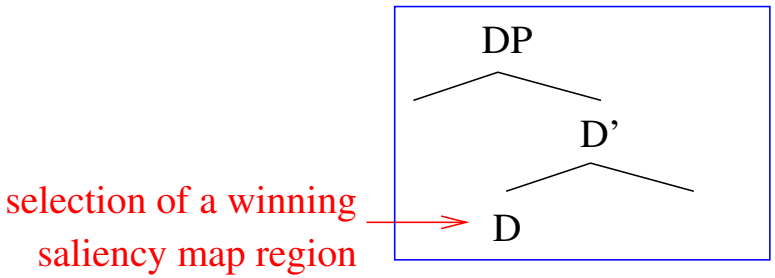

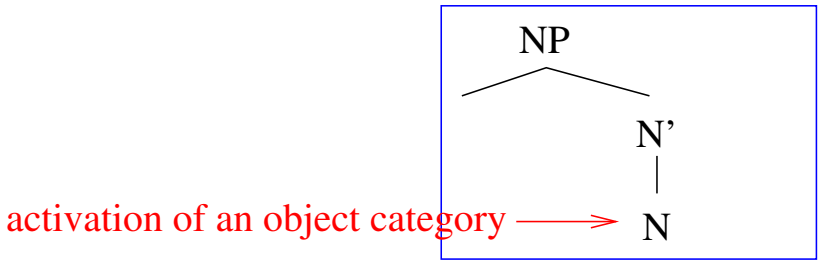

What about NumP?

We now need a perceptual interpretation for NumP.

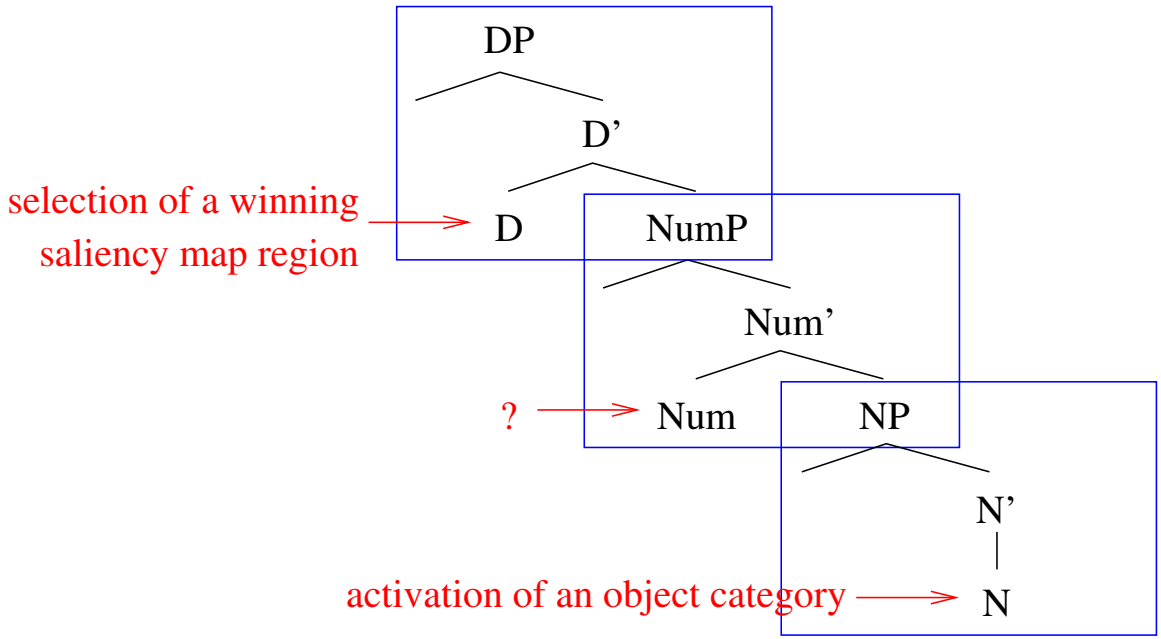

Outline of today's lecture

Group categorisation

An interesting thing about convolutional NNs: for a group of N objects of the same category, they are *blind to cardinality*. (Walles *et al.*, 2008)

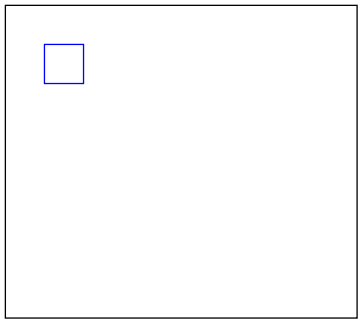

= 'square'

Group categorisation

An interesting thing about convolutional NNs: for a group of N objects of the same category, they are *blind to cardinality*. (Walles *et al.*, 2008)

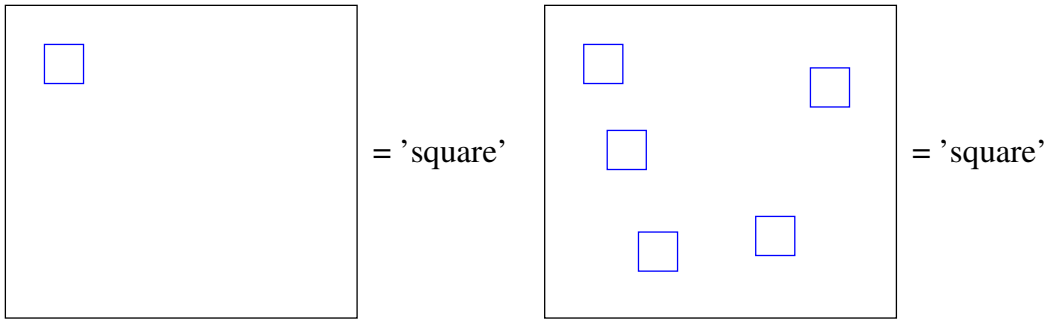

Group categorisation

An interesting thing about convolutional NNs: for a group of N objects of the same category, they are *blind to cardinality*. (Walles *et al.*, 2008)

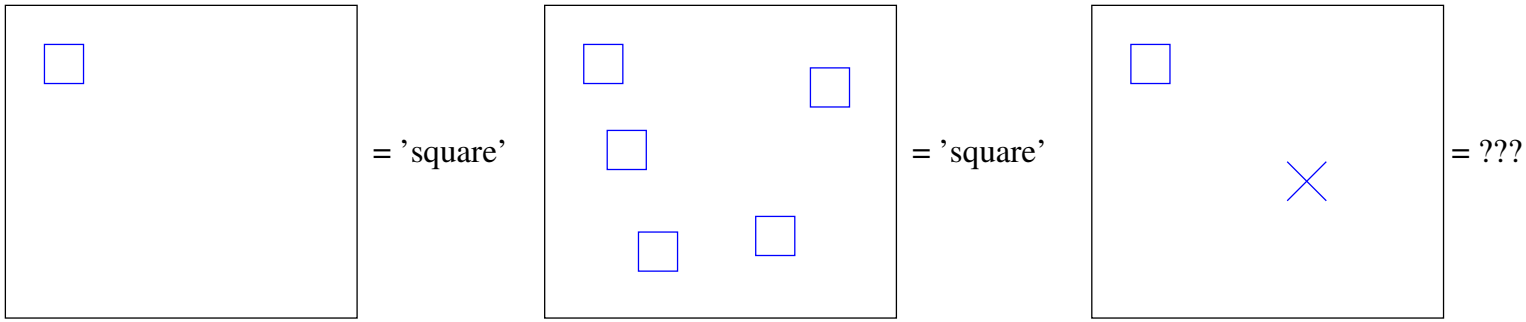

Group categorisation

An interesting thing about convolutional NNs: for a group of N objects of the same category, they are *blind to cardinality*. (Walles *et al.*, 2008)

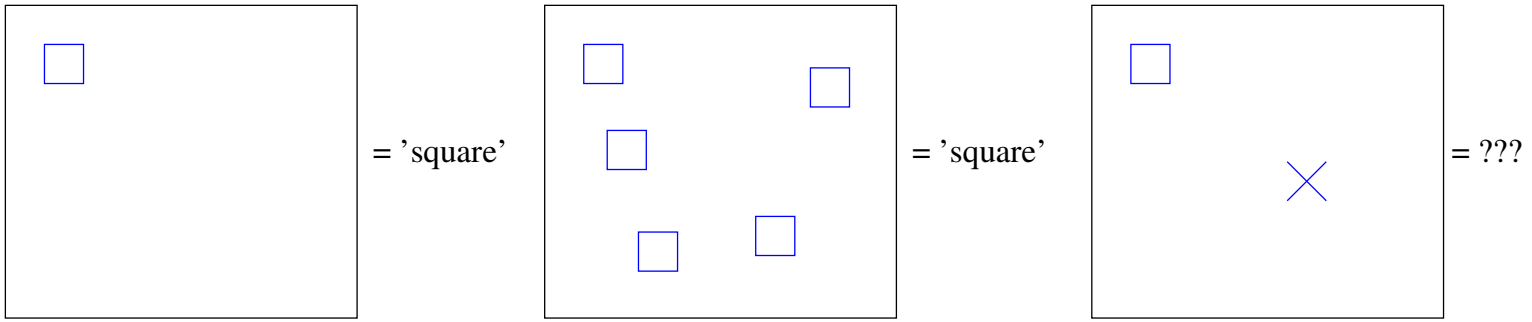

This property emerges from the spatial abstraction operation.

Group categorisation

An interesting thing about convolutional NNs: for a group of N objects of the same category, they are *blind to cardinality*. (Walles *et al.*, 2008)

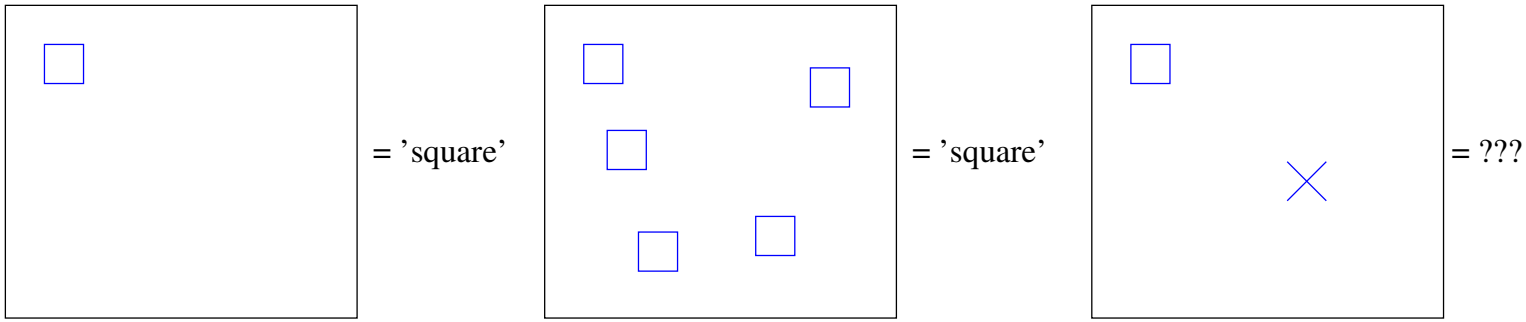

This property emerges from the spatial abstraction operation.

The categorisation system is thus very well placed to deliver the denotation of a noun stem.

One dog *Two dog-s*

Group categorisation

An interesting thing about convolutional NNs: for a group of N objects of the same category, they are *blind to cardinality*. (Walles *et al.*, 2008)

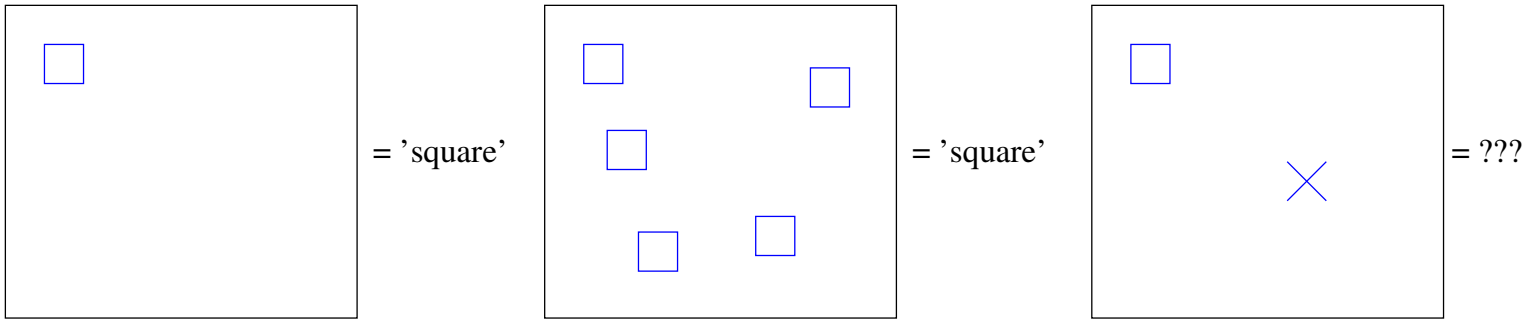

This property emerges from the spatial abstraction operation.

The categorisation system is thus very well placed to deliver the denotation of a noun stem.

One dog *Two dog-s*

We now need additional machinery to provide Number / cardinality.

The attentional and object categorisation systems

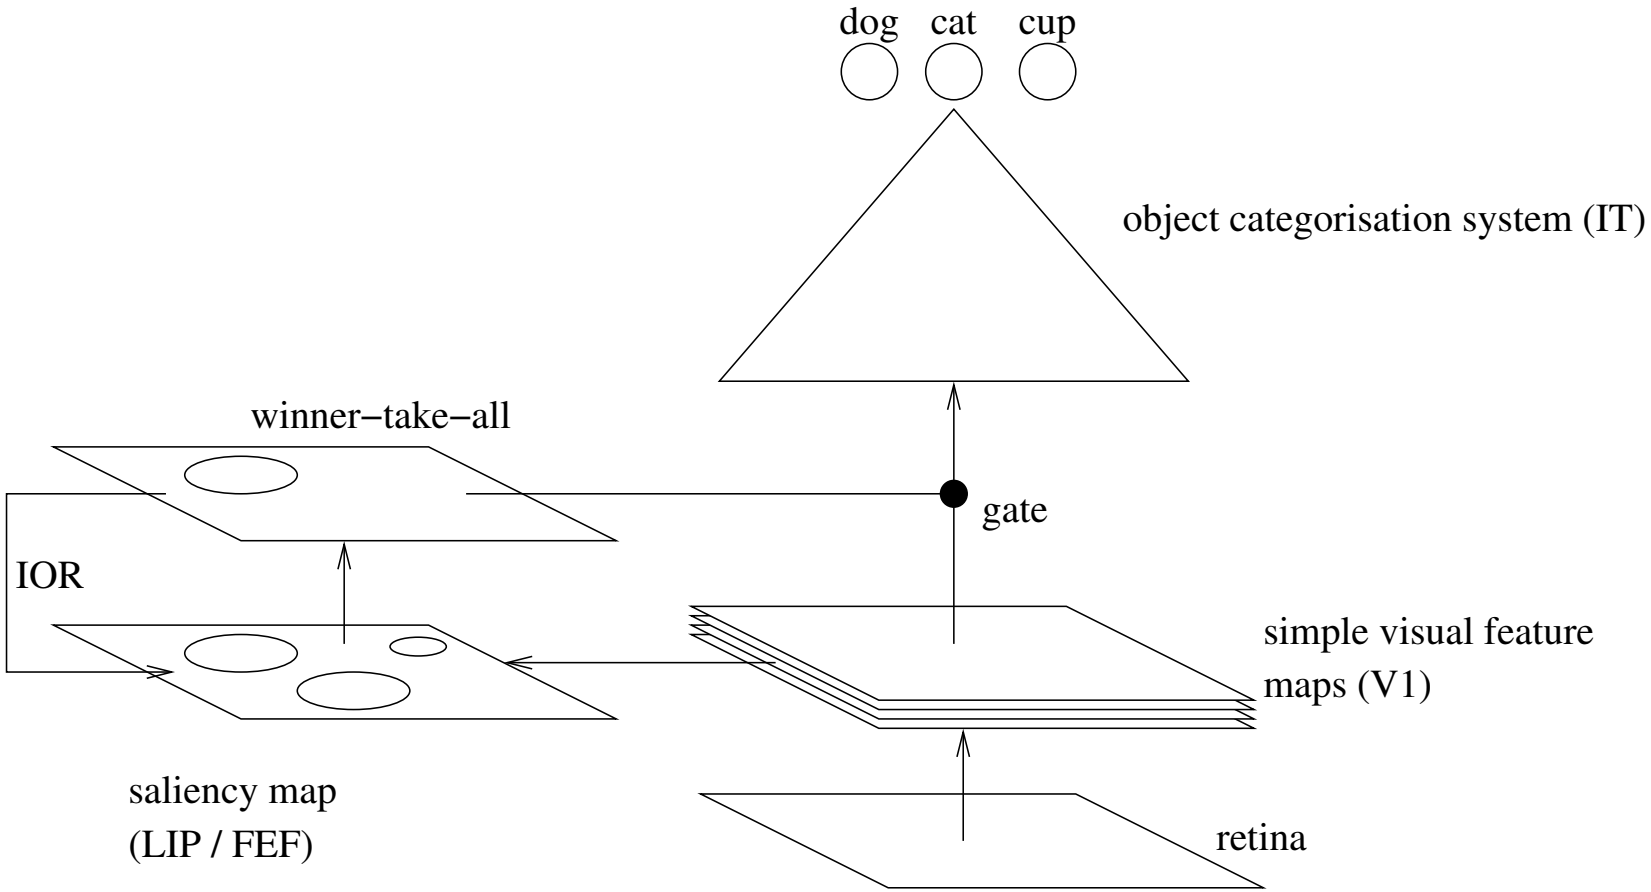

Configuring attention for a cardinality-blind categoriser

If the categoriser can recognise groups of type-identical objects, this changes the way attention should select its input.

- Salient regions needn't just contain single objects. They can also contain 'homogeneous groups'.
- (Note that 'similarity' is a well-known Gestalt principle.)

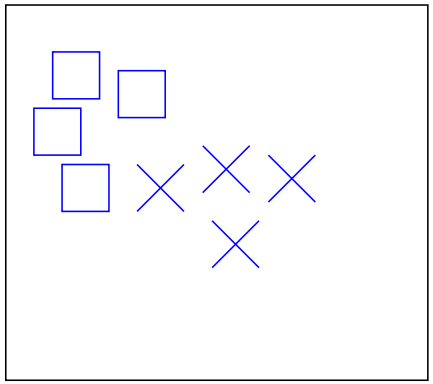

Configuring attention for a cardinality-blind categoriser

If the categoriser can recognise groups of type-identical objects, this changes the way attention should select its input.

- Salient regions needn't just contain single objects. They can also contain 'homogeneous groups'.
- (Note that 'similarity' is a well-known Gestalt principle.)

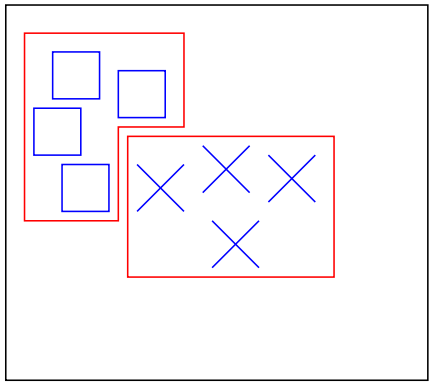

Another look at spatial scale

Note that there can be salient regions of different sizes.

What happens if there is a conflict between large- and small-scale category information?

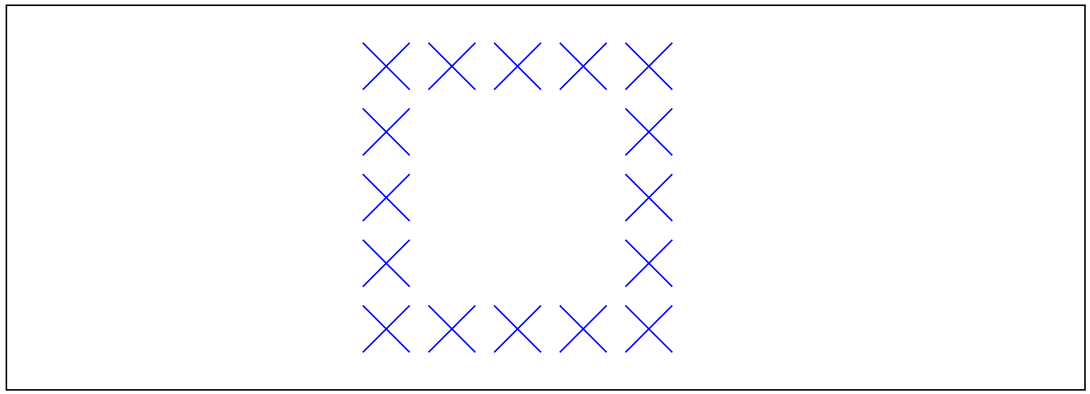

- There must be an attentional mechanism for *selecting a particular scale*.
- Navon (1977): large-scale stimuli are processed first.

Another look at spatial scale

Note that there can be salient regions of different sizes.

What happens if there is a conflict between large- and small-scale category information?

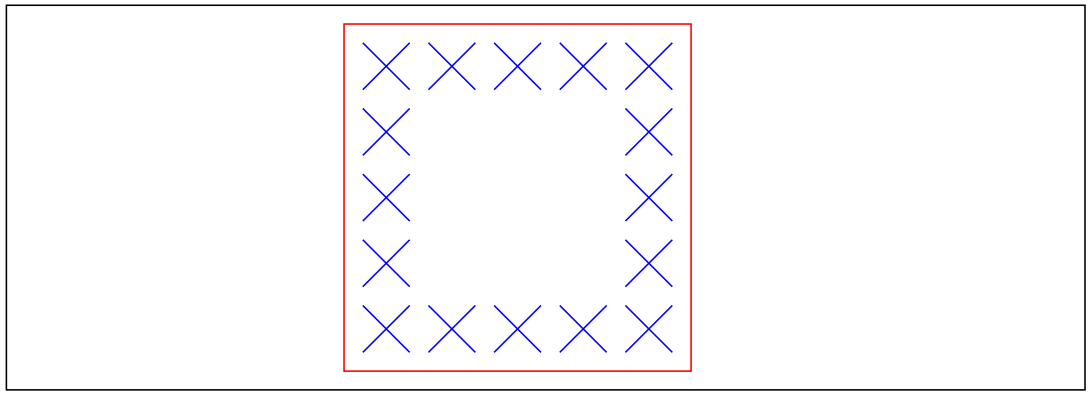

- There must be an attentional mechanism for *selecting a particular scale*.
- Navon (1977): large-scale stimuli are processed first.

Selection of spatial scales for categorisation

Idea: the saliency map supports location selection and scale selection.

- The most-salient region selects a corresponding retinal region for processing by the categorisation system.
- The *size* of this region selects an appropriate *scale* of feature map.

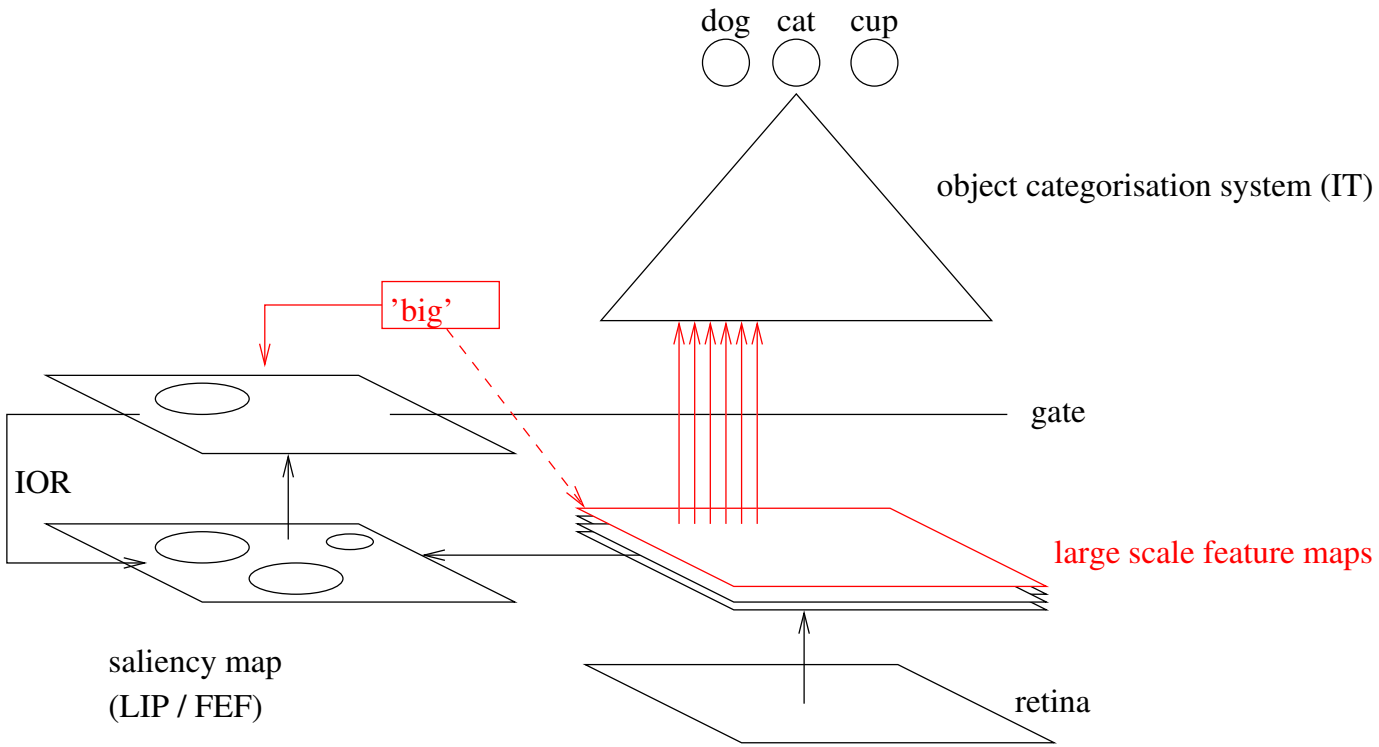

Selection of spatial scales for categorisation

Idea: the saliency map supports location selection and scale selection.

- The most-salient region selects a corresponding retinal region for processing by the categorisation system.
- The *size* of this region selects an appropriate *scale* of feature map.

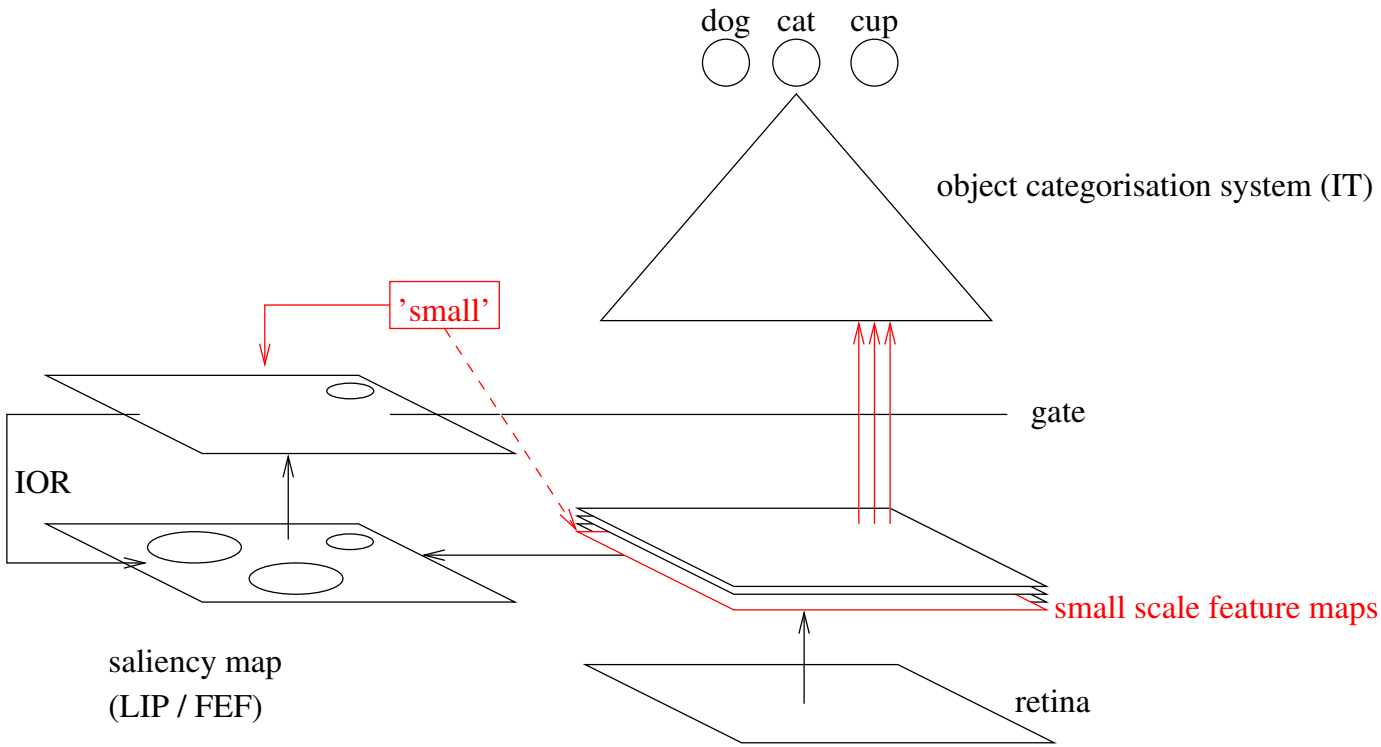

A perceptual correlate of singular and plural

To perform group categorisation, we need to change scale *without changing location*.

- We envisage an operation of ‘scale IOR’: switching to a spatial scale smaller than that given by default.

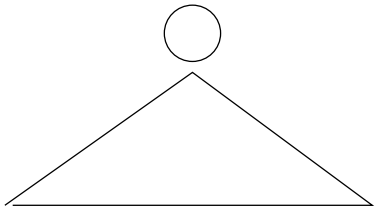

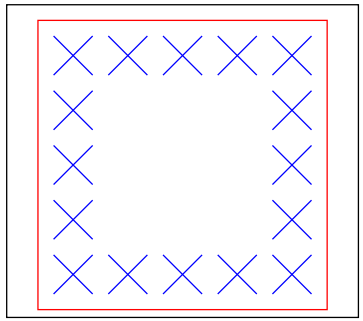

A perceptual correlate of singular and plural

To perform group categorisation, we need to change scale *without changing location*.

- We envisage an operation of 'scale IOR': switching to a spatial scale smaller than that given by default.

"A single square":

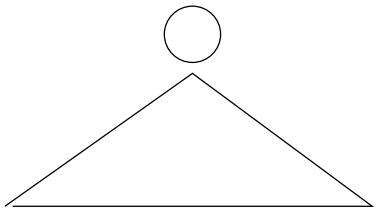

default scale feature maps

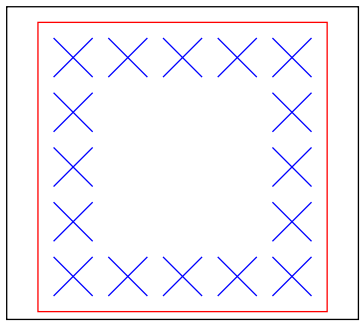

A perceptual correlate of singular and plural

To perform group categorisation, we need to change scale *without changing location*.

- We envisage an operation of 'scale IOR': switching to a spatial scale smaller than that given by default.

"A single square":

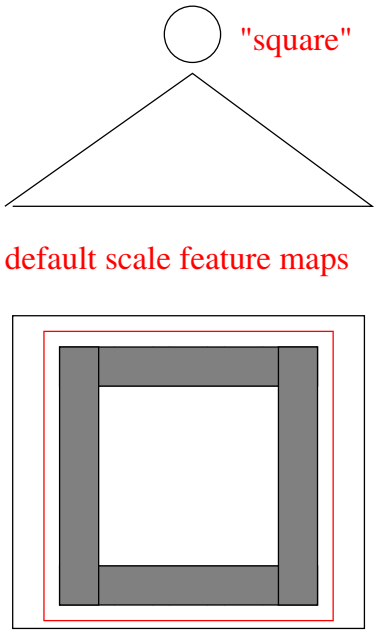

A perceptual correlate of singular and plural

To perform group categorisation, we need to change scale *without changing location*.

- We envisage an operation of ‘scale IOR’: switching to a spatial scale smaller than that given by default.

"More than one cross":

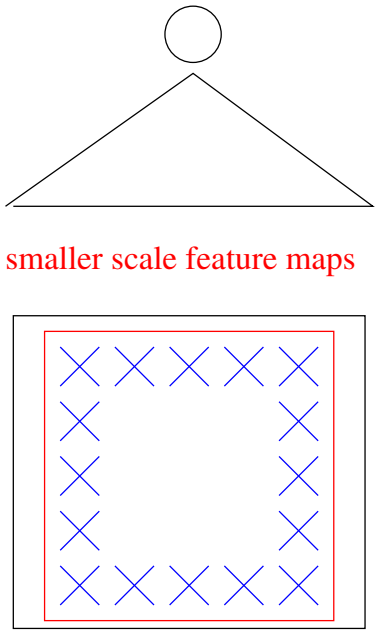

A perceptual correlate of singular and plural

To perform group categorisation, we need to change scale *without changing location*.

- We envisage an operation of ‘scale IOR’: switching to a spatial scale smaller than that given by default.

"More than one cross":

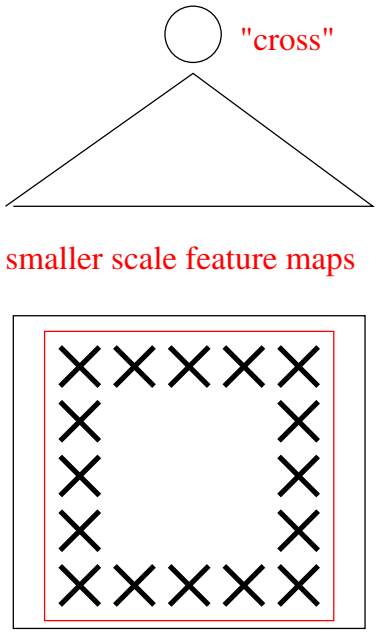

A perceptual correlate of singular and plural

To perform group categorisation, we need to change scale *without changing location*.

- We envisage an operation of 'scale IOR': switching to a spatial scale smaller than that given by default.

"More than one cross":

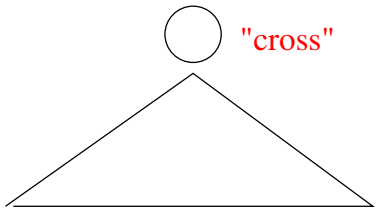

smaller scale feature maps

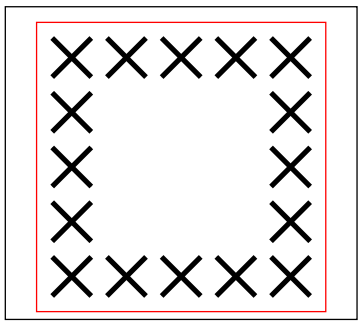

A perceptual definition of singular and plural:

- Singular = categorisation at default spatial scale
- Plural = Categorisation at a smaller spatial scale.

Outline of today's lecture

A perceptual correlate of NumP?

Note that the decision about spatial frequency has to come after a location has been selected, but before categorisation occurs.

Proposal: Num signals the operation of selecting a spatial frequency.

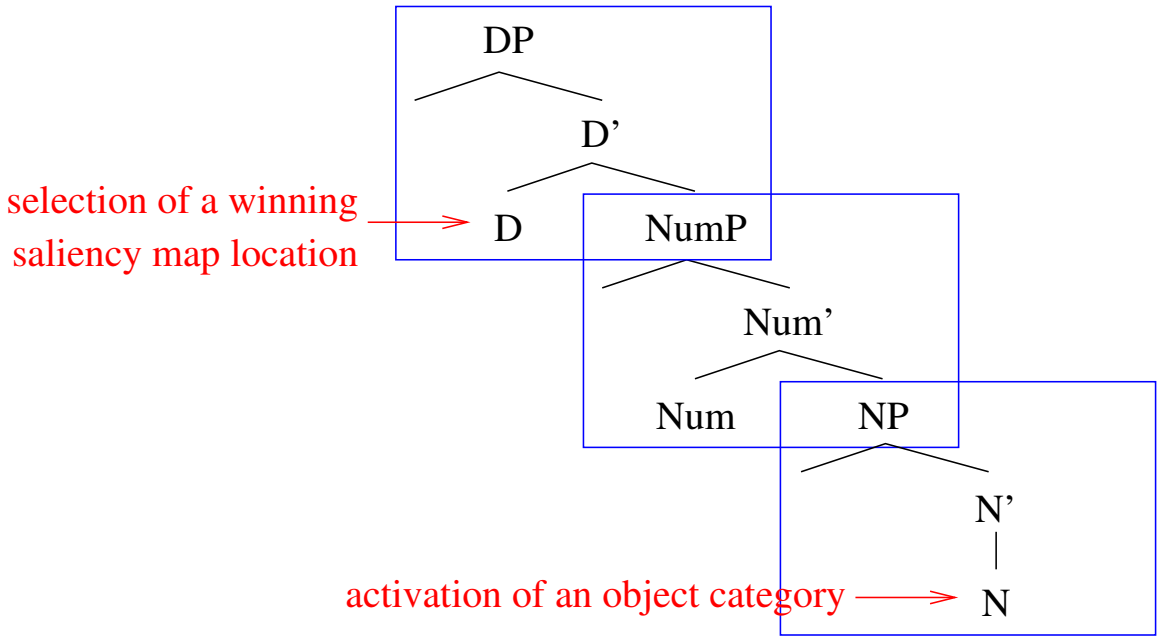

A perceptual correlate of NumP?

Note that the decision about spatial frequency has to come after a location has been selected, but before categorisation occurs.

Proposal: Num signals the operation of selecting a spatial frequency.

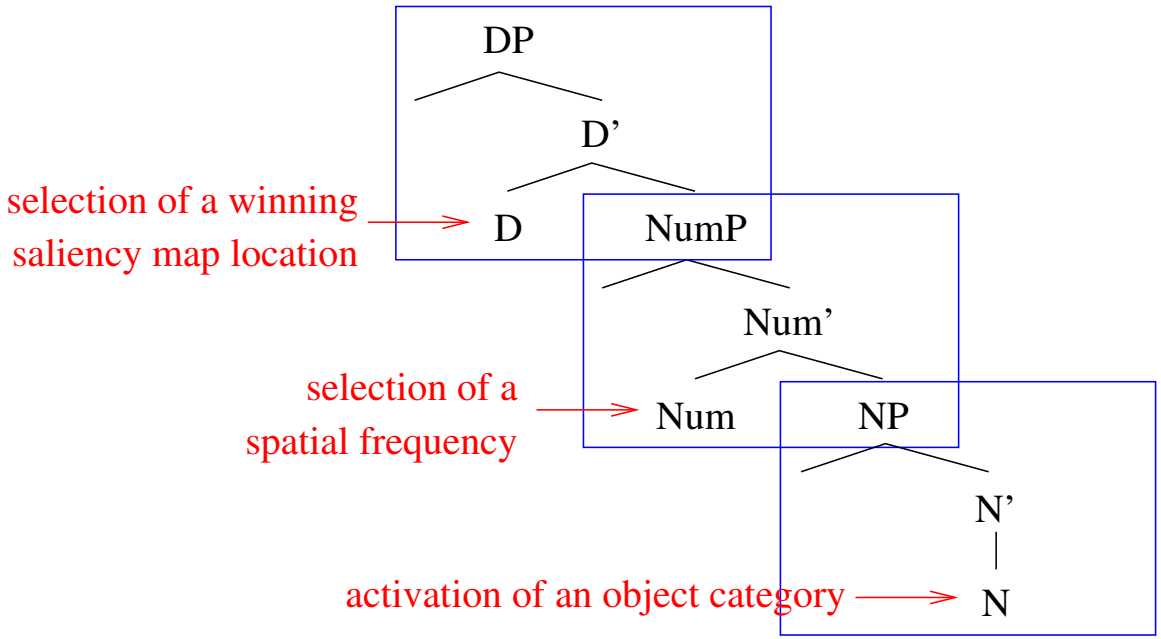

Head movement

Recall: determiners and nouns can be inflected for number, person and gender.

- In Minimalism, a fully inflected noun carrying a ‘number feature’ is ‘generated’ at N. It raises to Num to ‘check’ number, and then to D.
- What might the perceptual correlates of ‘N raising’ be?

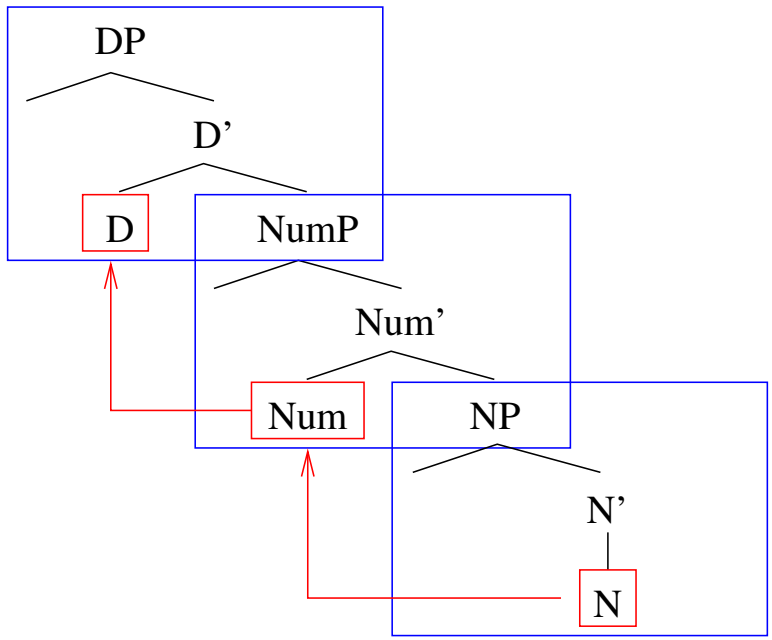

Working memory representations of objects

Proposal: a DP describes an attentional sequence as recalled from working memory (WM), rather than 'as it happens'.

1. Objects are represented in WM as *prepared attentional sequences*. (Call each object representation a WM individual.)
   - WM object representations are in prefrontal cortex (PFC). They include representations of location, category and cardinality (Rainer *et al.*, 1998; Nieder & Miller, 2004).
   - PFC is also where prepared action sequences are stored (Barone & Joseph, 1989).

2. DPs describe attentional sequences *as replayed from WM*.

WM replay and head movement

A prepared action sequence in PFC:

- is *tonically active* throughout execution of the sequence;
- includes representations of each component action.

(Averbeck *et al.*, 2002)

Proposal: heads of XPs denote tonically active *prepared* actions rather than transitory action signals.

- So at each stage in a replayed sequence, we have access to representations of all three (prepared) actions.

Interpreting X-bar structures as replayed sequences

Original interpretation:

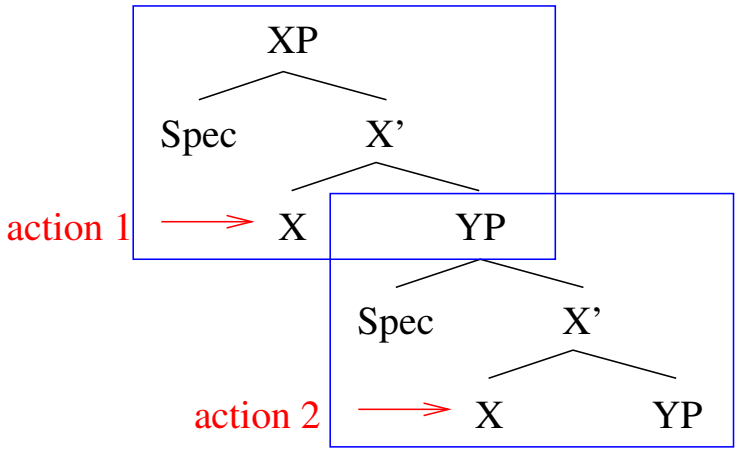

Interpreting X-bar structures as replayed sequences

New interpretation:

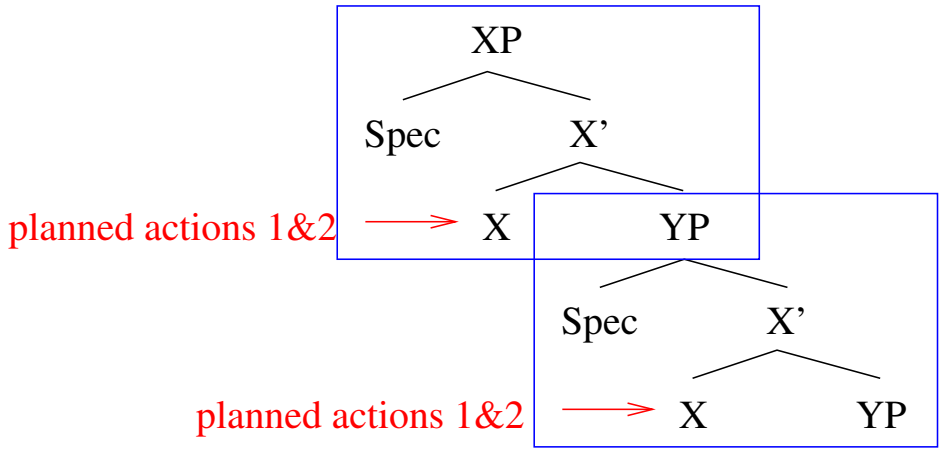

The DP-clause interface

Recall: DPs appear at positions within a clause.

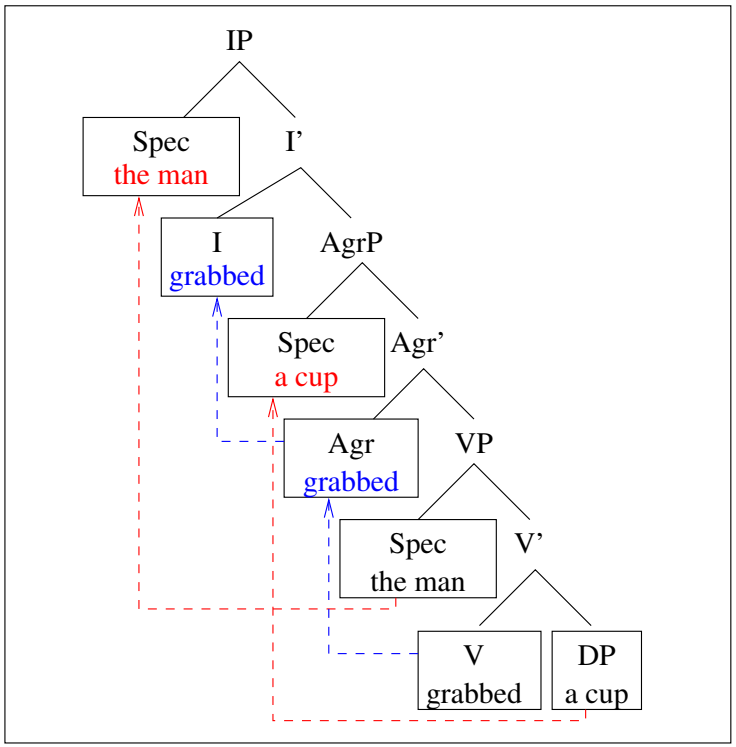

If a DP is to be understood as describing a replayed attentional sequence in its own right, then how do we interpret the appearance of a DP at a position within a clause?

The DP-clause interface

Proposal:

- Each attentional sequence in WM is associated with a referent.
- 'DP positions' in a clause represent moments when a particular referent becomes active.
- When a referent is active it is an *opportunity* for its associated attentional sequence to be rehearsed.
- There are two opportunities to rehearse both subject and object referents in a transitive clause.
- An infant has to learn which opportunity to take.

Extending the clause-level network

Here's the network for reading out the LF of a clause, from the previous lecture.

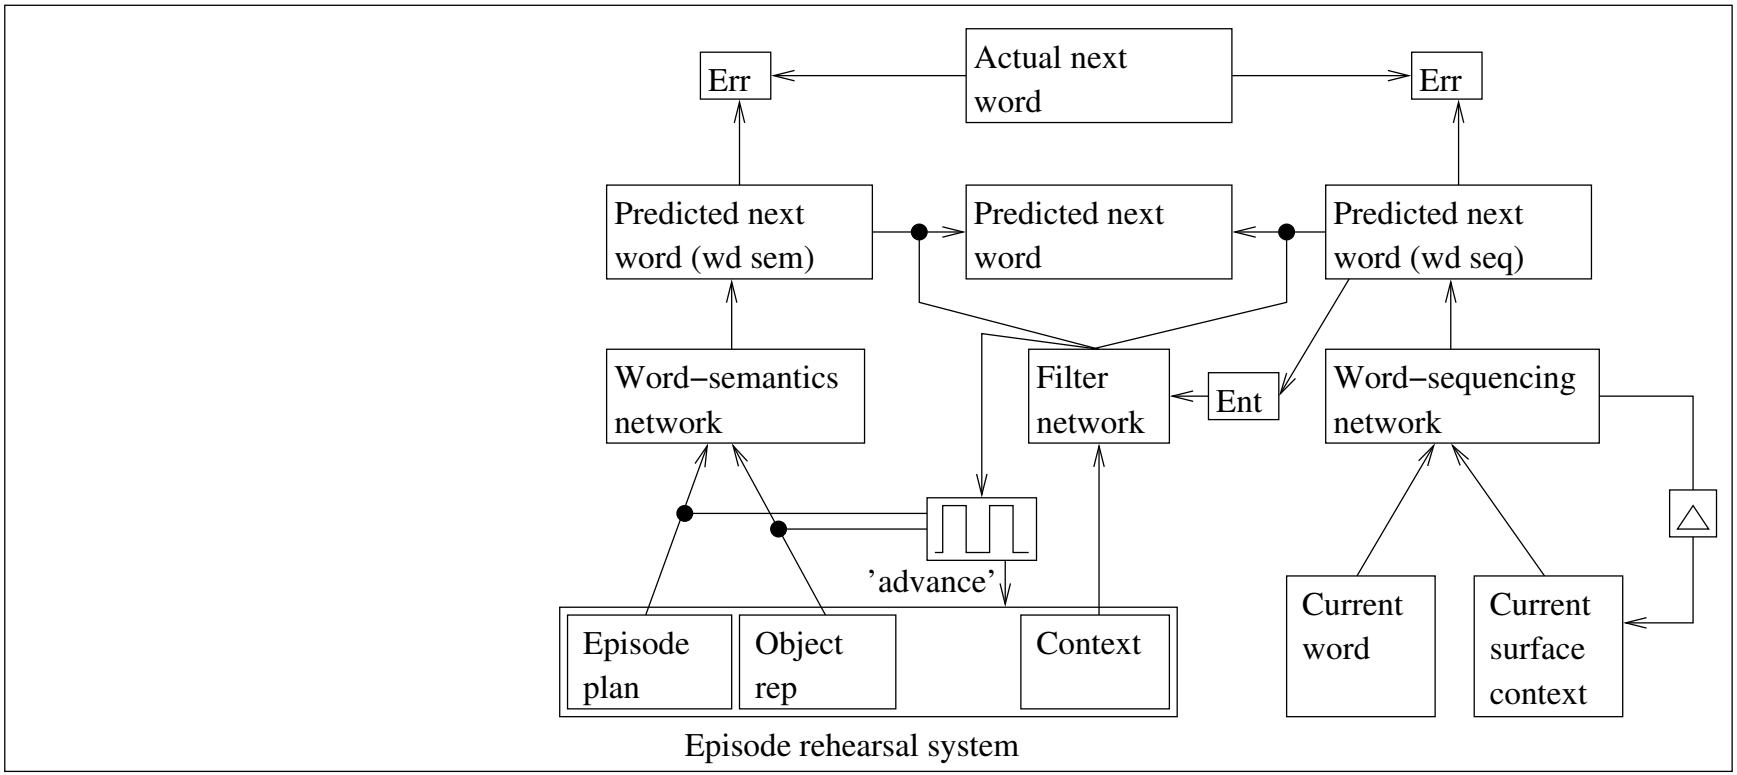

Extending the clause-level network

Proposal: as well as an episode rehearsal system, there's an attentional sequence rehearsal system.

- This system takes an attentional sequence plan, and generates a stream of SM representations.

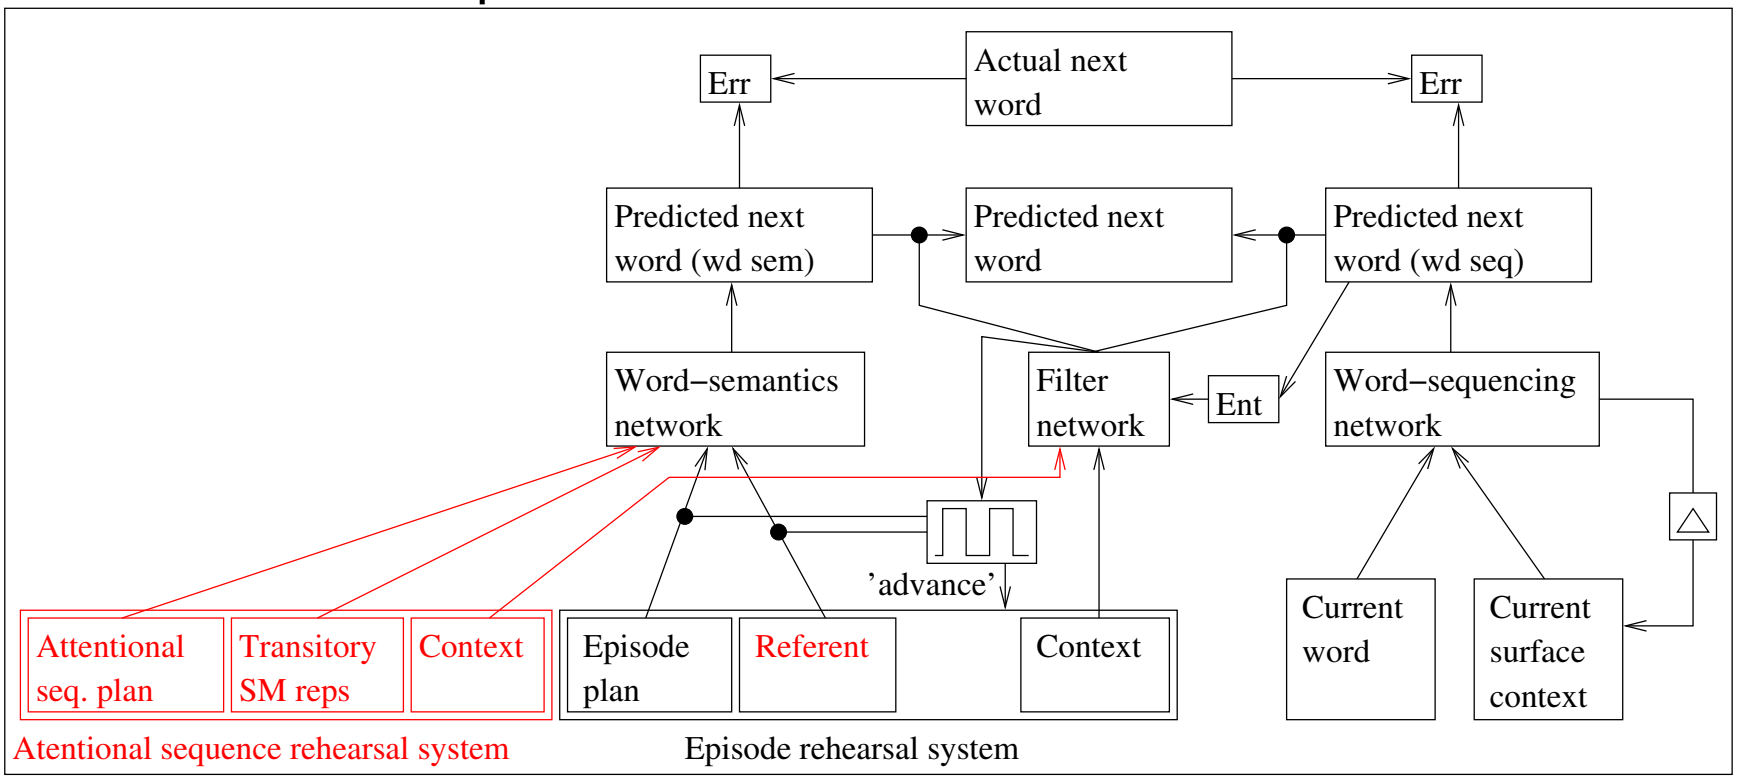

Extending the clause-level network

The episode rehearsal system generates opportunities to activate different attentional sequence plans.

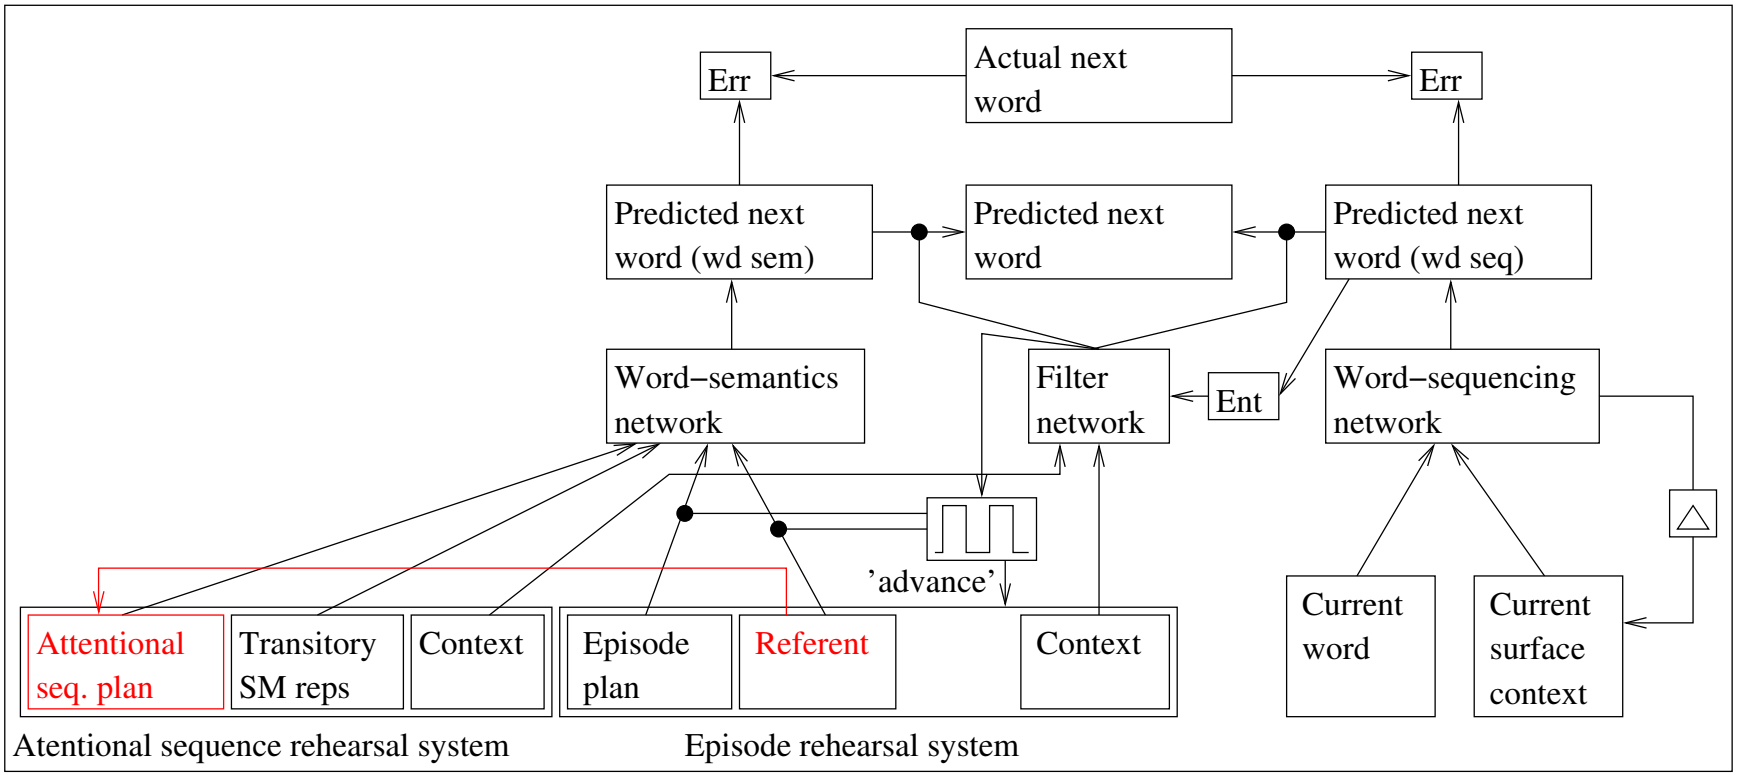

Extending the clause-level network

Proposal: the filter network can learn to interrupt the episode rehearsal system, and (temporarily) drive linguistic output from the attentional sequence rehearsal system.

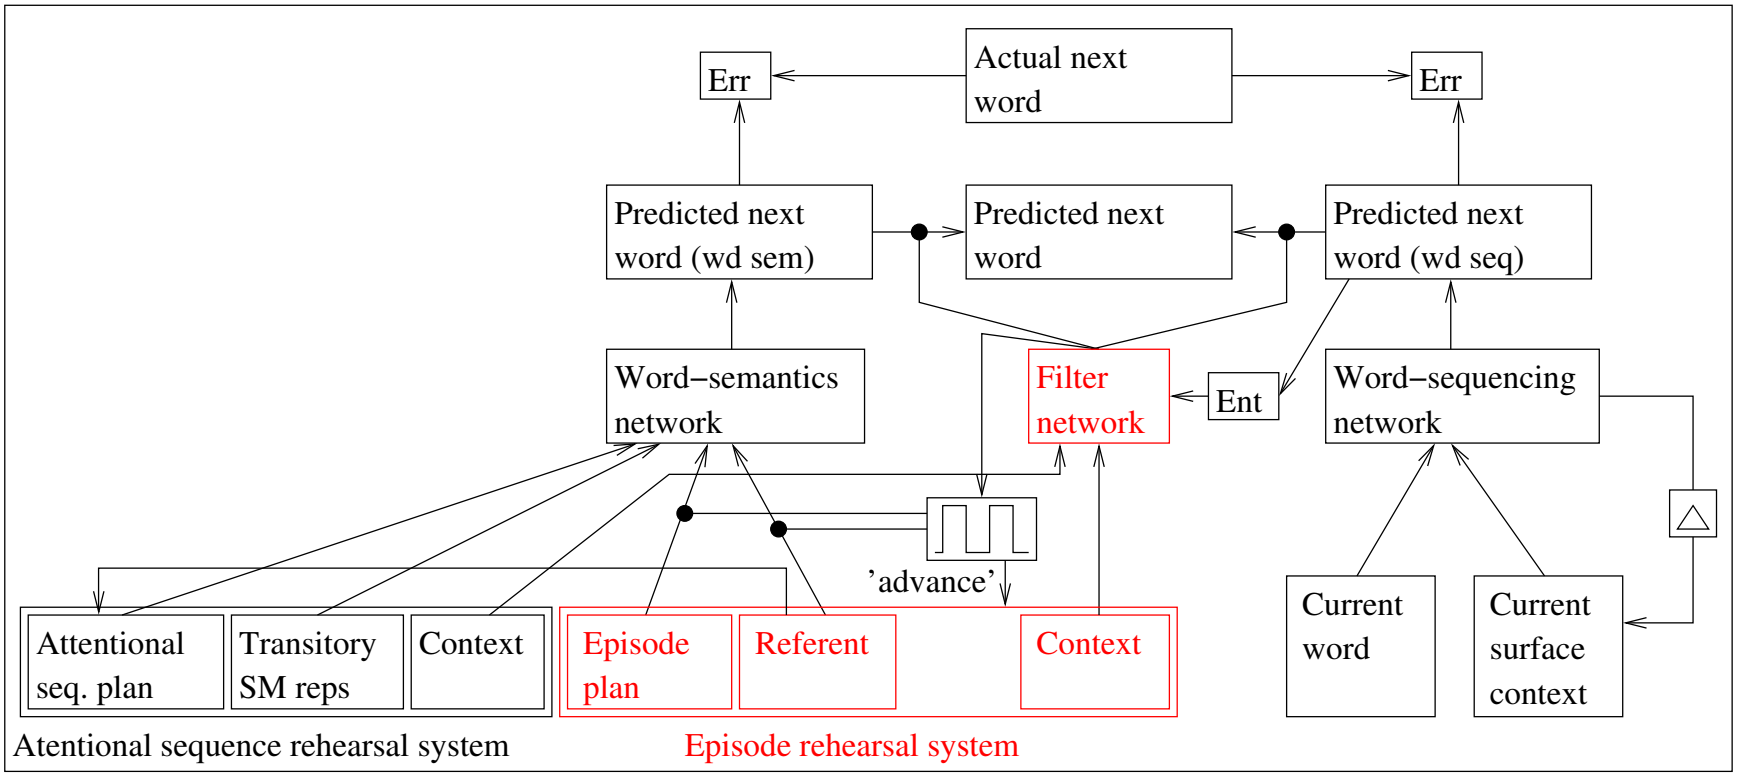

Extending the clause-level network

Proposal: the filter network can learn to interrupt the episode rehearsal system, and (temporarily) drive linguistic output from the attentional sequence rehearsal system.

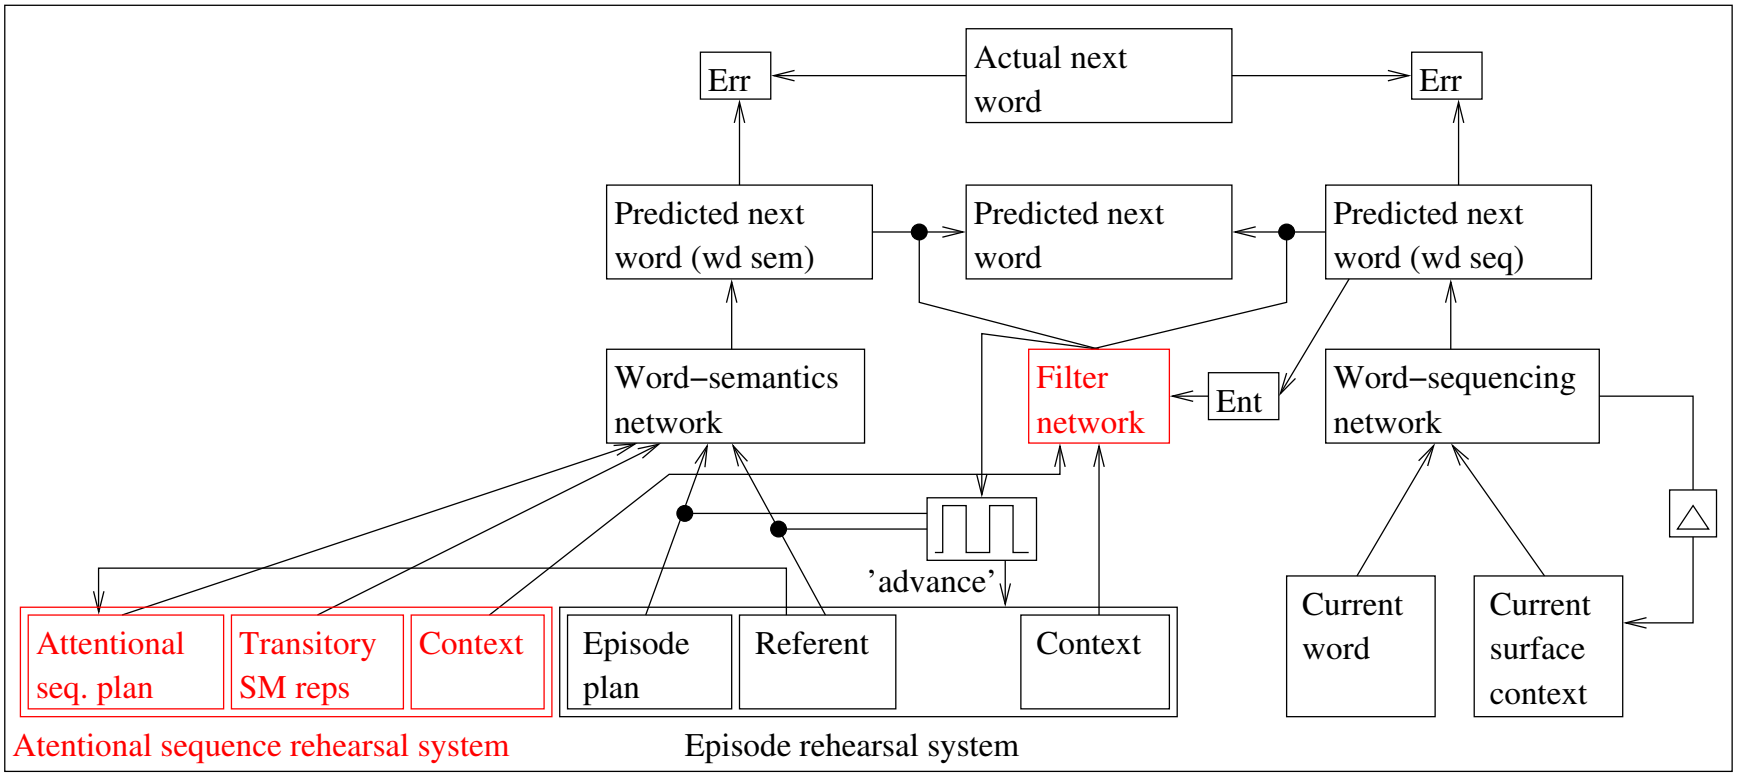

Perceiving properties of objects

What are the perceptual processes involved in noting that an object 'has a certain property'?

The dog was brown. *The dog was a dachshund.*

1. The categorisation system is likely to be involved. Because properties can be categories.

- (But not all properties are categories.)

2. Perceiving a property in an object probably has a sequential structure:

- First we establish the object.
- Then we 'attend to' a property.

Property complexes and object types

The categorisation system establishes a rich complex of properties.

- This complex goes a long way towards identifying an object as a *token* (e.g. Bar et al, 2001).
- So where are *object types* evoked in the categorisation system?

Proposal:

- Properties are linked in assemblies (through Hebbian learning).
- Types are assemblies.
- There is a level in the categorisation system where assemblies compete with one another.

Competition between property assemblies

The property complex layer holds a rich set of properties representing an individual.

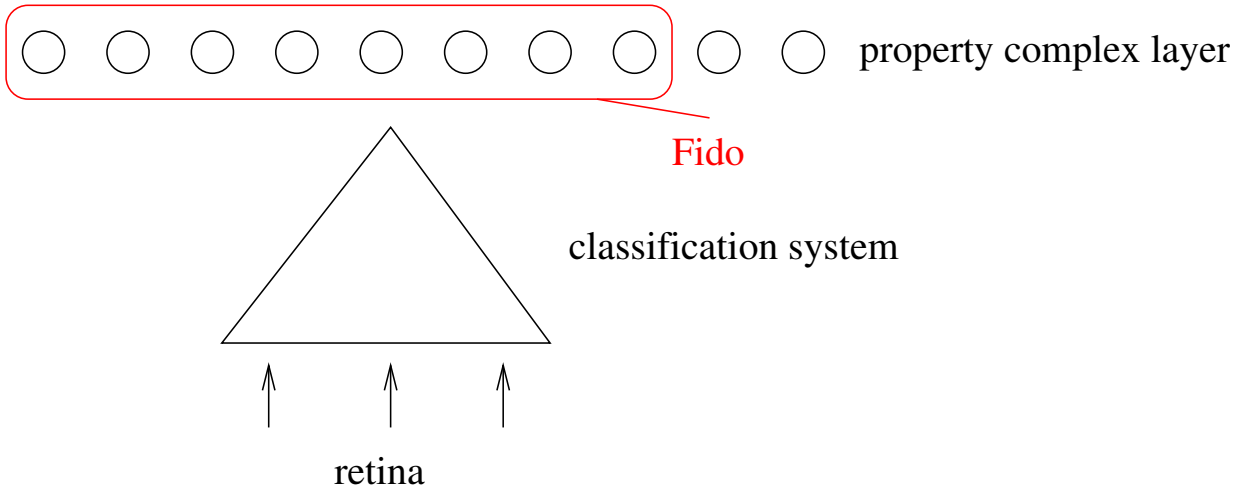

Competition between property assemblies

These properties are hierarchically organised into assemblies by Hebbian learning.

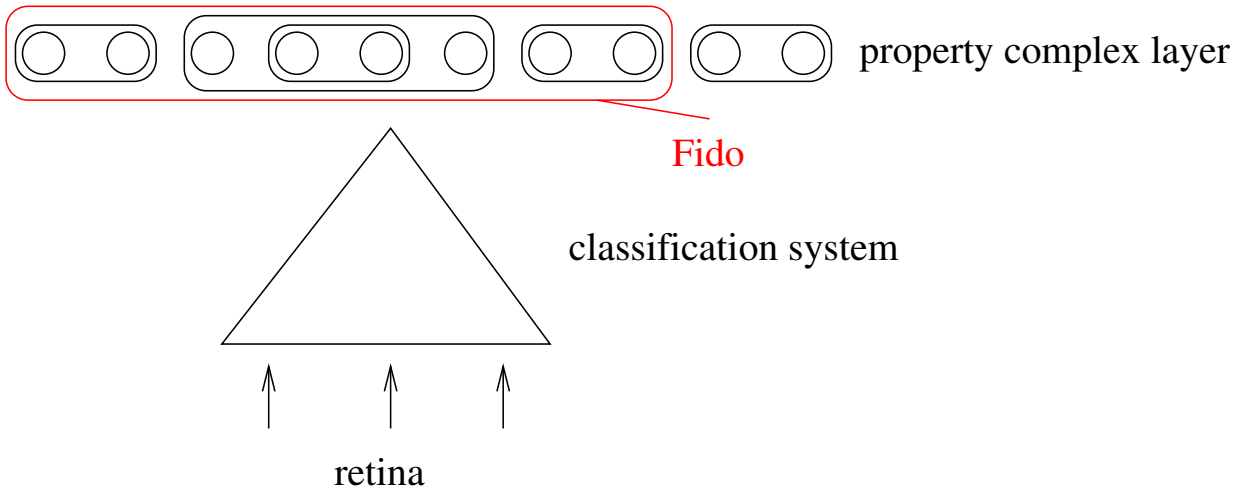

Competition between property assemblies

The property competition layer selects the most active assembly. This is the 'type' returned by the classification system.

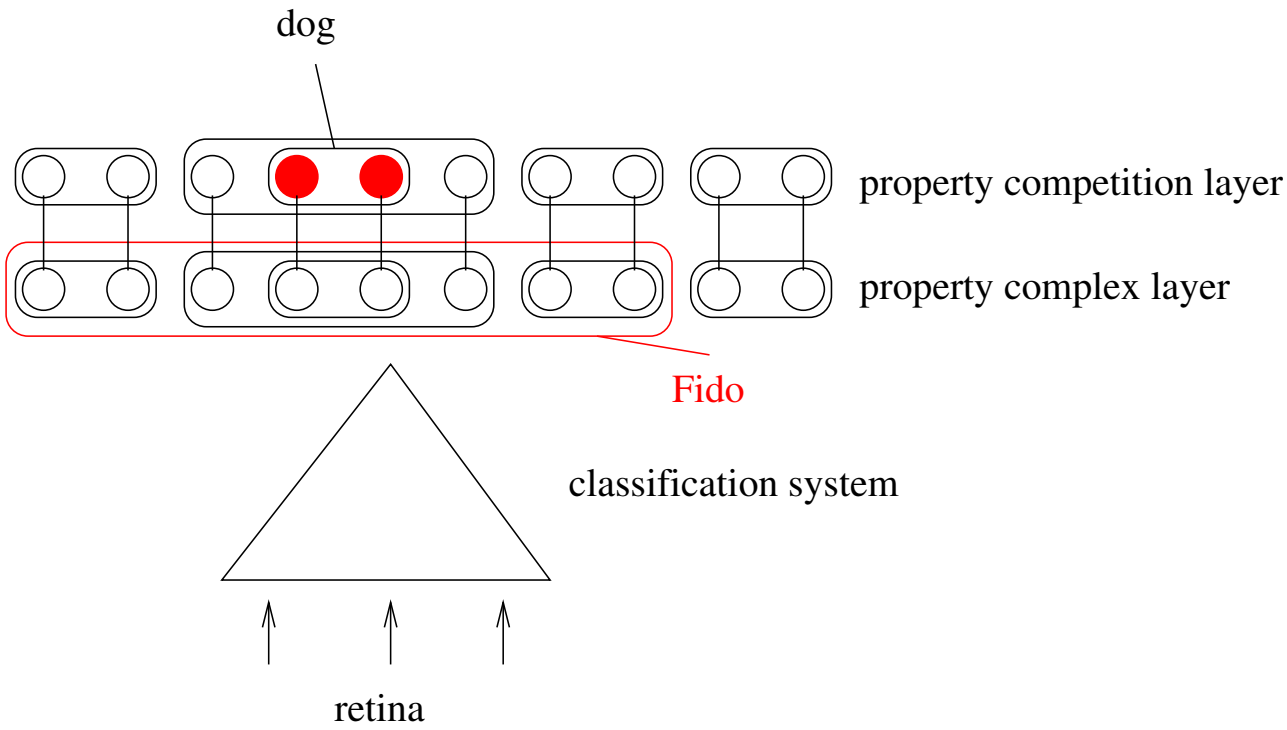

Property-level IOR

To 'attend to a property' of the currently established object, we can *inhibit* the most active assembly.

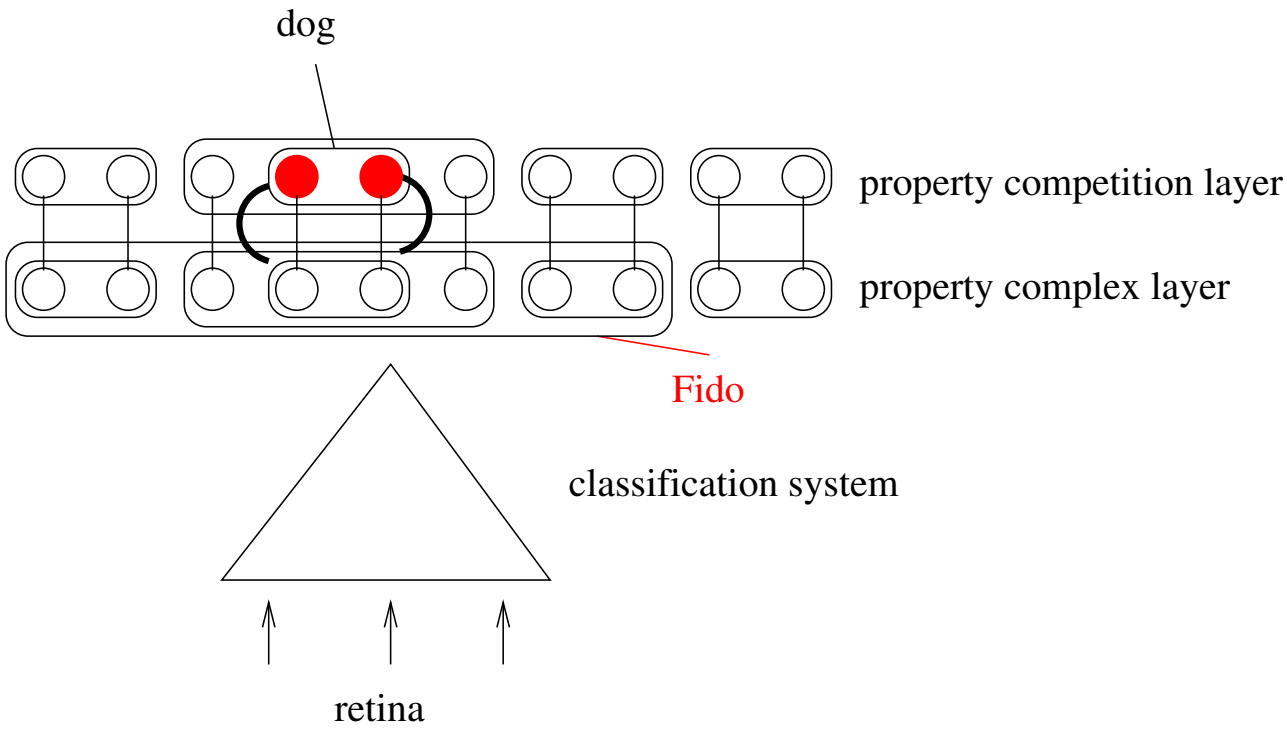

Property-level IOR

The assembly which dominates next reprepresents the property which most distinguishes it from the typical members of its type.

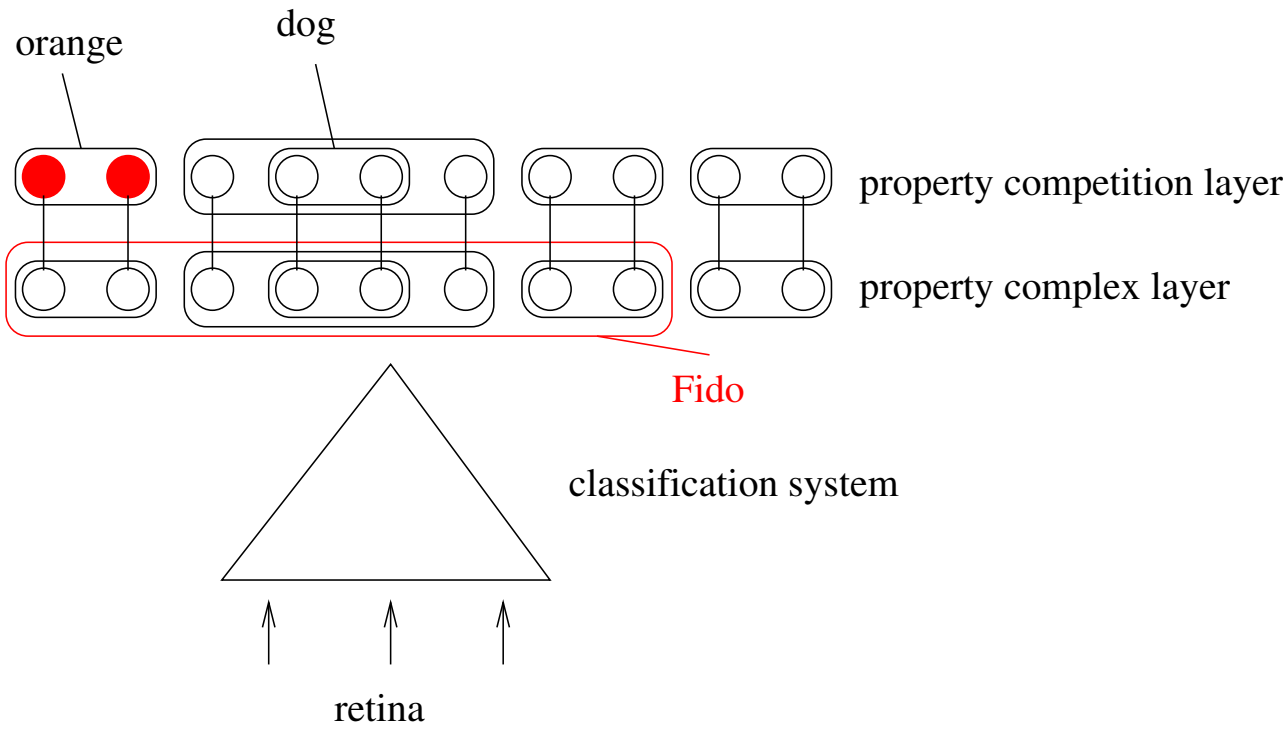

Property-level IOR

The assembly which appears after IOR can itself be a type.

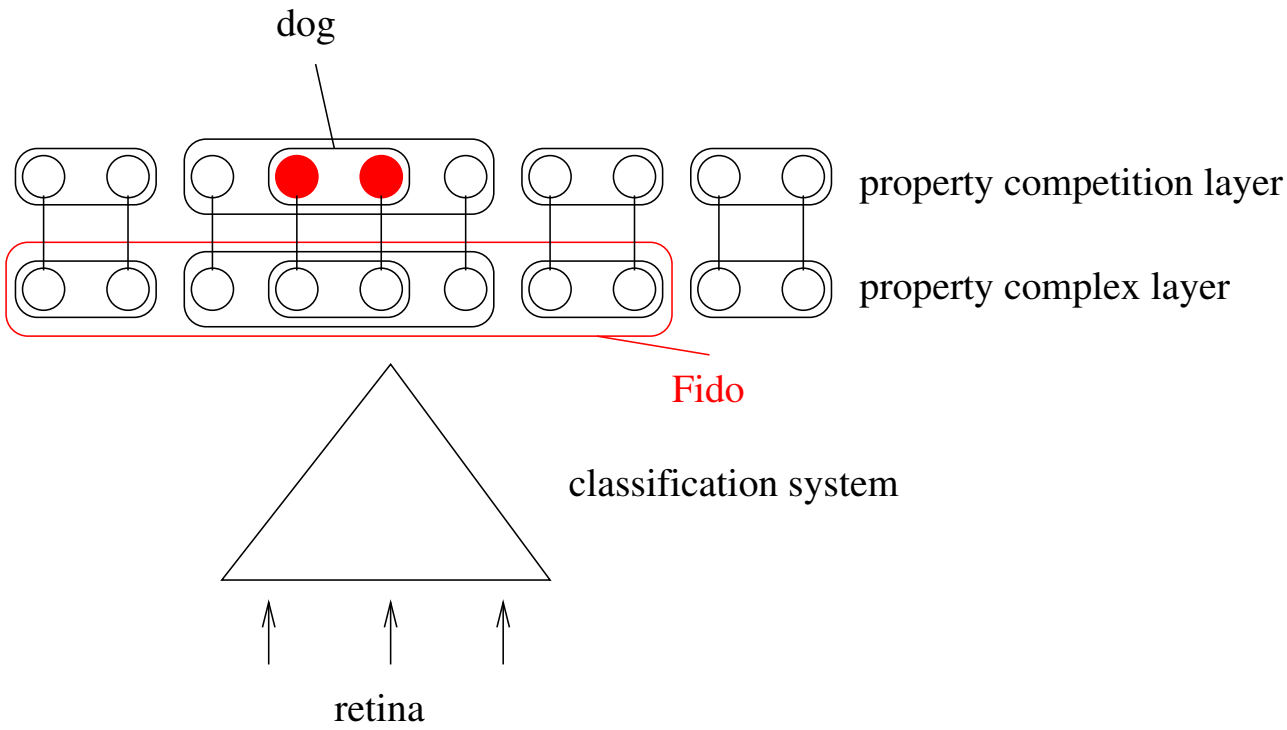

Property-level IOR

The assembly which appears after IOR can itself be a type.

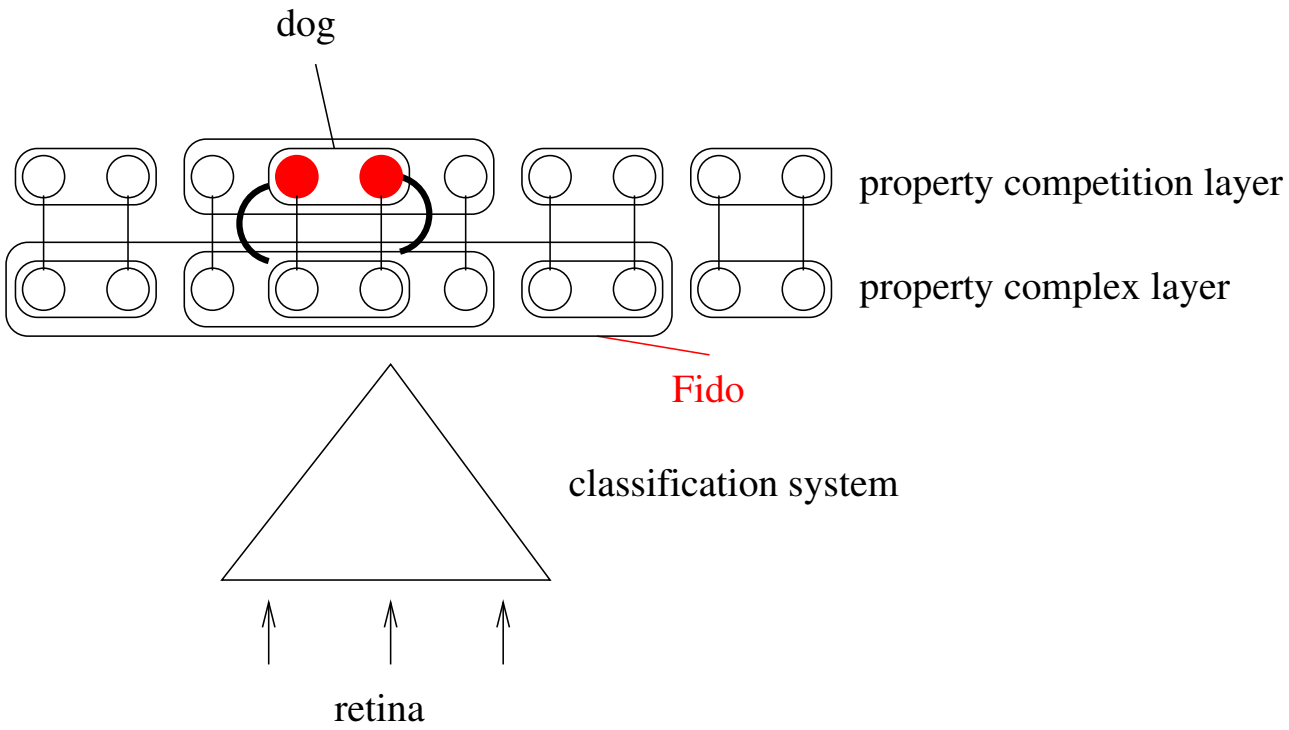

Property-level IOR

The assembly which appears after IOR can itself be a type.

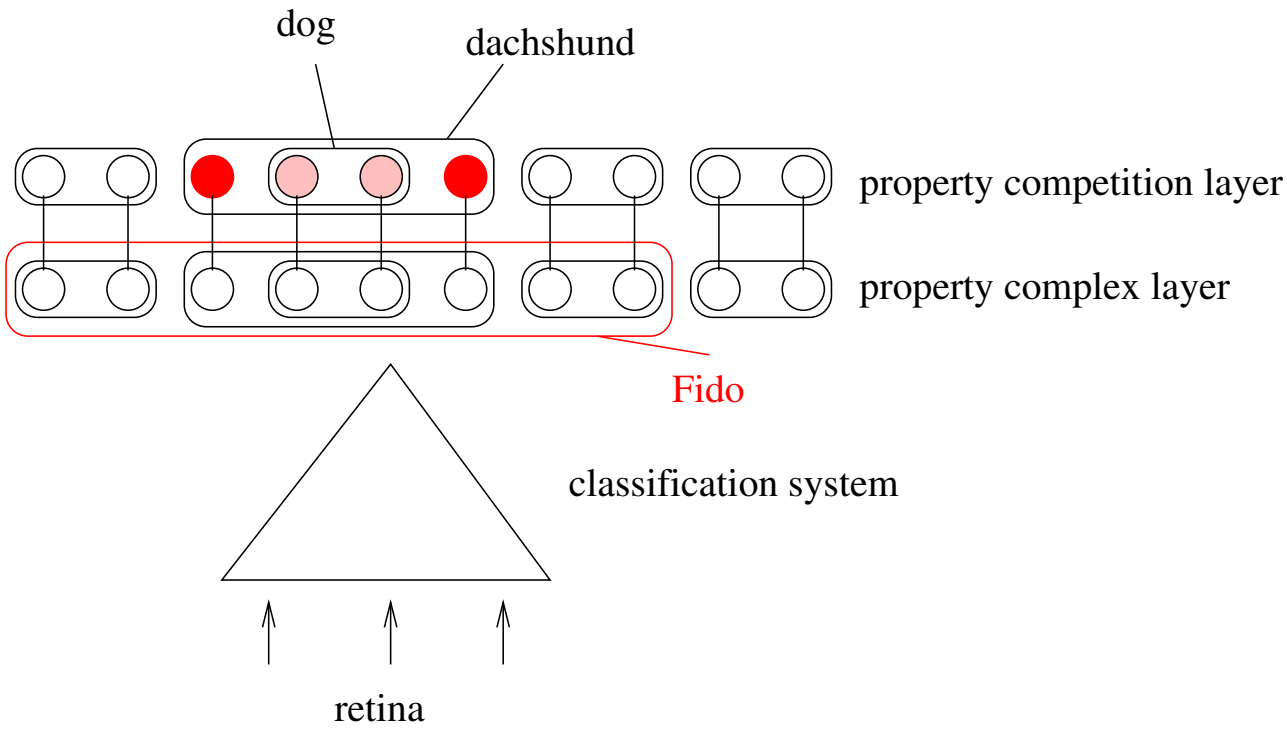

Properties and types

What's the difference between types (e.g. *dog*, *dachshund*) and regular properties (e.g. *orange*)?

Proposal:

- Types are assemblies which can be evoked in the property competition layer *when an object is first established*.
- Regular properties are assemblies which can only be evoked in the competition layer *after property-level IOR*.

Syntactic correlates of property-level IOR

Recall:

- DP signals selection of a location; NumP signals selection of a scale; NP signals activation of a 'category'.
- A predicate nominal has no DP layer.

Note that property-level IOR maintains the currently selected *location*.

- Proposal: a predicate nominal can convey the results of property-level IOR.

  The dog is a dachshund.

- Proposal: *be* can signal property-level IOR.

Existentials

Recall:

- An existential sentence can be analysed as predicating a property of ‘a point in space’.

There is a dog in the garden.

Proposal: we can represent the establishment of a new object within the clause system or within the DP system.

- Expletive ‘there’ signals the selection of a new region in the saliency map. (By itself.)
- The effects of this operation are similar to the effects of property-level IOR, and can be described in the clause system.

WM individuals and reattention

Each attended object is represented as a WM individual: a prepared attentional sequence.

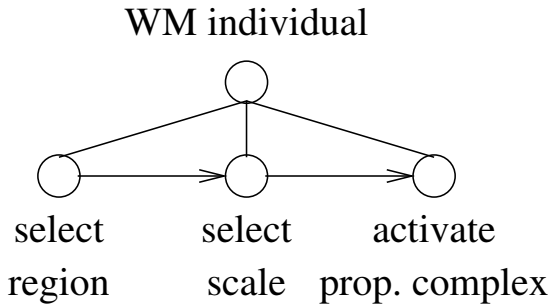

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

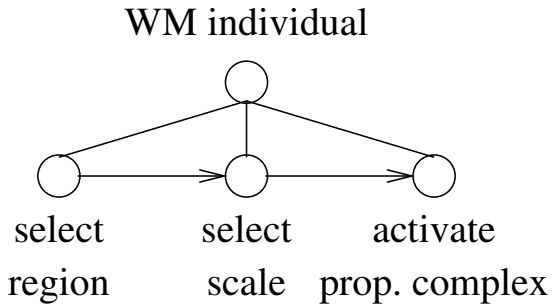

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

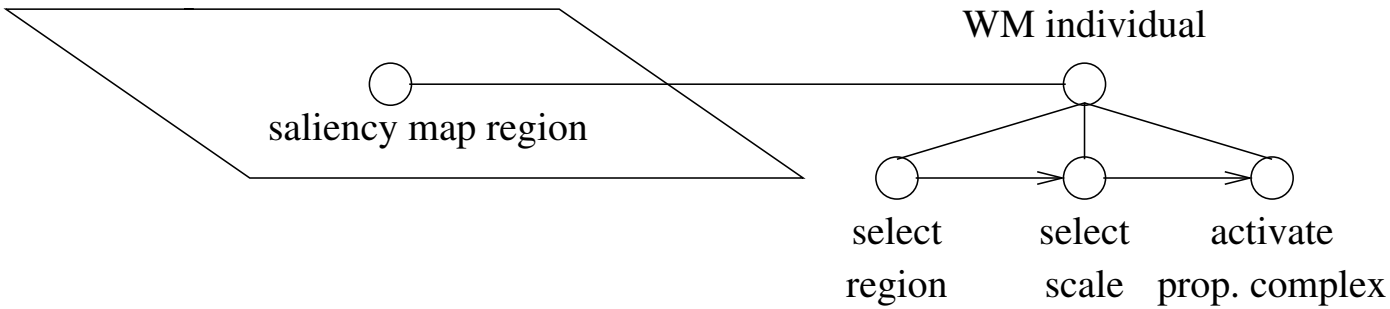

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

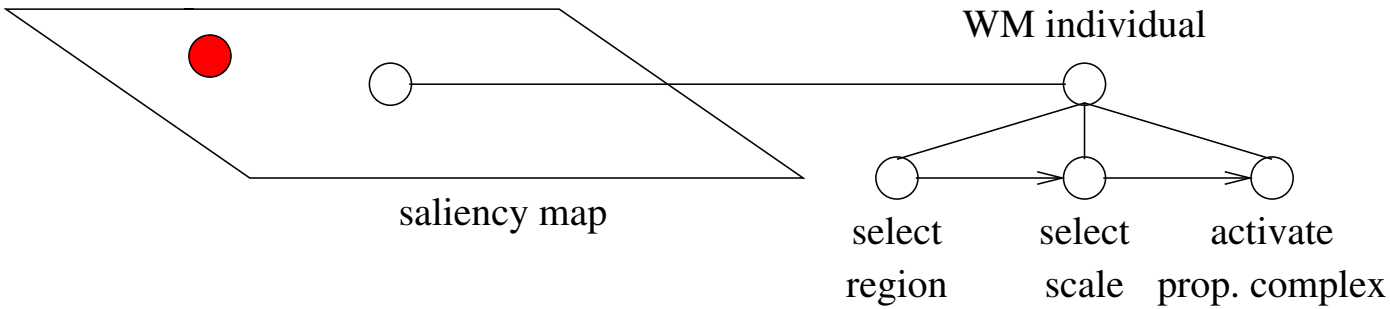

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

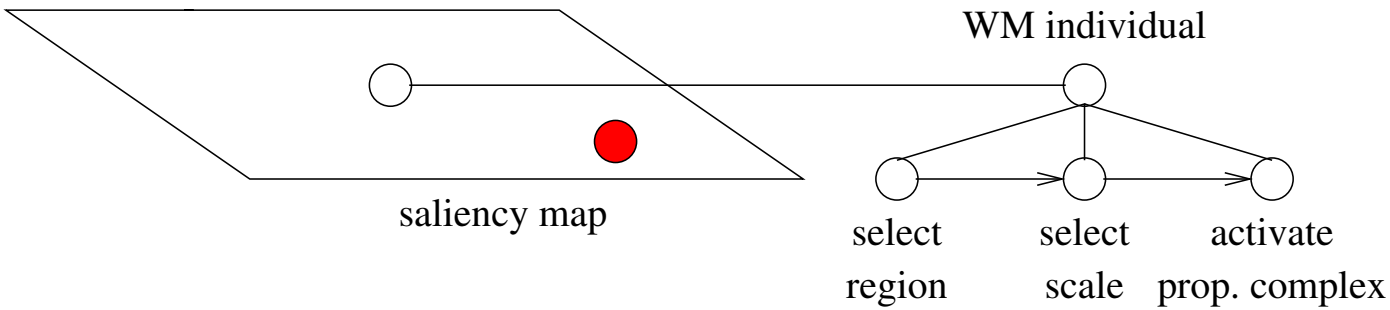

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

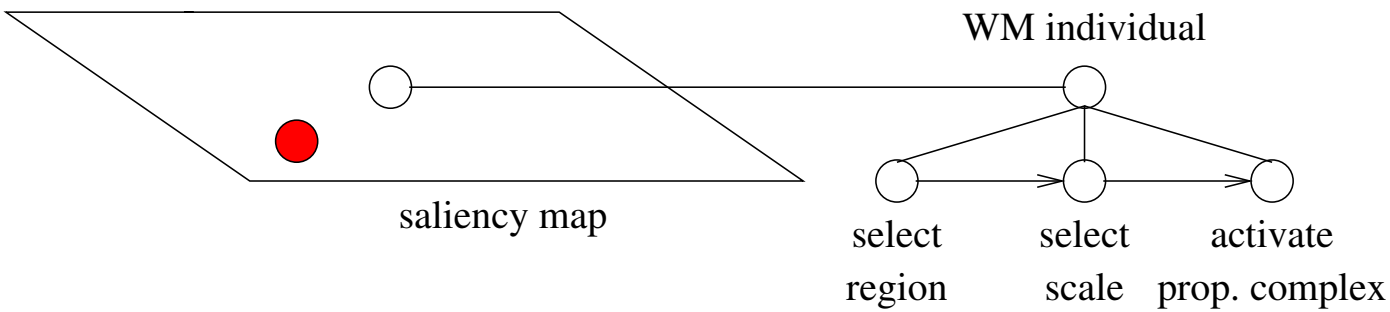

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

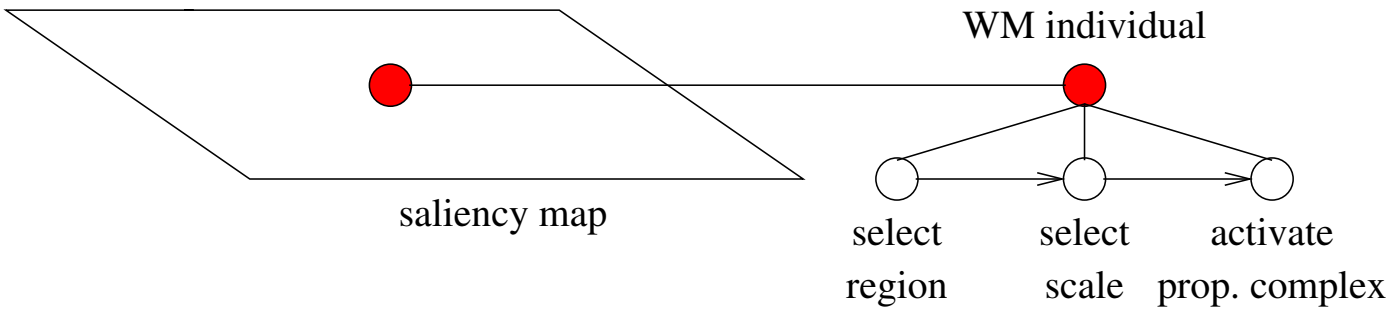

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

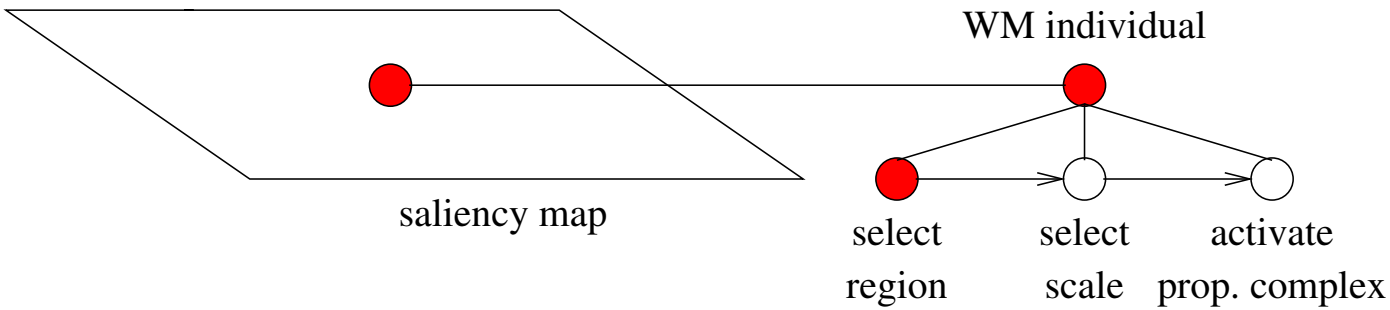

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

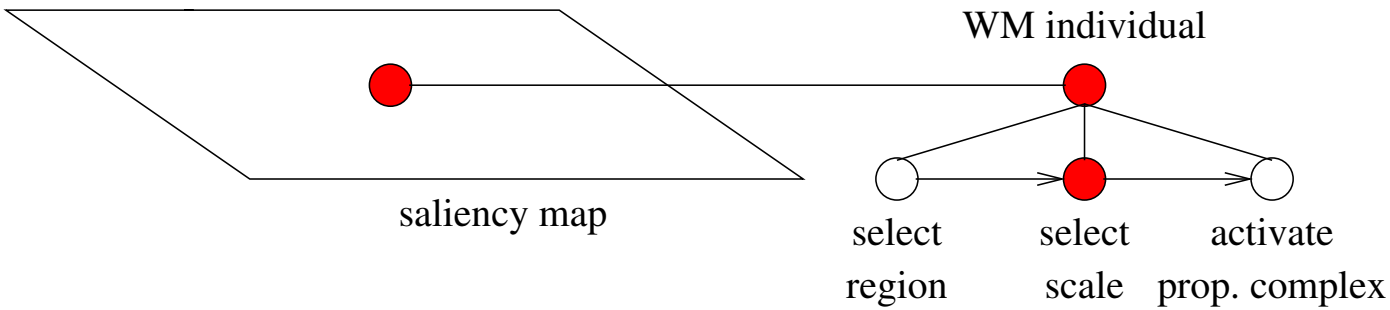

WM individuals and reattention

Proposal: the purpose of WM individuals is to aid *reattention* to objects.

- Each WM individual is linked to a saliency map region.
- If our attention is drawn back to this region, we know what to expect.

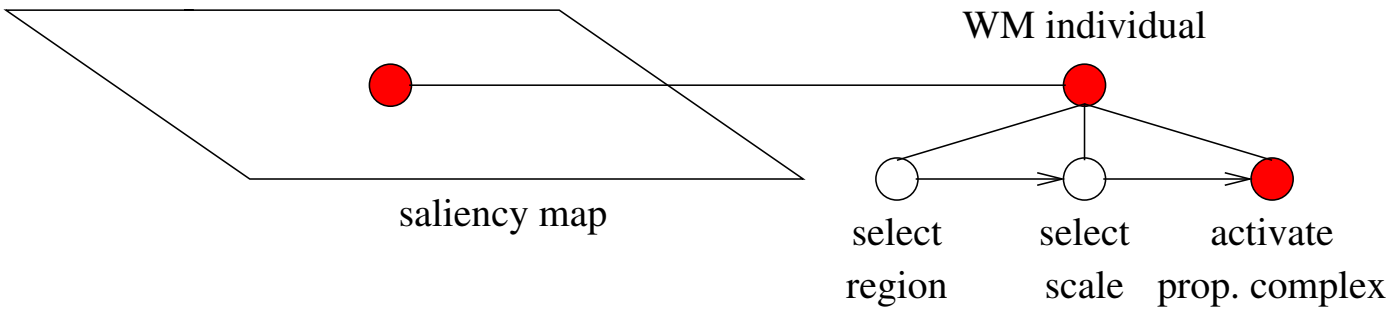

LTM individuals

Each WM individual is also linked to a LTM individual.

- An LTM individual represents an object known to the observer.
- It activates when an observed object is *recognised*.
- LTM individuals are in MT (e.g. Eichenbaum *et al.*, 2007)

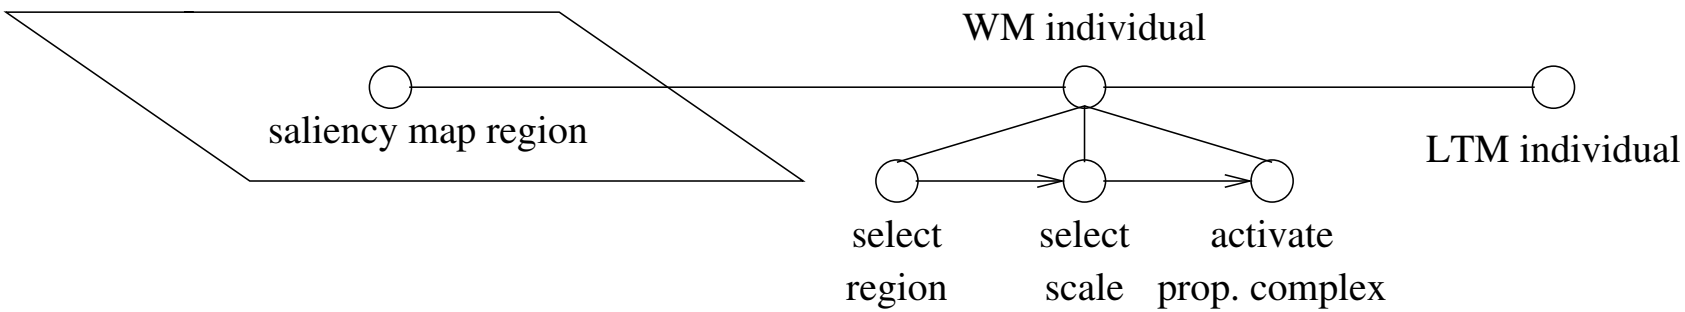

WM individuals and definiteness

Recall:

- An indefinite DP introduces a new referent.
- A definite DP presupposes an existing referent.

WM individuals are well suited to function as referents.

Proposal:

- An indefinite DP signals the creation of a new WM individual.
- A definite DP signals reactivation of an existing WM individual.

The position of definite DPs

Recall:

- The definite determiner *the* is associated with D.
- D denotes a particular saliency map region.

If *the* signals reactivation of a WM individual, we can explain why it sits at D:

- Reactivation of a WM individual is *triggered* by attention to a saliency map region.

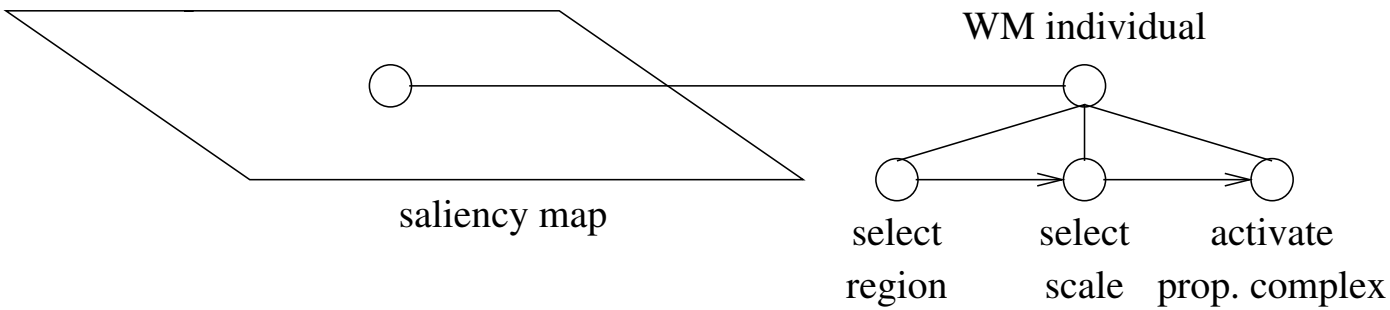

The position of definite DPs

Recall:

- The definite determiner *the* is associated with D.
- D denotes a particular saliency map region.

If *the* signals reactivation of a WM individual, we can explain why it sits at D:

- Reactivation of a WM individual is *triggered* by attention to a saliency map region.

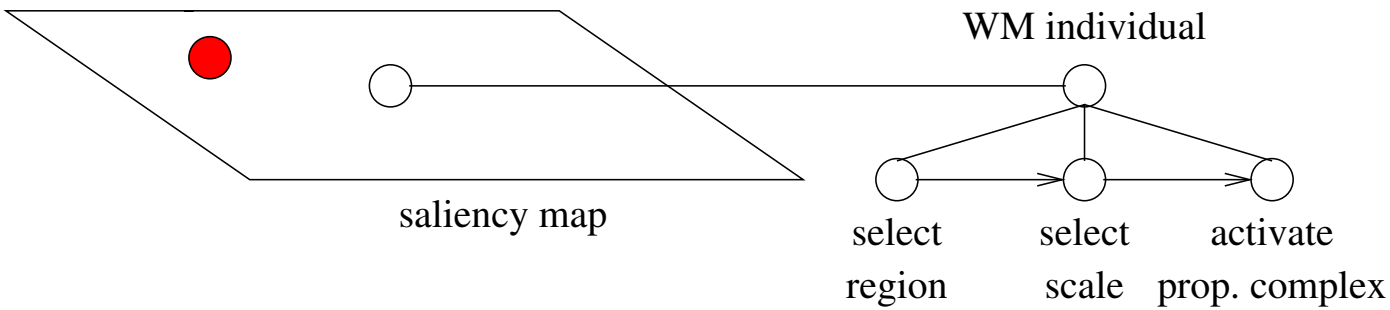

The position of definite DPs

Recall:

- The definite determiner *the* is associated with D.
- D denotes a particular saliency map region.

If *the* signals reactivation of a WM individual, we can explain why it sits at D:

- Reactivation of a WM individual is *triggered* by attention to a saliency map region.

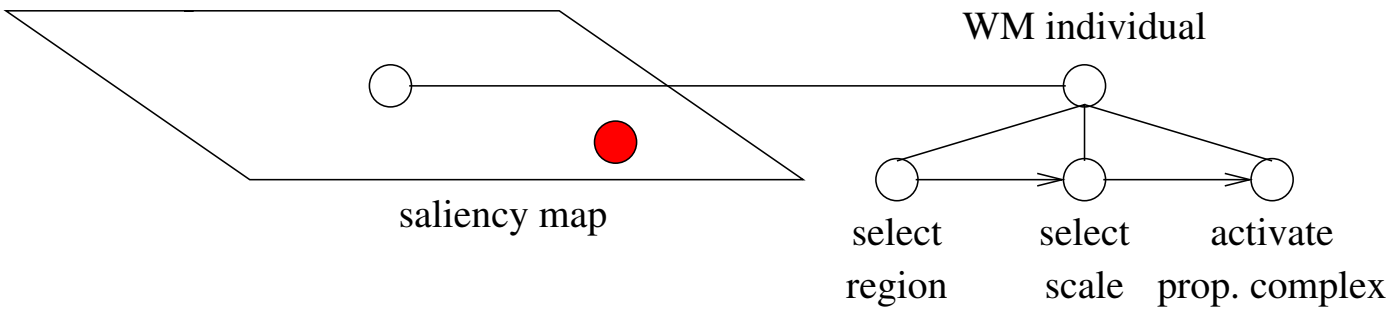

The position of definite DPs

Recall:

- The definite determiner *the* is associated with D.
- D denotes a particular saliency map region.

If *the* signals reactivation of a WM individual, we can explain why it sits at D:

- Reactivation of a WM individual is *triggered* by attention to a saliency map region.

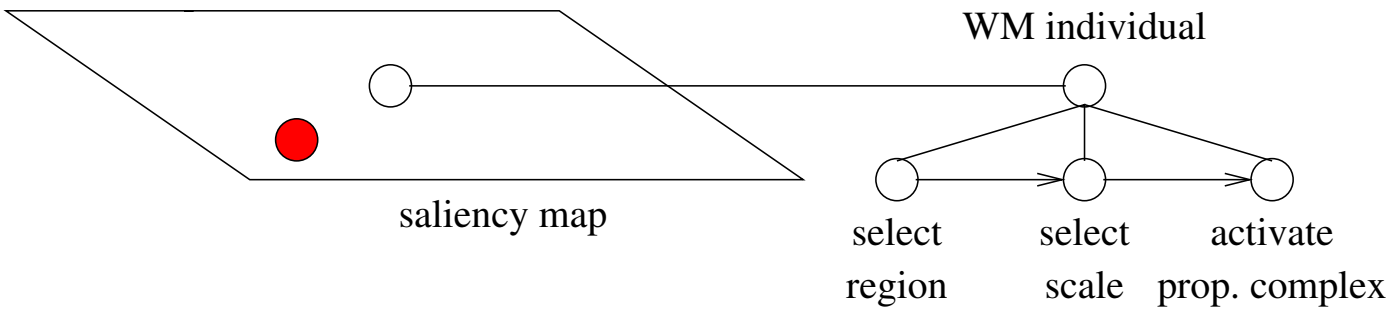

The position of definite DPs

Recall:

- The definite determiner *the* is associated with D.
- D denotes a particular saliency map region.

If *the* signals reactivation of a WM individual, we can explain why it sits at D:

- Reactivation of a WM individual is *triggered* by attention to a saliency map region.

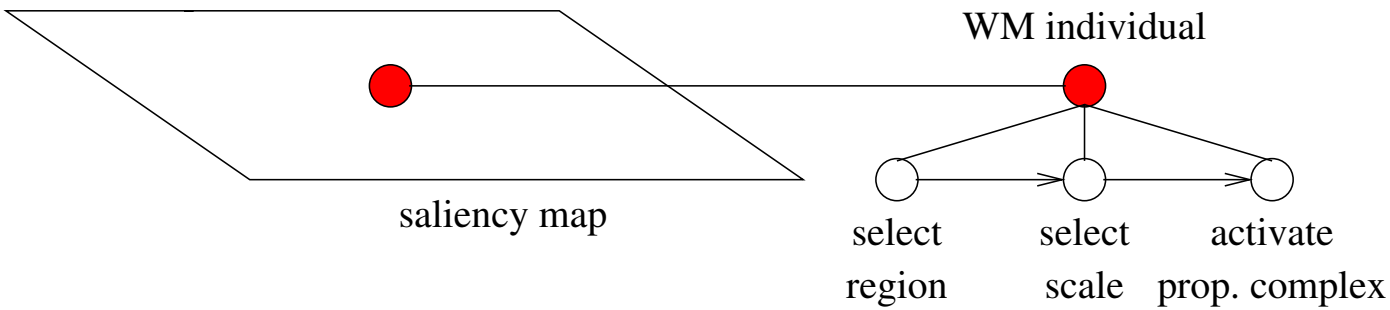

Indefinite determiners

Recall: Indefinite determiners are associated with Num, not D.

However, they can also introduce *new referents*.

A cat walked in. *John grabbed a cup.*

Indefinite determiners

Recall: Indefinite determiners are associated with Num, not D.

However, they can also introduce *new referents*.

A cat walked in. *John grabbed a cup.*

Proposal: a 'referential' indefinite DP describes the creation of a new WM individual.

saliency map

Indefinite determiners

Recall: Indefinite determiners are associated with Num, not D.

However, they can also introduce *new referents*.

A cat walked in. *John grabbed a cup.*

Proposal: a 'referential' indefinite DP describes the creation of a new WM individual.

- First, a 'new' saliency map region is attended to.

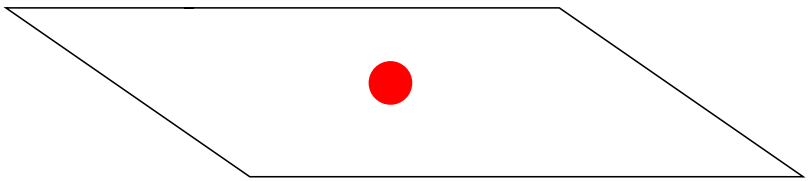

saliency map

Indefinite determiners

Recall: Indefinite determiners are associated with Num, not D.

However, they can also introduce *new referents*.

A cat walked in. *John grabbed a cup.*

Proposal: a 'referential' indefinite DP describes the creation of a new WM individual.

- First, a 'new' saliency map region is attended to.
- Then a WM individual is created.

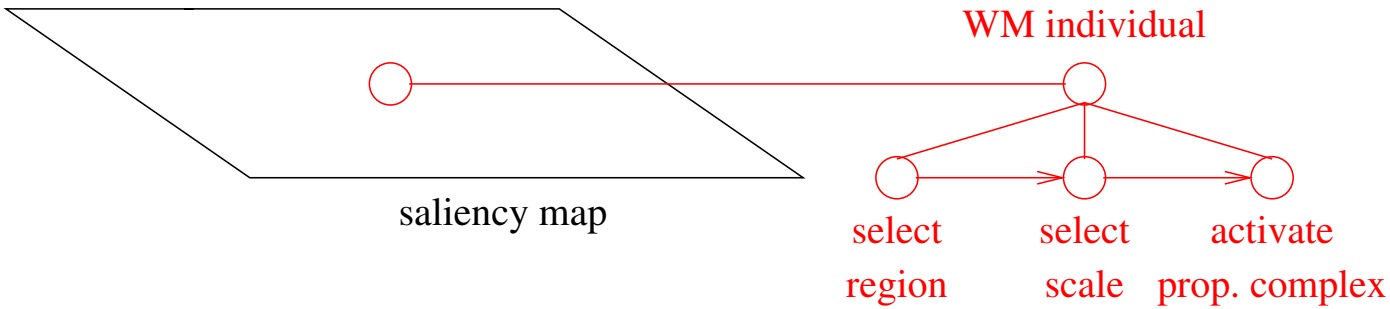

Definiteness and attention to properties

The subject of a sentence predicating an ‘individual-level’ property of an object can only be definite.

The dog was brown

Definiteness and attention to properties

The subject of a sentence predicating an 'individual-level' property of an object can only be definite.

? A dog was brown

If it's indefinite, the sentence is interpreted generically.

Definiteness and attention to properties

The subject of a sentence predicating an 'individual-level' property of an object can only be definite.

The dog was brown

Note:

- Individual-level properties (e.g. brown) don't capture attention: they're established by property-level IOR.
- So visually establishing a property can only happen by reattending to a known object.

Definiteness and existentials

An existential sentence can't introduce a definite DP.

There is a dog outside. **There is the dog outside.*

Definiteness and existentials

An existential sentence can't introduce a definite DP.

There is a dog outside. **There is the dog outside.*

Note:

- *There* signals attention to a point in space.
- If this point were already associated with a WM individual, this operation would evoke a full (definite) DP.
- So it must be a 'new' location.

Quantified sentences

E.g. *Many dogs were brown*

Very briefly:

- Semantic memory is modelled as a set of links between LTM individuals and property complexes.
- A category can be used to query semantic memory, and activate all the LTM individuals (in a certain context) which are of this category.
- The cardinality/numerosity of the set of individuals is recorded.
- The activated individuals all activate their associated property complexes.
- Property-level IOR can be performed on the resulting complex, picking out a subset of individuals.

Quantification over episodes

E.g. *Every man grabbed a cup*

Very briefly:

- When we observe an episode, we create a WM episode to represent it.
- We can turn WM episodes into properties, which can be associated with individuals in semantic memory.
- Assume that WM episodes can refer to individuals as WM individuals.
- When an episode features in a generalisation in semantic memory, any references to WM individuals function like references to bound variables.

Summary

- A DP has a right-branching syntactic structure.
- Establishing an object or group has a sequential structure.
- A DP can be understood as describing a *replayed* attentional sequence.
- This interpretation suggests there is a perceptual basis for the syntactic concept of Number.
- The attentional model can be extended to allow perceptual interpretations of predicate nominals, existentials, definiteness, and the interpretation of DPs in quantified sentences.

A model of saliency featuring homogeneity

Saliency is typically associated with *local contrast*.

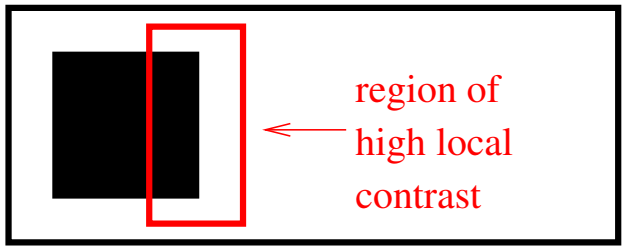

A model of saliency featuring homogeneity

Saliency is typically associated with *local contrast*.

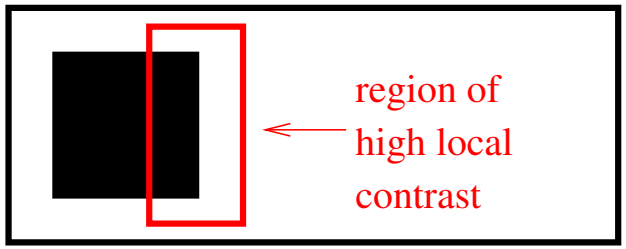

Prima facie, homogeneity and saliency are conflicting cues.

A model of saliency featuring homogeneity

Saliency is typically associated with *local contrast*.

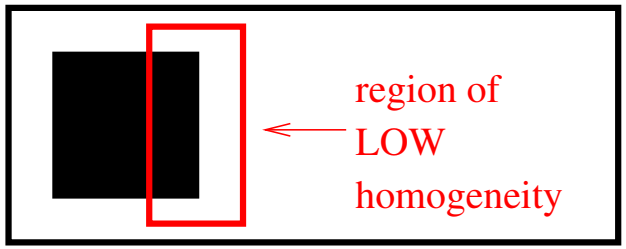

Prima facie, homogeneity and saliency are conflicting cues.

A model of saliency featuring homogeneity

Saliency is typically associated with *local contrast*.

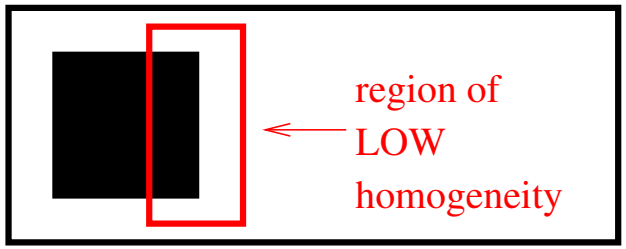

Prima facie, homogeneity and saliency are conflicting cues.

Solution: homogeneity should apply at a different spatial scale.

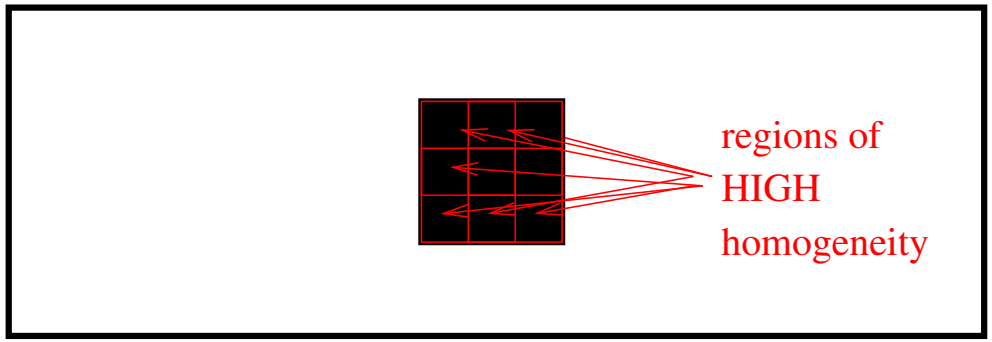

A model of saliency featuring homogeneity

Saliency is typically associated with *local contrast*.

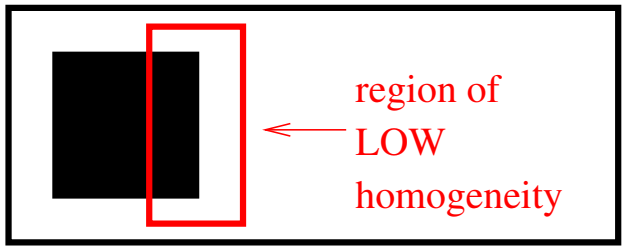

Prima facie, homogeneity and saliency are conflicting cues.

Solution: homogeneity should apply at a different spatial scale.

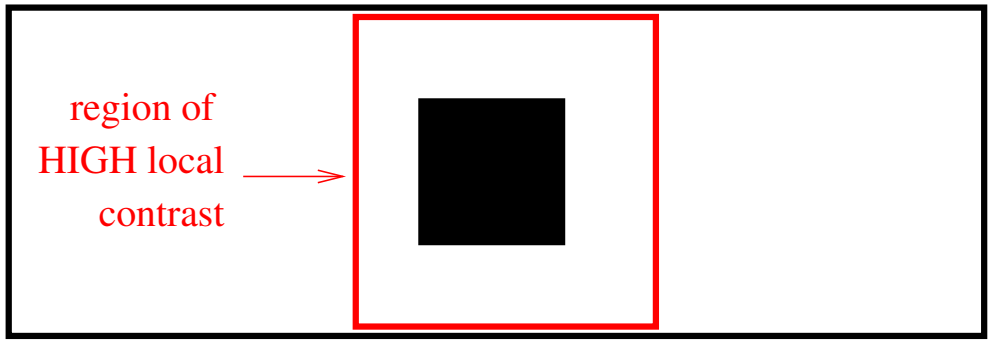

A model of saliency featuring homogeneity

Saliency is typically associated with *local contrast*.

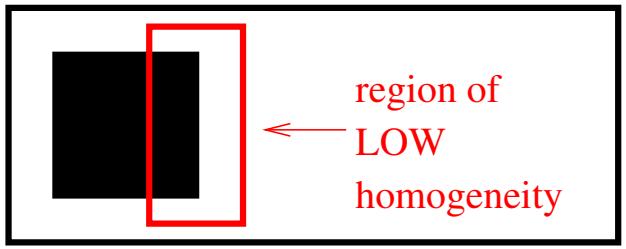

Prima facie, homogeneity and saliency are conflicting cues.

Solution: homogeneity should apply at a different spatial scale.

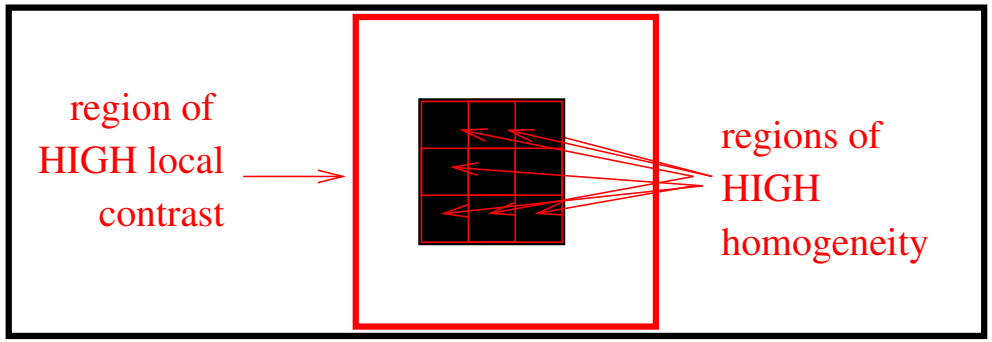

Examples from the saliency map

Some screenshots, showing the tradeoff between proximity and homogeneity.

If stimuli are close enough, they are grouped, even if they are dissimilar.

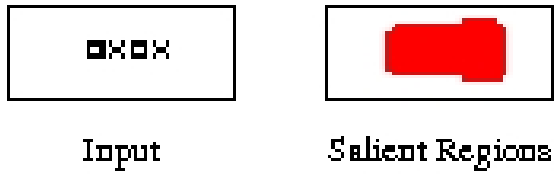

Examples from the saliency map

Some screenshots, showing the tradeoff between proximity and homogeneity.

At an intermediate separation, homogeneity influences grouping.

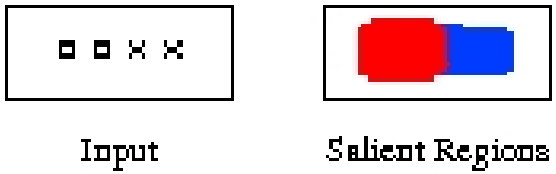

Examples from the saliency map

Some screenshots, showing the tradeoff between proximity and homogeneity.

At an intermediate separation, homogeneity influences grouping.

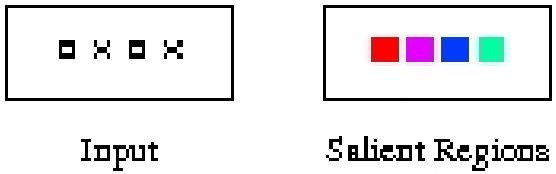

Examples from the saliency map

Some screenshots, showing the tradeoff between proximity and homogeneity.

If stimuli are far enough apart they are separate regions, even if they are homogeneous.

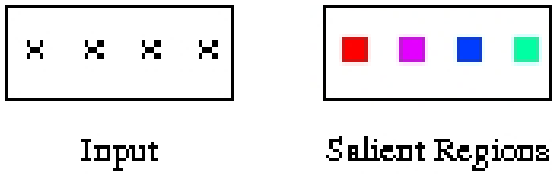

Visual search and stimulus similarity

Duncan and Humphreys (1989): visual search for a target is fast if

- the distractors are similar to one another; and
- the target is dissimilar to the distractors.

Distractors similar, target different from distractors:

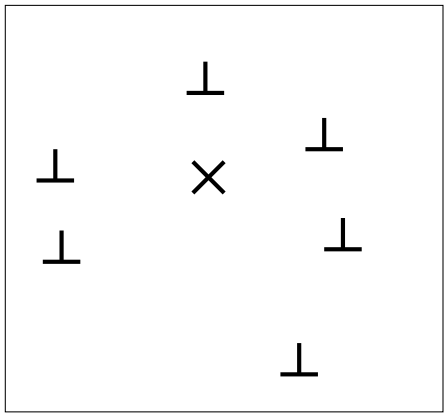

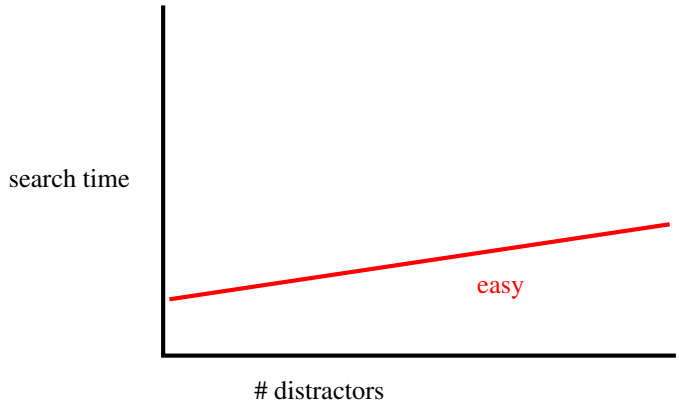

Visual search and stimulus similarity

Duncan and Humphreys (1989): visual search for a target is fast if

- the distractors are similar to one another; and
- the target is dissimilar to the distractors.

Distractors different, target different from distractors:

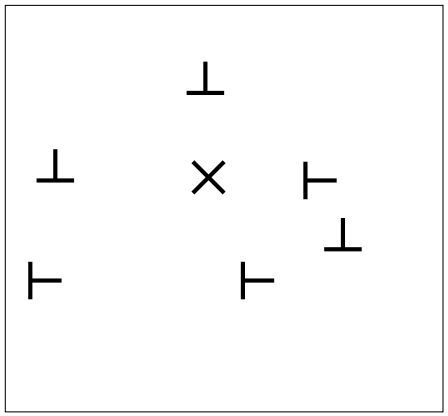

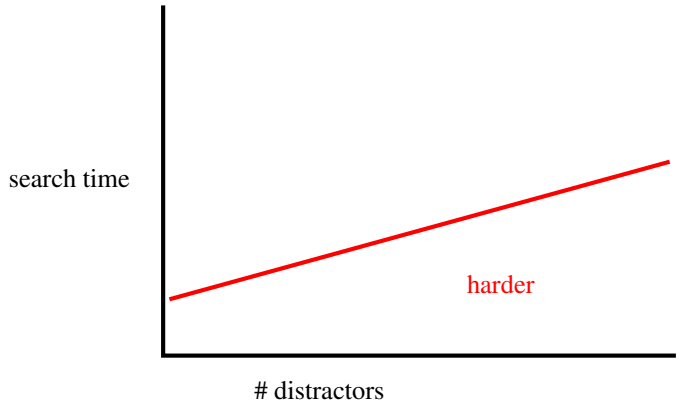

Visual search and stimulus similarity

Duncan and Humphreys (1989): visual search for a target is fast if

- the distractors are similar to one another; and
- the target is dissimilar to the distractors.

Distractors similar, target similar to distractors:

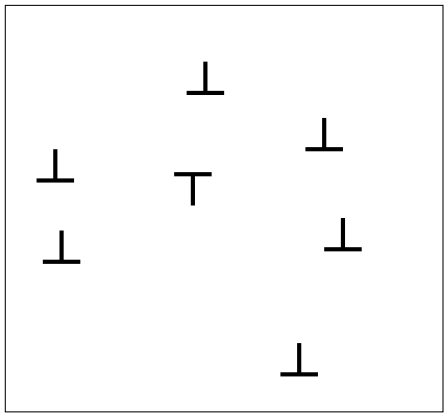

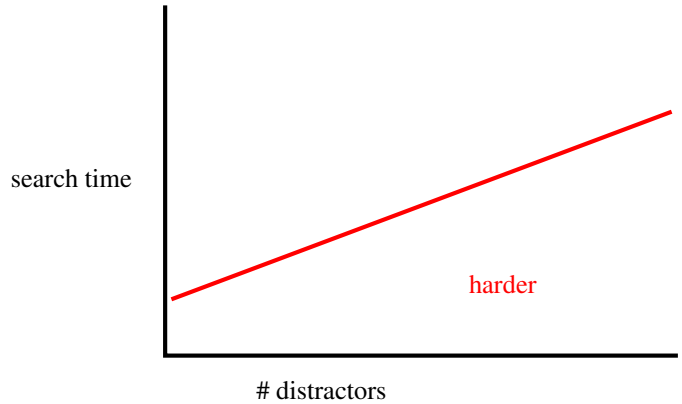

Visual search and stimulus similarity

Duncan and Humphreys (1989): visual search for a target is fast if

- the distractors are similar to one another; and
- the target is dissimilar to the distractors.

Distractors different, target similar to distractors:

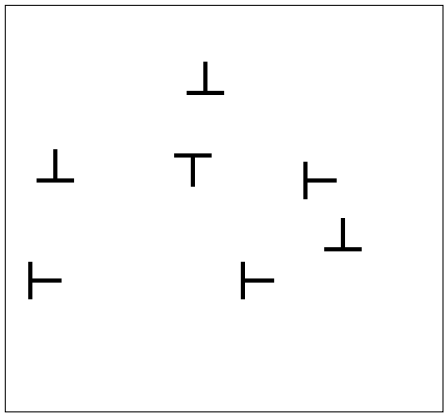

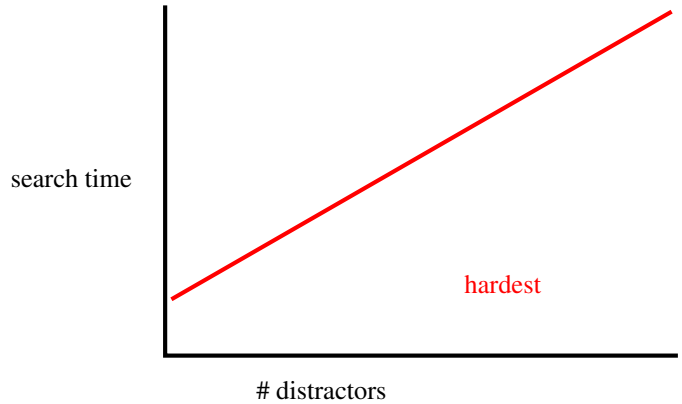

Visual search and stimulus similarity

Duncan and Humphreys (1989): visual search for a target is fast if

- the distractors are similar to one another; and
- the target is dissimilar to the distractors.

Distractors different, target similar to distractors:

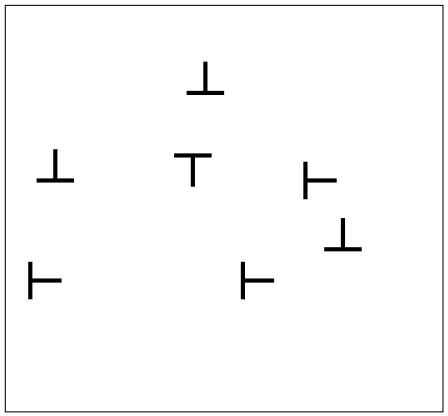

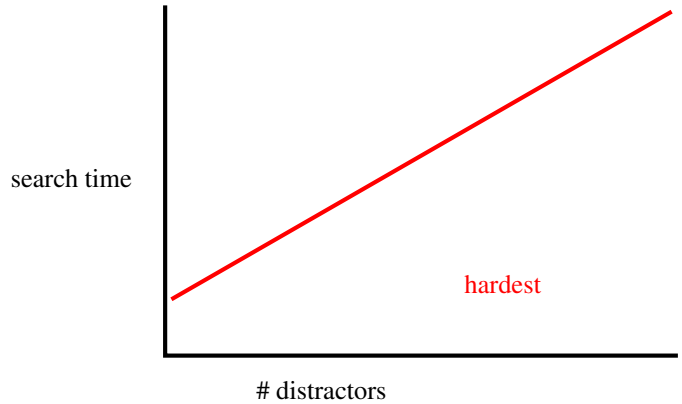

Our model of group attention/categorisation can model this effect.

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

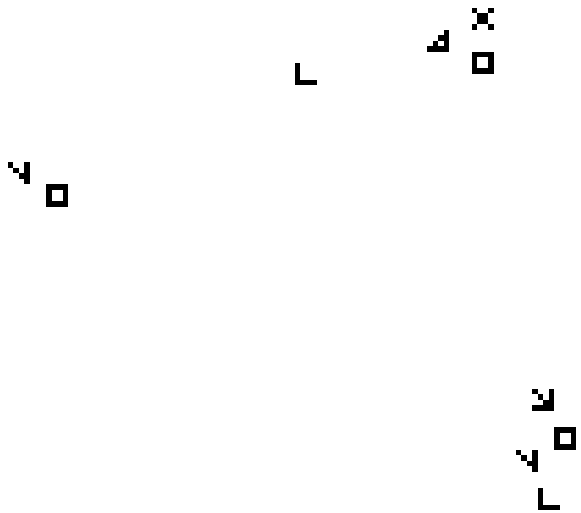

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

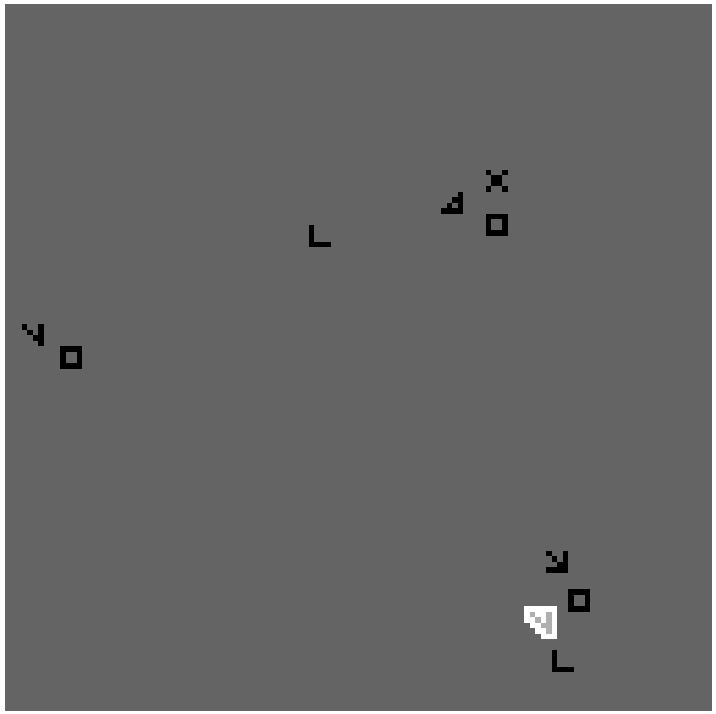

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

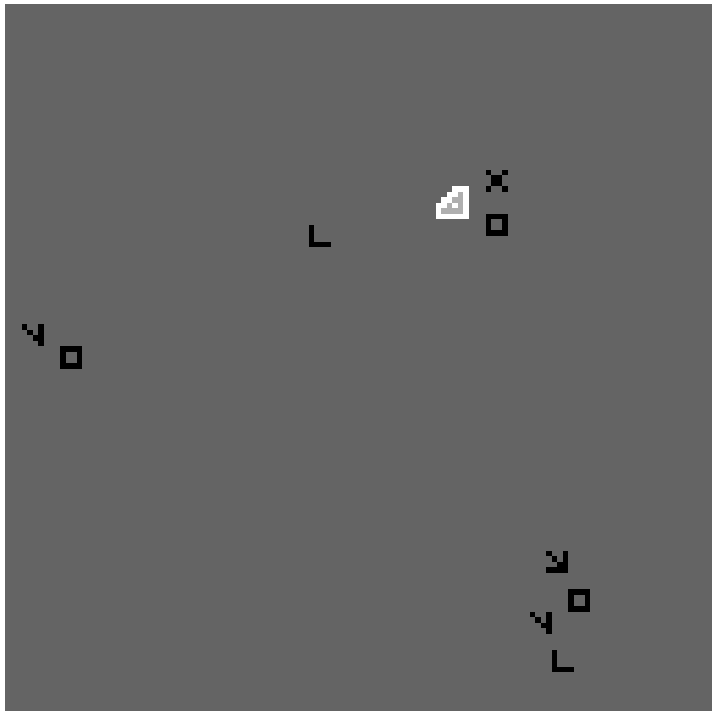

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

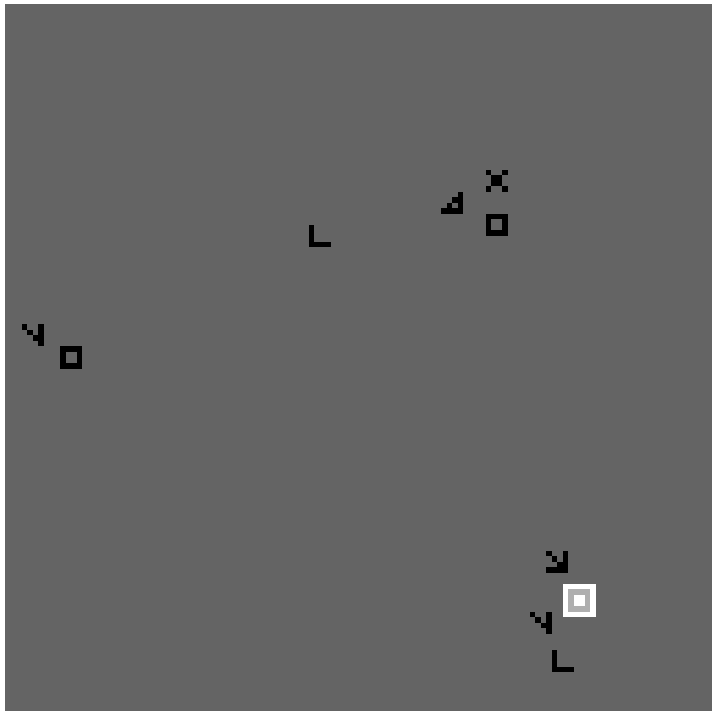

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

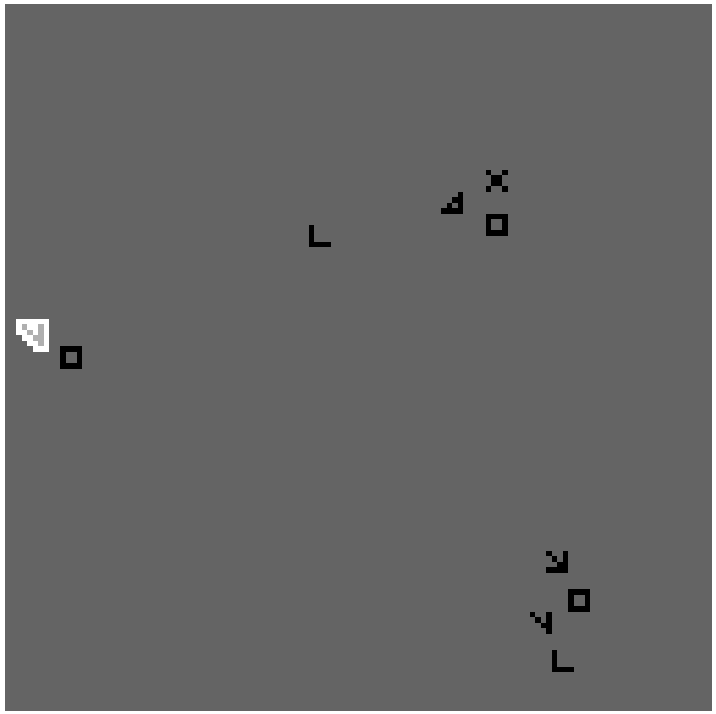

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

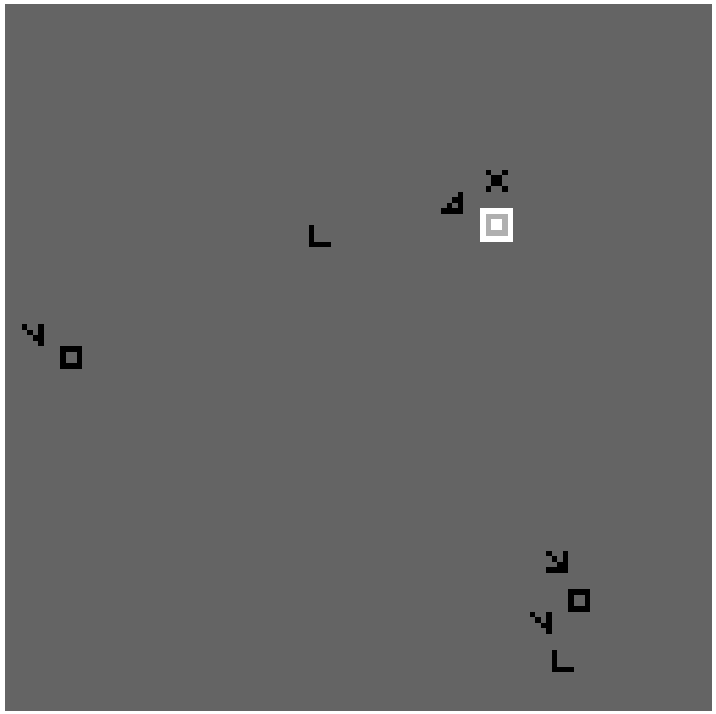

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

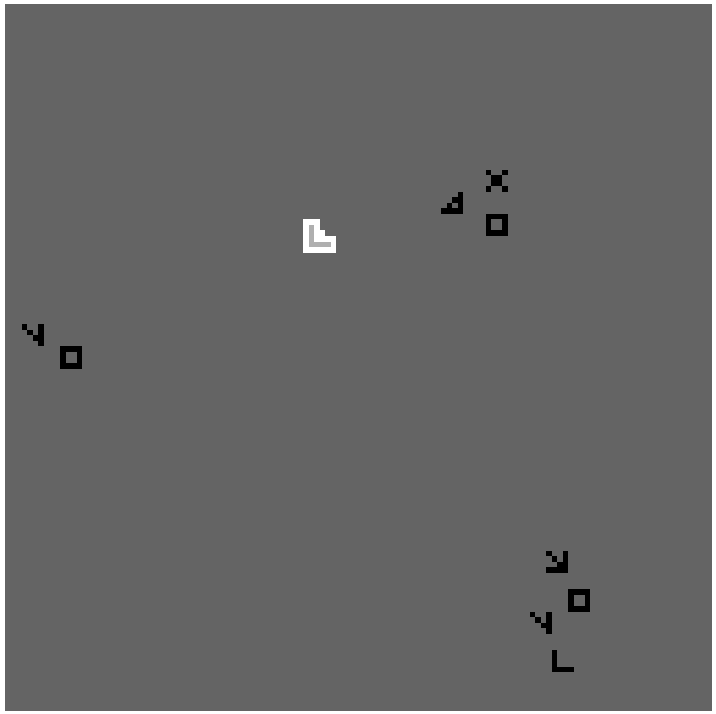

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

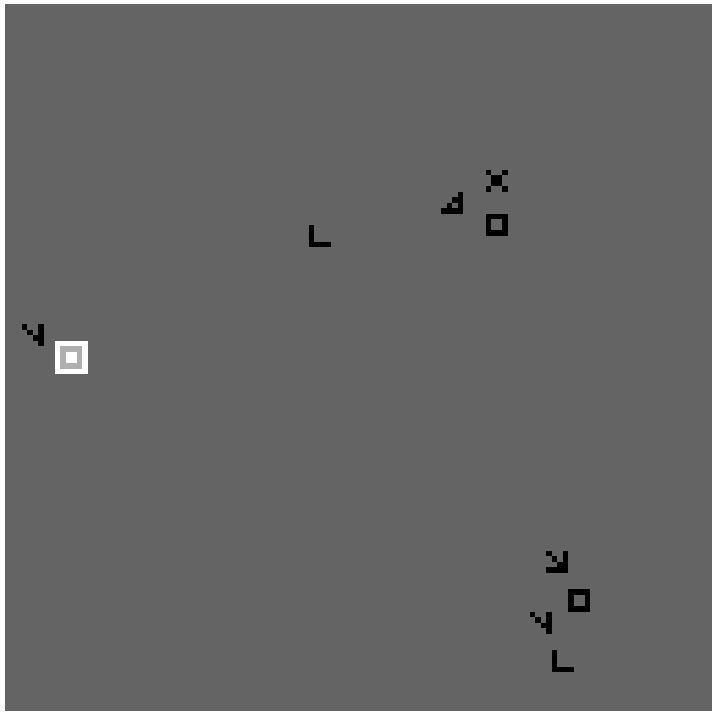

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

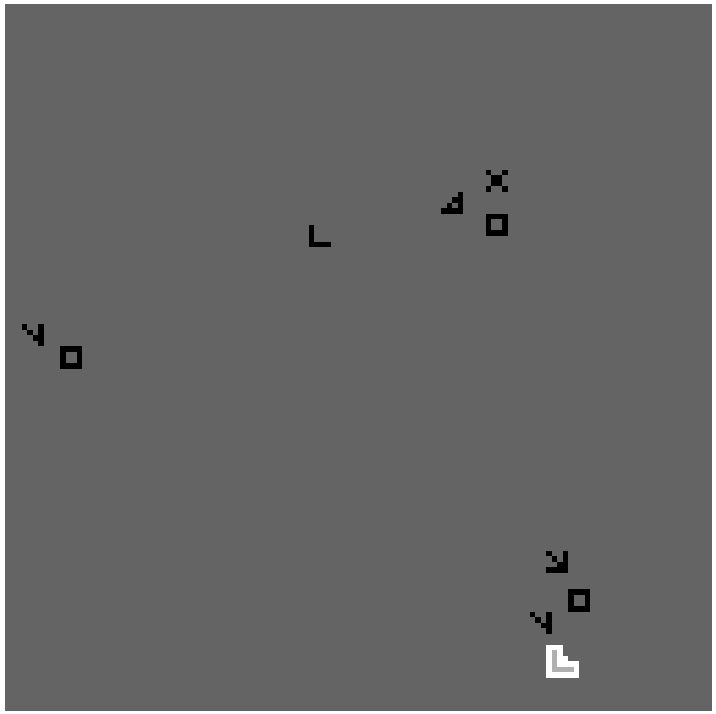

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

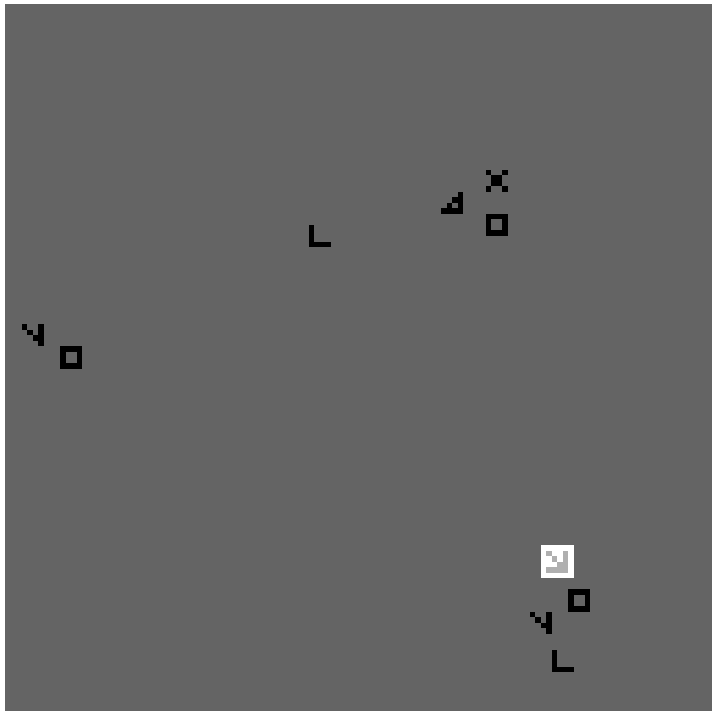

Dissimilar distractors

- Distractors can't be grouped (so search is linear).

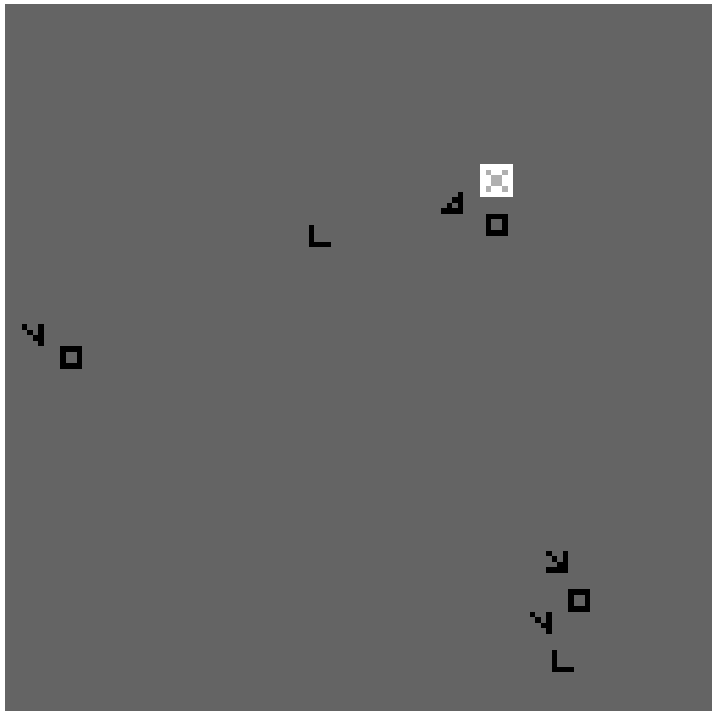

Distractors similar; target different from distractors

- Distractors can be grouped (which speeds up search).
- Targets aren't grouped with distractors.

Distractors similar; target different from distractors

- Distractors can be grouped (which speeds up search).
- Targets aren't grouped with distractors.

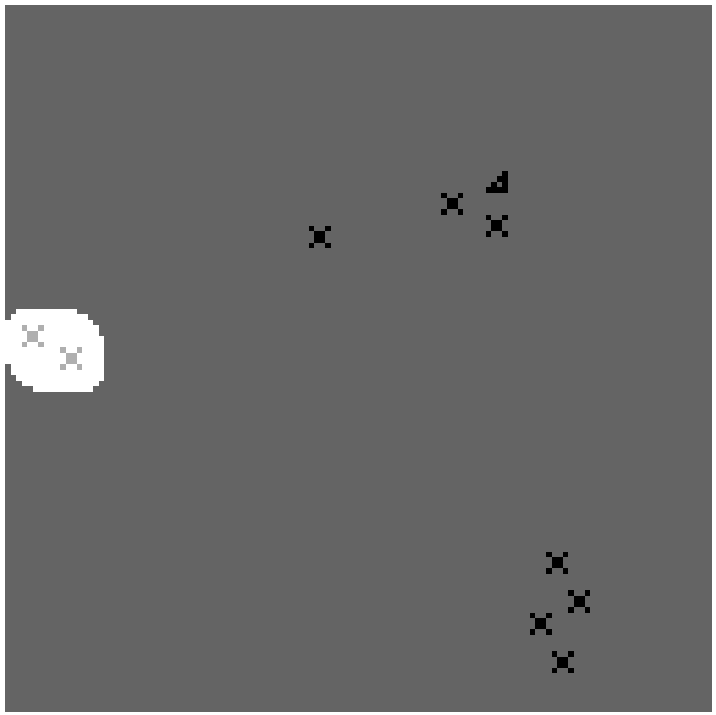

Distractors similar; target different from distractors

- Distractors can be grouped (which speeds up search).
- Targets aren't grouped with distractors.

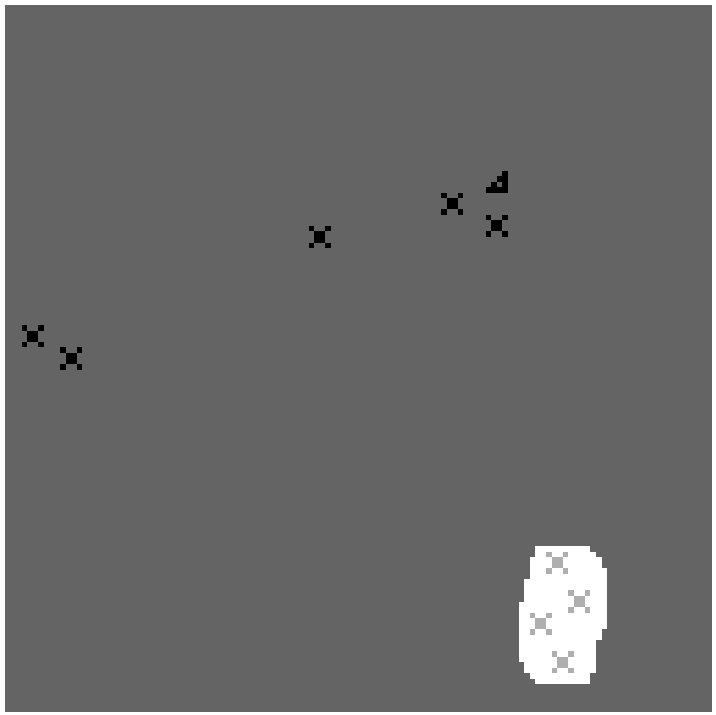

Distractors similar; target different from distractors

- Distractors can be grouped (which speeds up search).
- Targets aren't grouped with distractors.

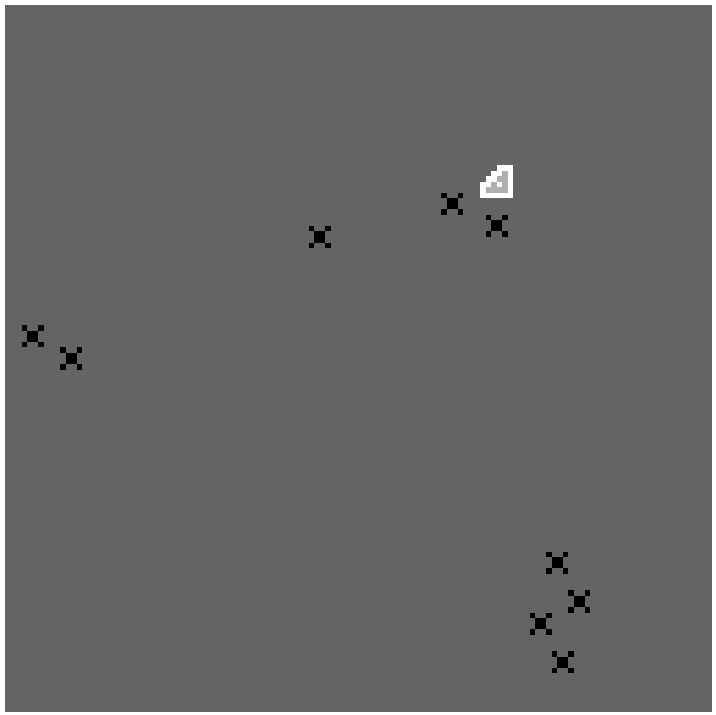

Distractors similar; target similar to distractors

- Distractors can be grouped (which speeds up search).
- But targets can be grouped with distractors, requiring backtracking.

Distractors similar; target similar to distractors

- Distractors can be grouped (which speeds up search).
- But targets can be grouped with distractors, requiring backtracking.

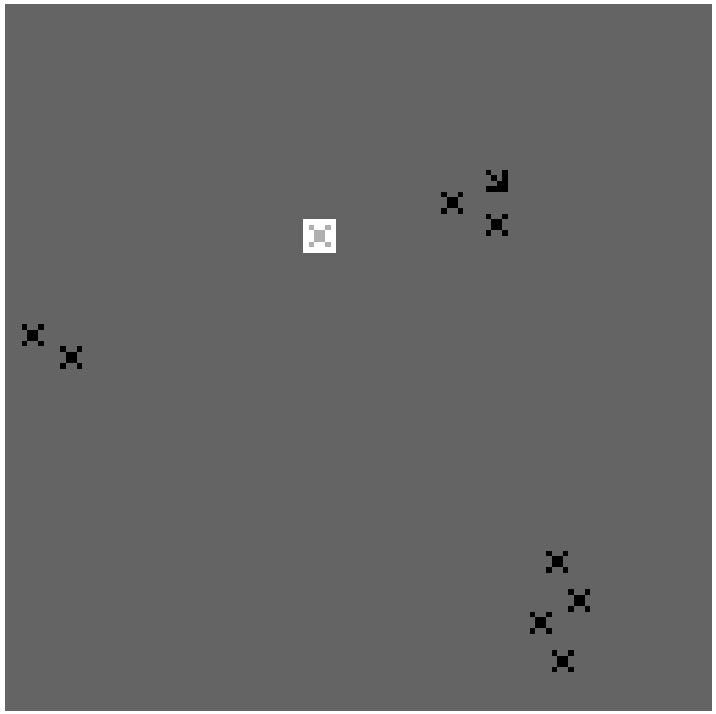

Distractors similar; target similar to distractors

- Distractors can be grouped (which speeds up search).
- But targets can be grouped with distractors, requiring backtracking.

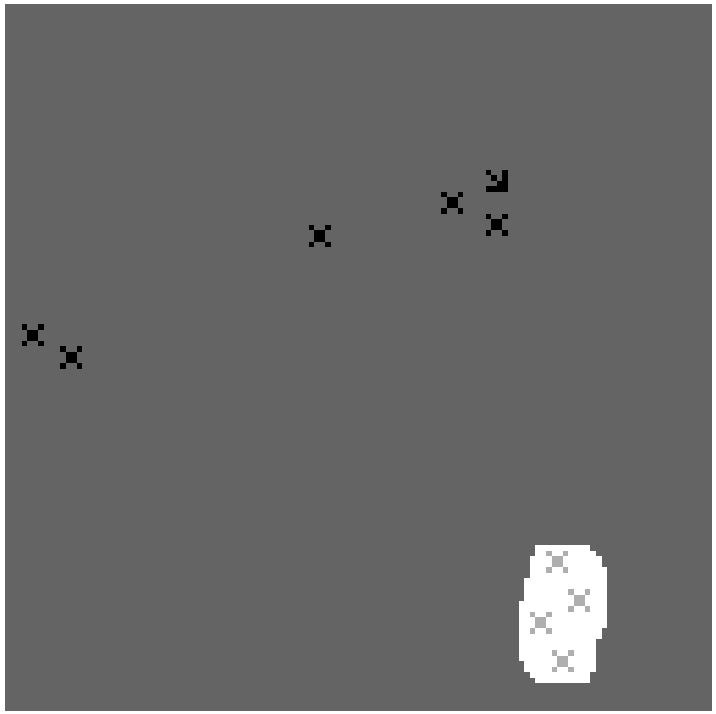

Distractors similar; target similar to distractors

- Distractors can be grouped (which speeds up search).
- But targets can be grouped with distractors, requiring backtracking.

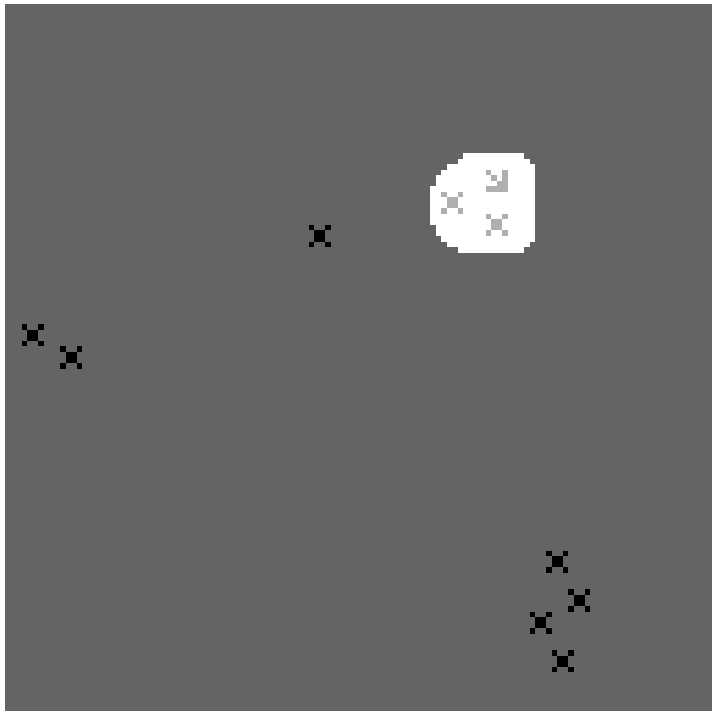

Distractors similar; target similar to distractors

- Distractors can be grouped (which speeds up search).
- But targets can be grouped with distractors, requiring backtracking.

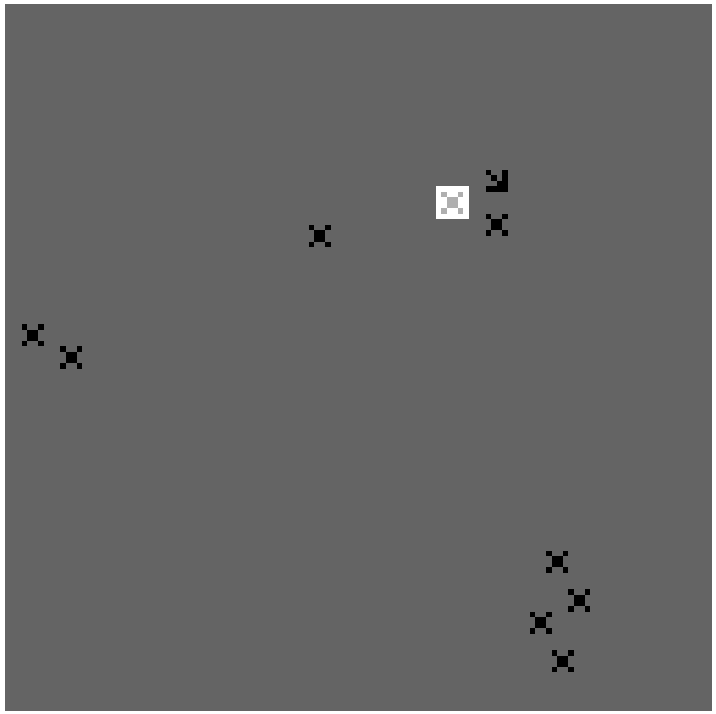

Distractors similar; target similar to distractors

- Distractors can be grouped (which speeds up search).
- But targets can be grouped with distractors, requiring backtracking.