Dude, where's my Warthog?

From **Pathfinding** to **General Spatial Competence**

Adapted from talk by
Damián Isla
Bungie Studios

The Grand Question

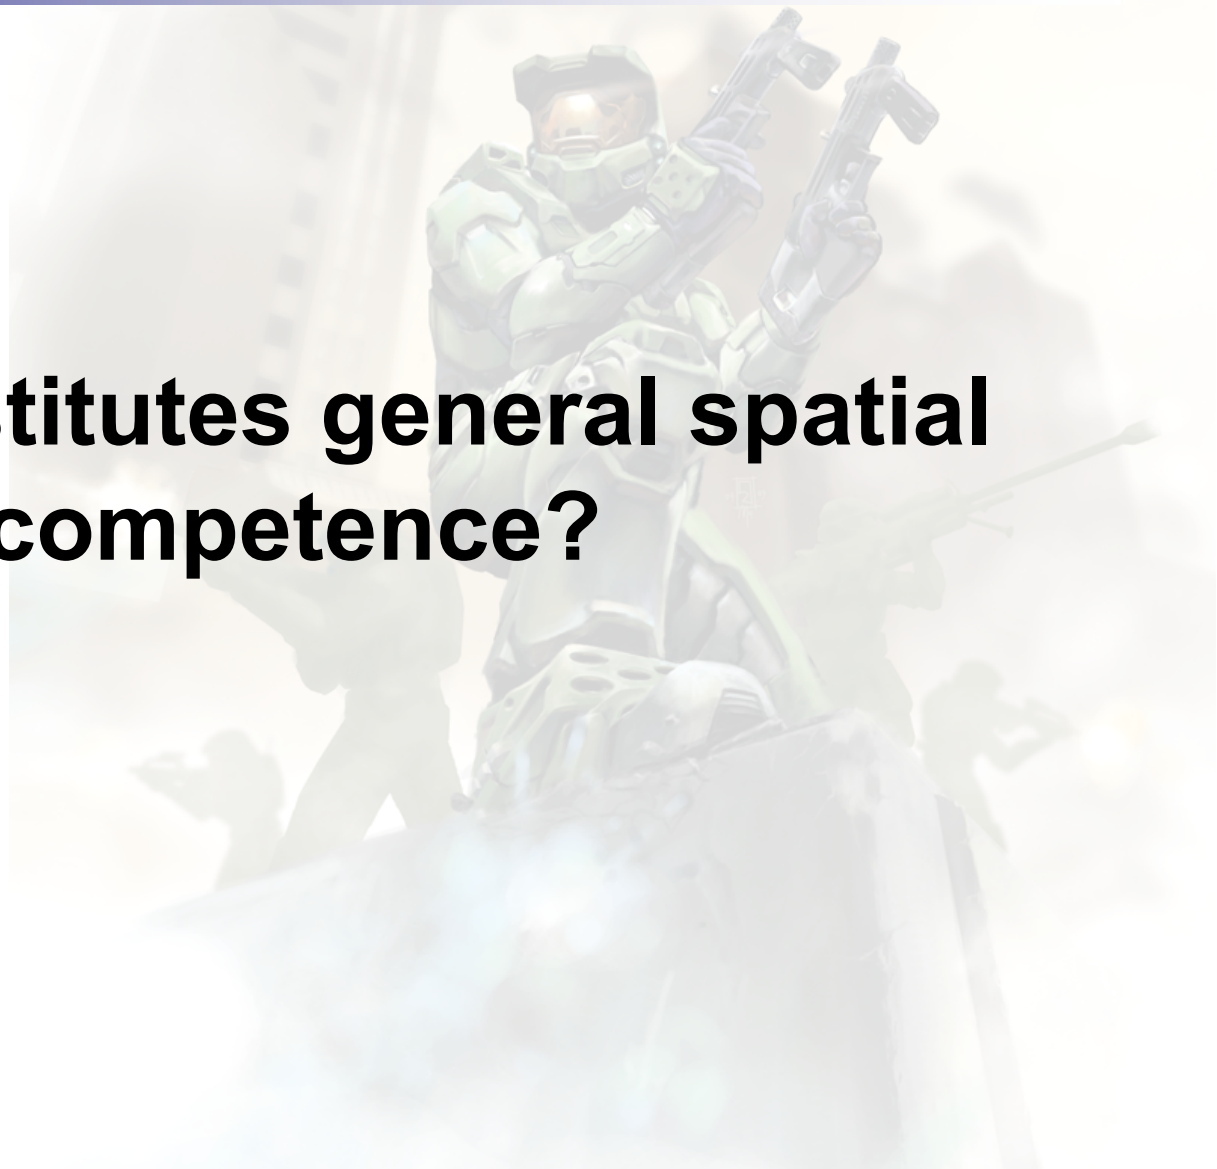

What constitutes general spatial competence?

The Halo Approach

- AIs are given a “playground”, within which they are allowed to do whatever they want
- The designer defines the flow of battle by moving the AI from one playground to another
- **The designer’s time is precious**
- Relatively little spatial information is explicitly entered by the designers

Problems Solved in Halo2

- Static Pathfinding
  - Navigation mesh (ground)
  - Waypoint network (airborne)
  - Raw pathfinding
  - Path-smoothing
  - Hint integration (jumping, hoisting, climbing)
  - Static scenery-based hints
  - Static scenery carved out of environment mesh
- Static feature extraction
  - Ledges and wall-bases
  - Thresholds
  - Corners
  - Local environment classification
- Object features
  - Inherent properties (size, mass)
  - Oriented spatial features
  - Object behaviors (mount-to-uncover, destroy cover)
- Dynamic Pathfinding
  - Perturbation of path by dynamic obstacles
  - "Meta-search" / Thresholds / Error stages
  - Obstacle-traversal behaviors
    - Vaulting, hoisting, leaping, mounting, smashing, destroying
- Path-following
  - Steering on foot (with exotic movement modes)
  - Steering a vehicle (e.g. ghost, warthog, banshee)
- Interaction with behavior
  - What does behavior need to know about the way its requests are being implemented?
  - How can pathfinding impact behavior?
- Body configuration
  - Flying, landing, perching
  - Cornering, bunkering, peeking
- Spatial analysis
  - Firing position selection
  - Destination evaluation based on line-of-sight, range-to-target, etc.
- "Local spatial behaviors"
  - Line-tracing (e.g., for diving off cliffs)
  - Not facing into walls
  - Crouch in front of each other
  - Don't walk into the player's line of fire
  - Curing isolation
  - Detecting blocked shots
- Reference frames
  - The viral nature of the reference frame
- Cognitive model / Object persistence
  - Honest perception
  - Simple partial awareness model
- Search
  - Simple by design
  - Group search
- Spatial conceptualization
  - DESIGNER-PROVIDED
  - Zones, Areas (areas), Firing positions (locations)

Problems Solved in Halo2

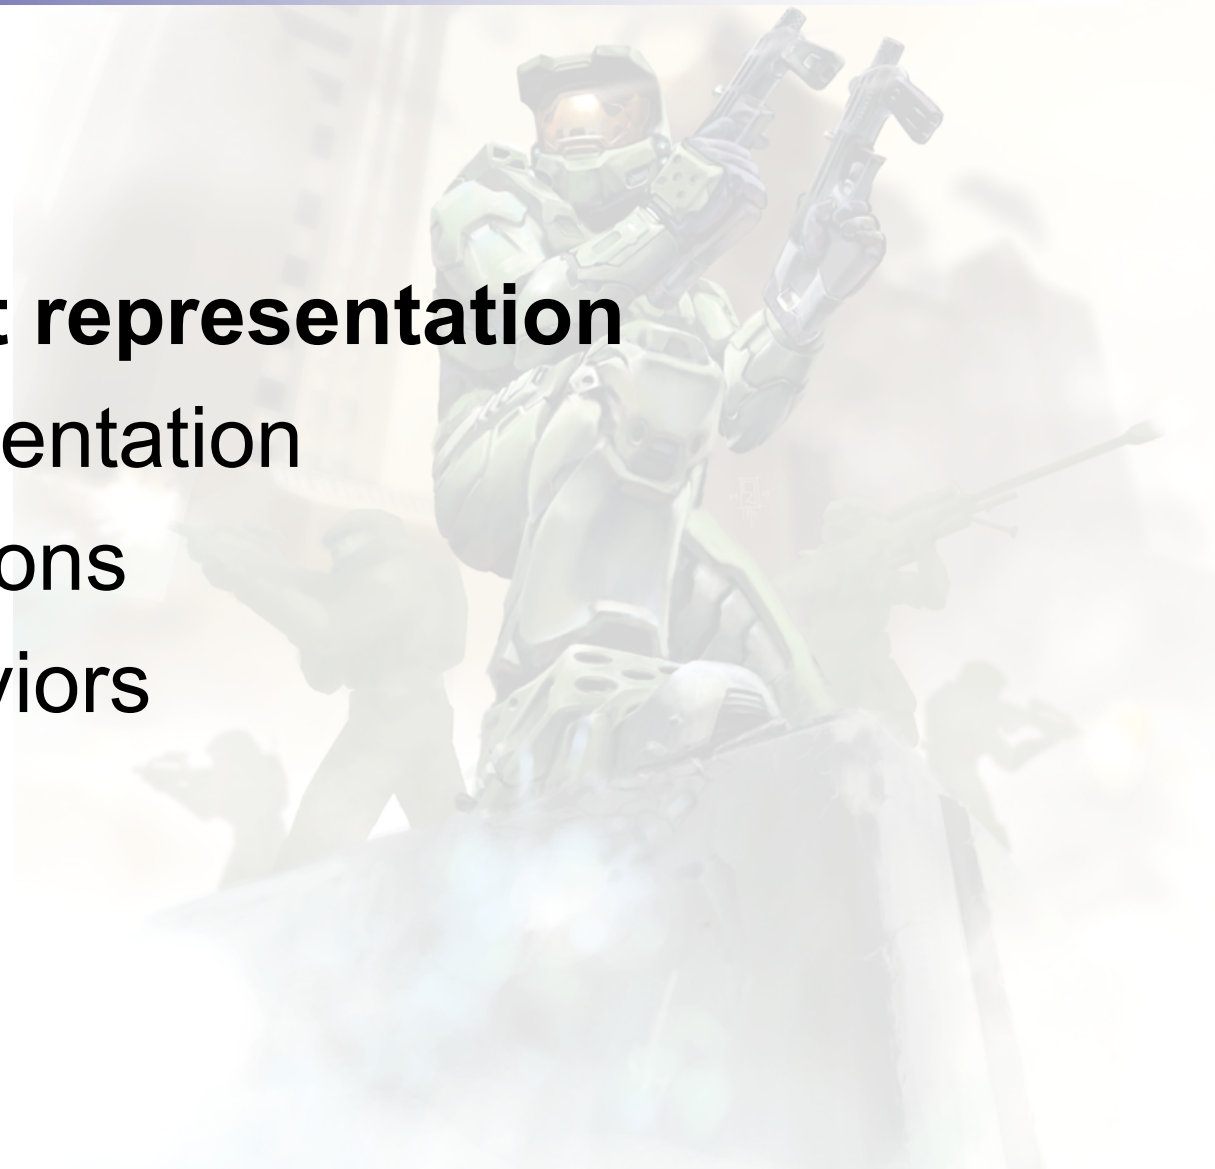

- **Environment representation**
- Object representation
- Spatial Relations
- Spatial Behaviors

Environment Representation

How do we represent the environment to the AI?

An important constraint: as few restrictions as possible on the form the geometry can take

- The environment artist's time (and artistic freedom) is precious

Environment Representation

Halo2: **navigation mesh** constructed from the raw environment geometry

- CSG (Constructive Solid Geometrry) “stitching in” of static scenery
- Optimization
- “sectors”: convex, polygonal, but not planar

Spatial Feature Extraction

A lot of features we're interested in can be extracted automatically …

- Surface categorization / characterization
- Surface connectivity
- Overhang detection
- Interior/exterior surfaces
- Ledges
- Wall-bases
- "Leanable" walls
- Corners
- "Step" sectors
- Thresholds
- Local environment classification
  - Captures the "openness" of the environment at firing positions

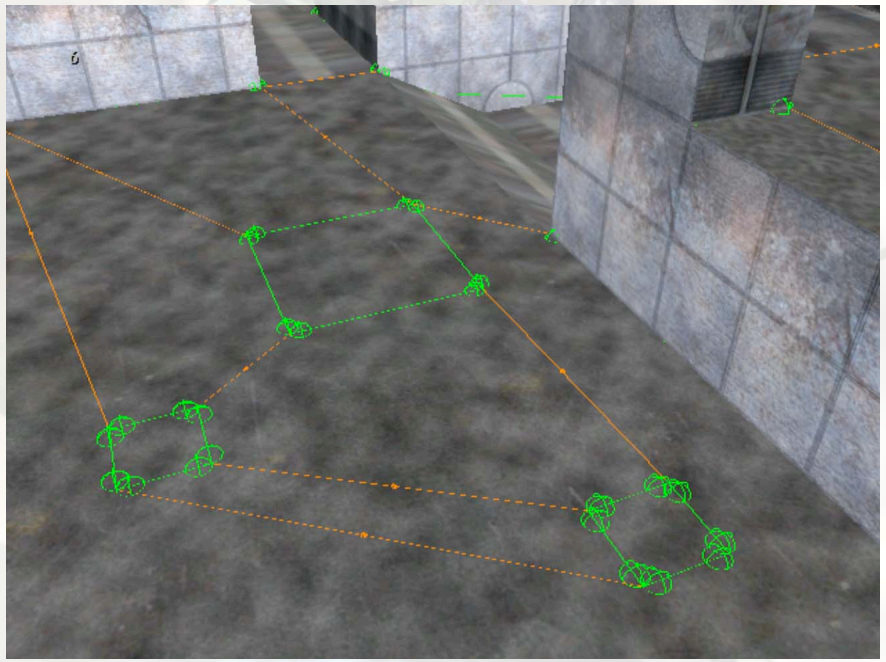

Spatial Feature Extraction

… and a lot can't. So we make the designers do it.

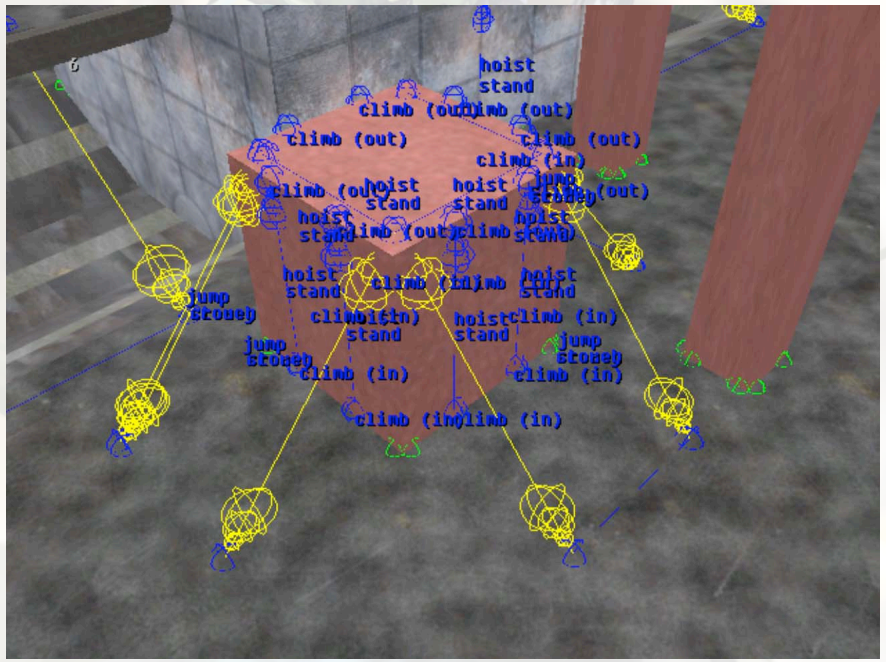

Designer "hints":

- Jumping
- Climbing
- Hoisting
- "Wells"
- Manual fix-up for when the automatic processes fail:
  - Cookie-cutters
  - Connectivity hints

Place

But that's not enough.

The navigation graph is good for metric queries (e.g., would I run into a wall if I were to move 10 feet in this direction?)

… but not a good representation for reasoning about space [I want to go behind the desk]

Place

Psychologists talk about *cognitive maps* as the internal representation of behaviorally-relevant **places** and how they relate.

A couple of interesting properties:

- Not metric
- Fuzzy
- Hierarchically organized

Useful for:

- Landmark navigation
- Dead-reckoning
- Place-learning
- Self-localization

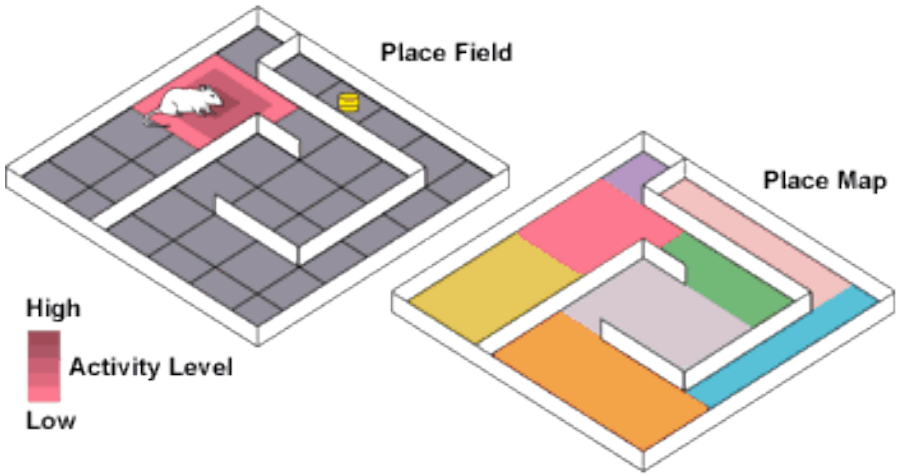

From http://www.brainconnection.com

Place

In the ideal world, we would be able to automatically construct some kind of spatial semantic network

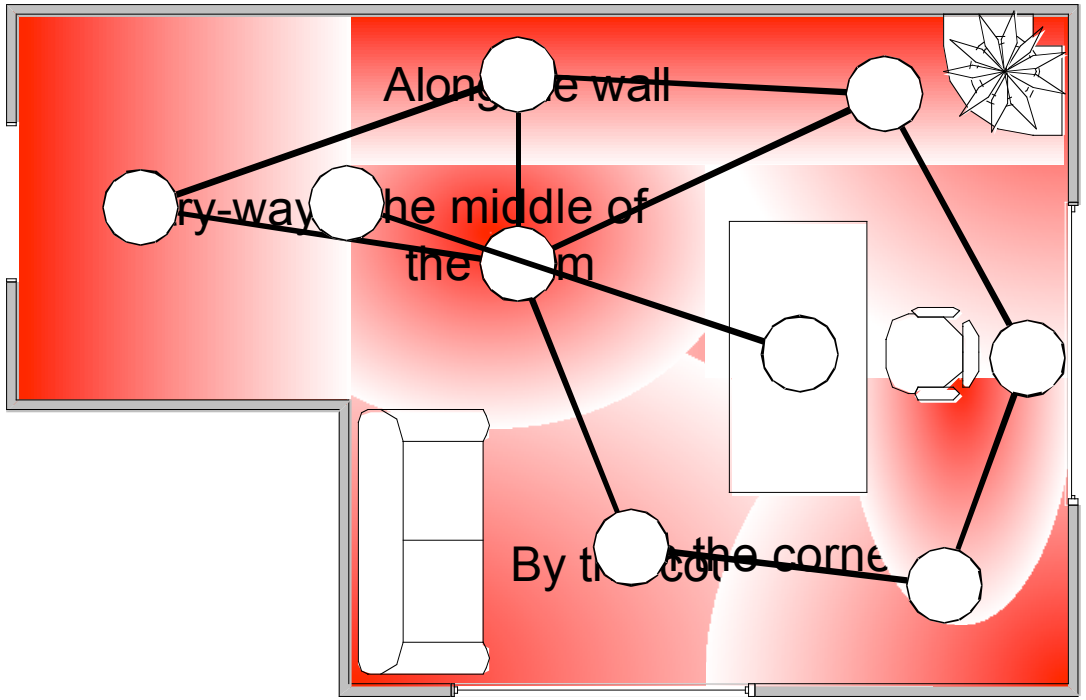

Place

The Halo place representation:

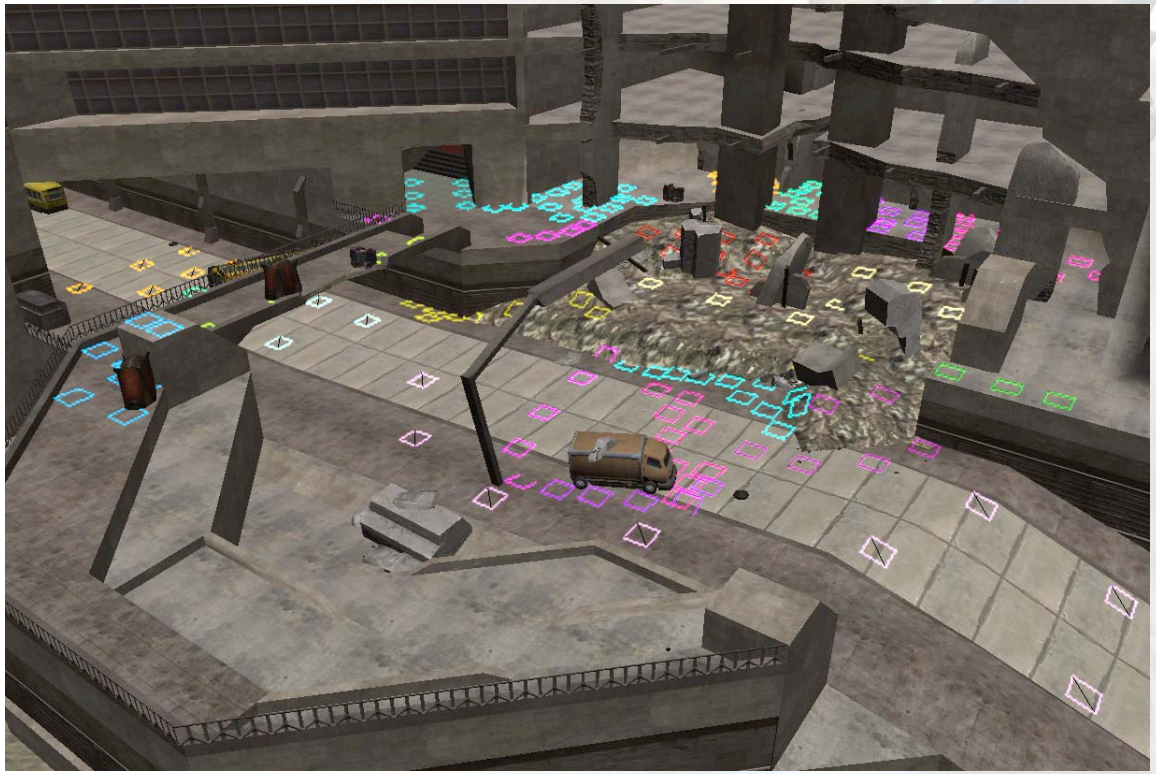

A shallow hierarchy of spatial groupings:

Zones → Areas → Positions

Place

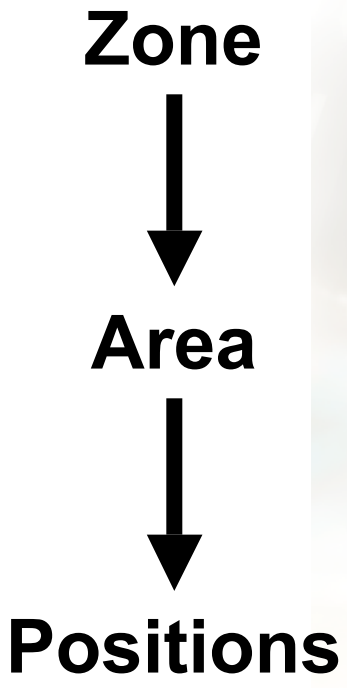

Organizational
Attached to a bsp

Defines the playground
Following

Tactical analysis

Place

But we lose something from taking a designer-authored approach to place:

- No *relational* information
  - A LOT of work for the designers to enter
- Very little *semantic* information
- The Designer has to do it

Place

Note, in any case, the dichotomy between our cognitive map and our navigation mesh:

Navigation mesh is continuous, metric

Cognitive map is discrete, relational

Problems Solved in Halo2

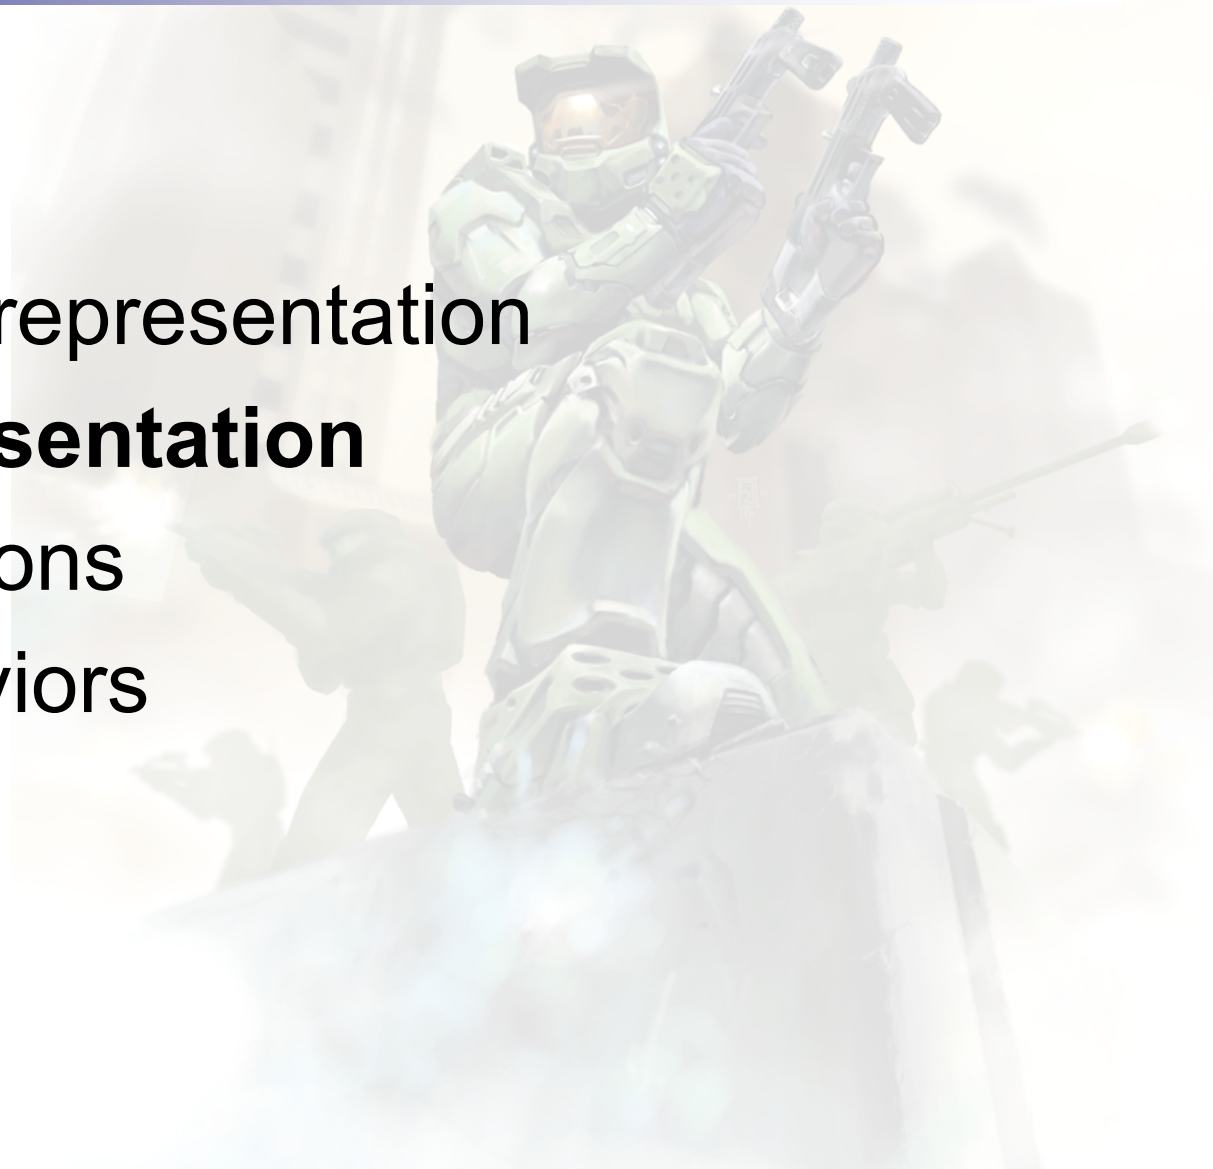

- Environment representation
- **Object representation**
- Spatial Relations
- Spatial Behaviors

Object Representation

How do we represent objects in a useful way to the AI?

Assume that static objects are part of the environment

Dynamic object representation: three ways to see an object:

- Inherent properties
- Volume
- Spatial features

- **Size**
- **Leap-speed**
- **Destructible**
- **Custom behavior X**

Volume

- Rough approximation using *pathfinding spheres*
- Spheres projected to AI's ground-plane at pathfinding time (to become pathfinding *discs*)
- A perturbation of the smoothed path

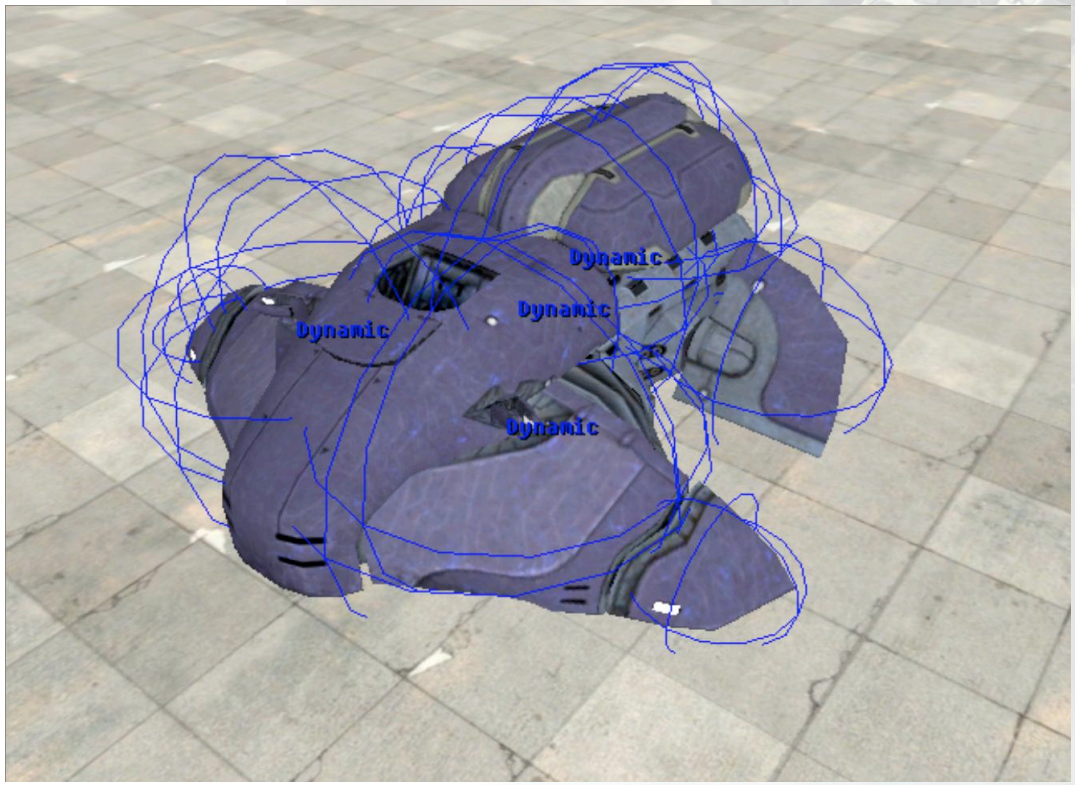

Spatial Object Features

- An object advertises its “affordances”, i.e., the things that can be done with / to it
- But they must do so in a geometrically precise way in order to be useful

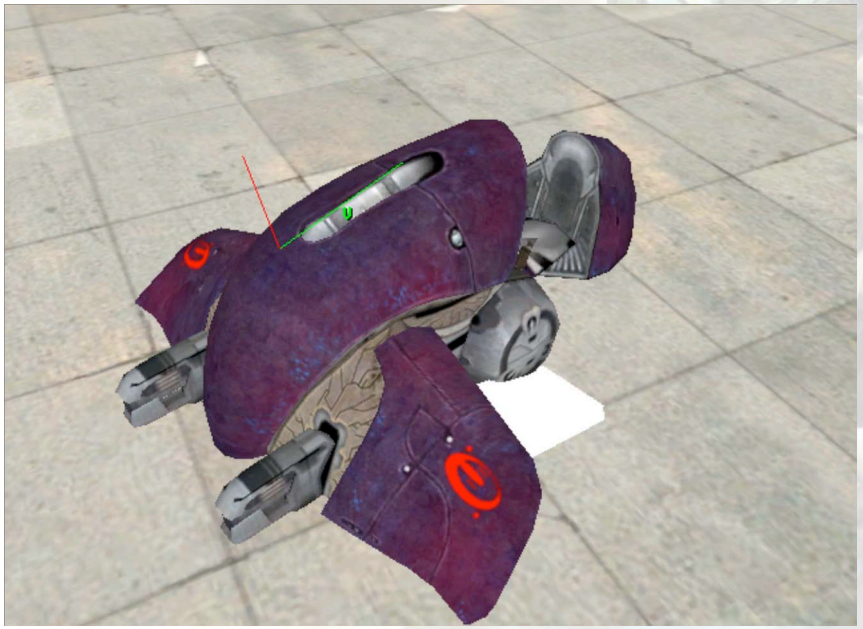

Implementation: “object markers”

- Rails or points
- Orientation vector indicates when the affordance is active
- An object has different properties at different orientations

Object Representation

Volume + Features = How the AI understands shape

Adding rich AI information becomes a fundamental part of the modeling of the object (just like authoring collision and physics models)

Used for

- Explicit behavior
  - Cornering (corner feature)
  - Mount-to-uncover (mount feature)
  - Destroy cover (destructible property)
- Pathfinding obstacle-traversal
  - Vault (vault feature)
  - Mount (mount feature)
  - Smash (size property)
  - Destroy obstacle (destructible property)

Problems Solved in Halo2

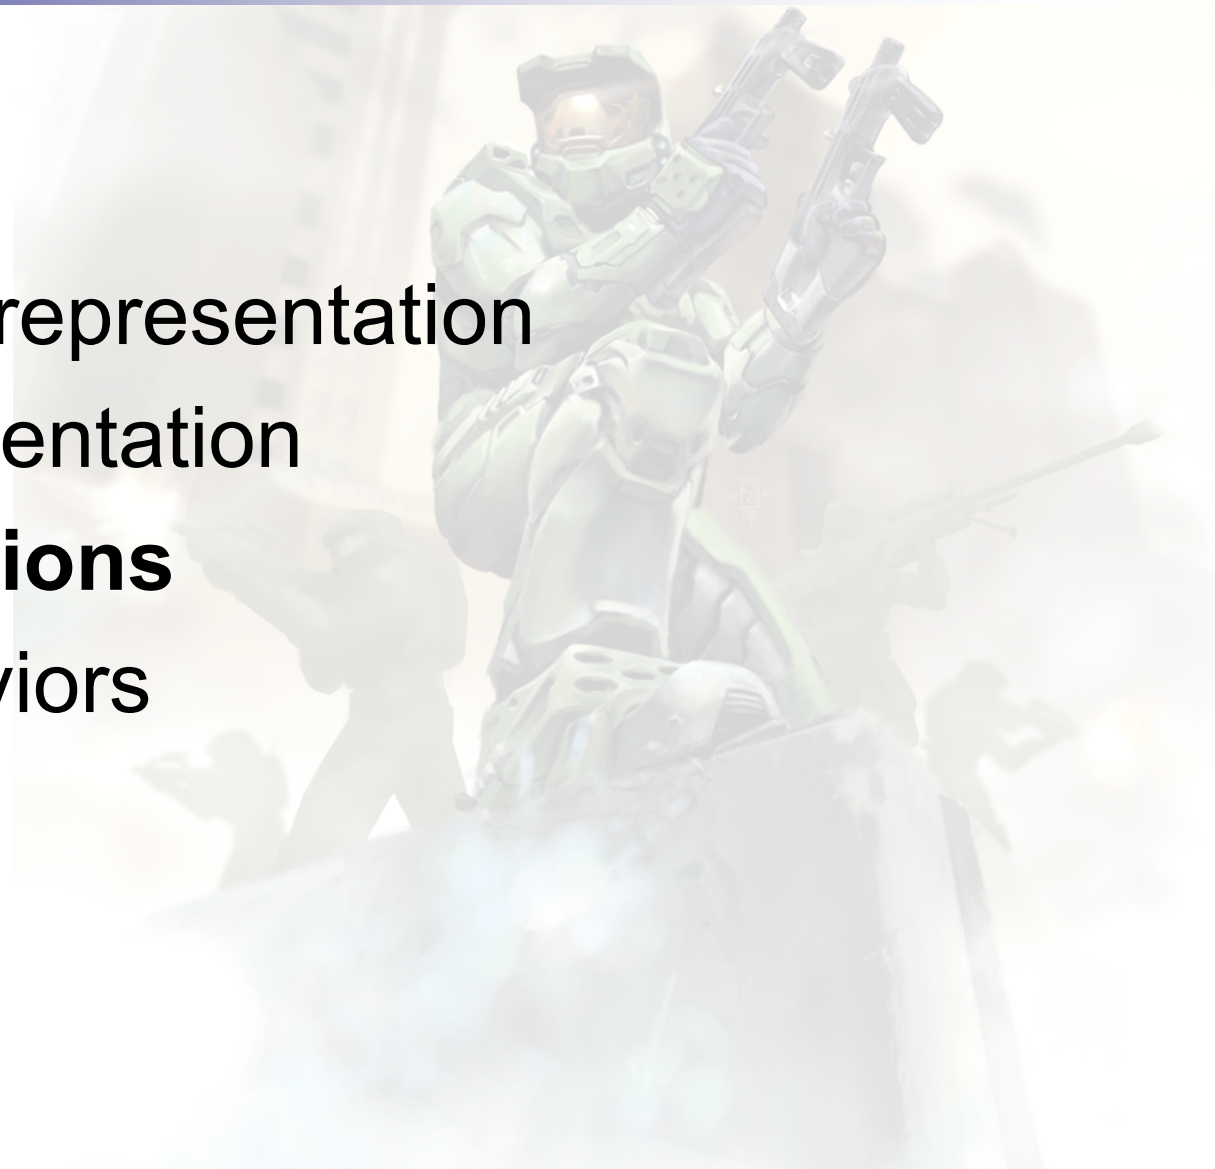

- Environment representation
- Object representation
- **Spatial Relations**
- Spatial Behaviors

Spatial Relations

How do the objects in the AI's knowledge model relate to each other spatially?

Well first of all, what's IN the knowledge model?

In Halo2:

- Potential targets (enemies)
- Player(s)
- Vehicles
- Dead bodies
- **And that's it.**

Spatial Relations

What the Knowledge Representation (KR) people think…

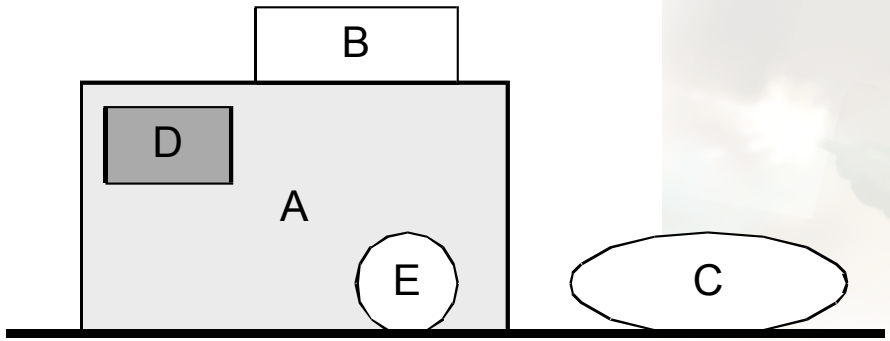

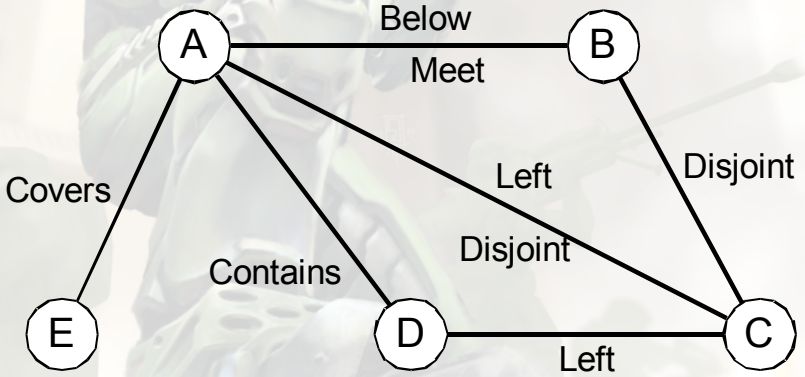

From Papadias et. al Acquiring, Representing and Processing Spatial Relations, Proceedings of the 6th International Symposium on Spatial Data Handling, Edinburgh, 1994

Spatial Relations

Some rudimentary Halo2 examples:

- Grenade-throwing
  - Find clusters of nearby enemies
- Blocked shots
  - Recognize “I can see my target, and I wanted my bullets to go X meters, but they only went 0.6X meters. I must be blocked.”
- Destroy-cover
  - Recognize that my target is behind destructible cover
- Mount-to-uncover
  - Recognize that my target is behind a mountable object

Behind the Space Crate

The notion of “behind” could happen at multiple levels

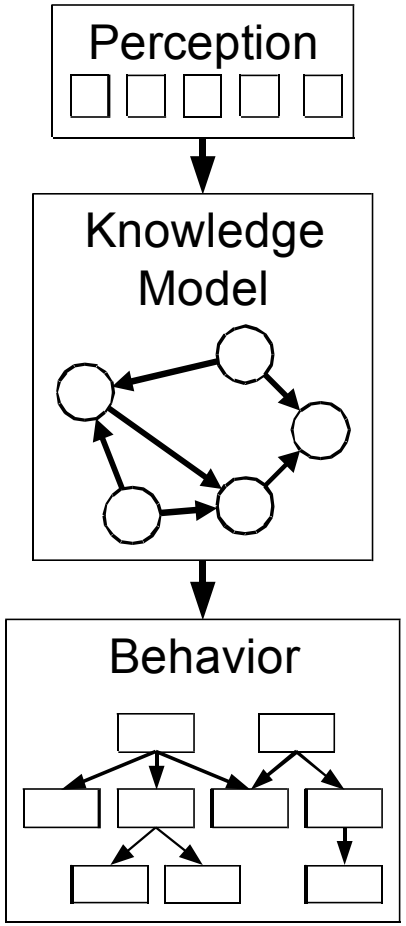

For each target, store the object, if any, that is blocking my view of it

Halo2

Make note of spatial relationship between objects in my knowledge model (semantic network)

When I lose track of my target, perform an in-line spatial computation to determine if I should try destroying an object in front of it.

Behind the Space Crate

The notion of “behind” could happen at multiple levels

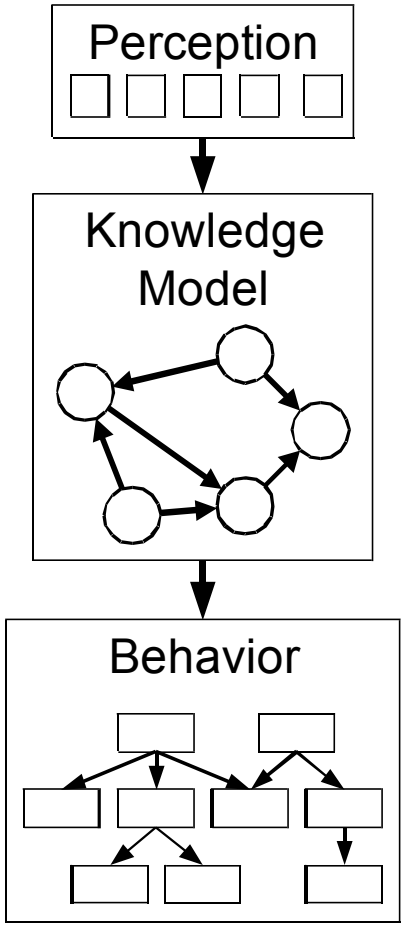

Find the appropriate obstacle, tell the KM to instantiate a representation for it

Store the “behind” information in a general, public (semantic net-like) format

When we lose track of our target, tell perception we’re interested in objects that are in front of it

Behind the Space Crate

All of which is just to say:

- “Behind” is not an entirely trivial concept
- The collection of spatial-relation information and the management of their representation structures are not trivial either!

Spatial Groupings

E.g.,:

- Clusters of enemies
- Battle fronts
- Battle vectors

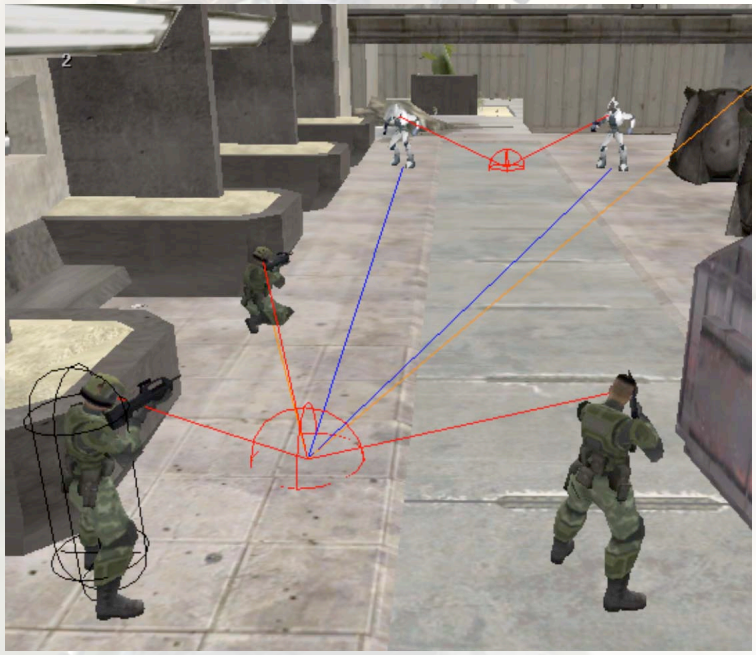

In Halo2: perform dynamic clumping of nearby allies, for:

- Joint behavior
- Call-response combat dialogue
- Shared perception

BUT, not a perceptual construct!

Spatial Groupings

Cognitive Efficiency

- One, two, many
- Give groupings first-class representation in the AI's knowledge model?
- Another hierarchy
  - See the many as one
  - Or, instantiate individuals as necessary

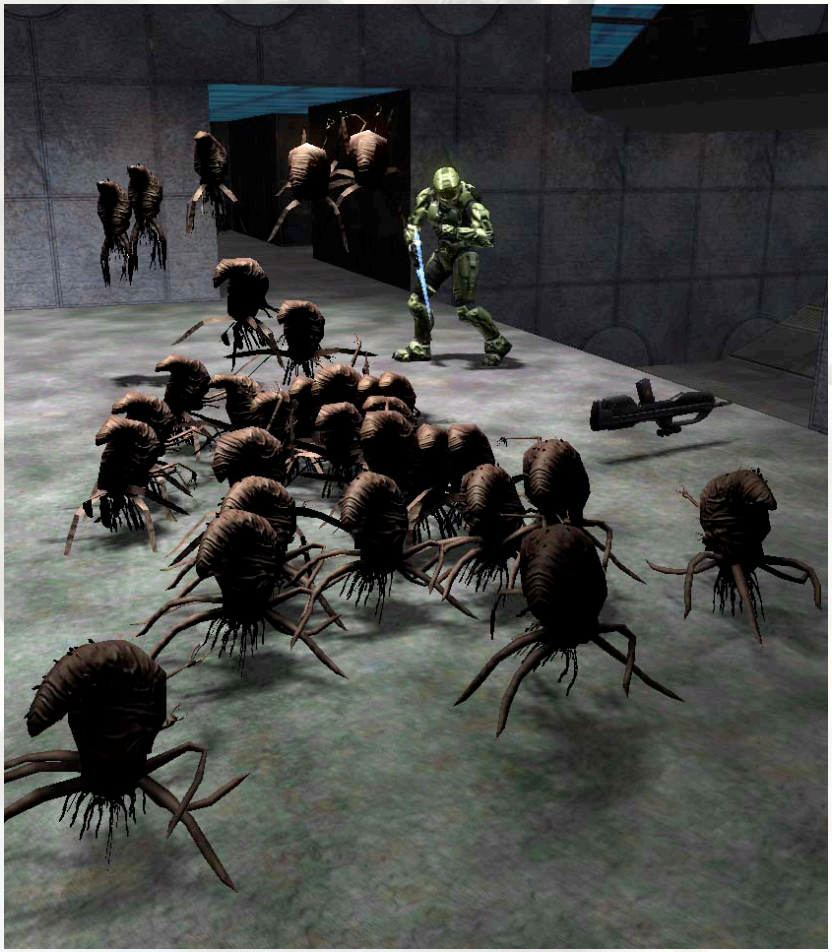

Problems Solved in Halo2

- Environment representation
- Object representation
- Spatial Relations
- **Spatial Behaviors**

Spatial behavior

Two types:

- World-relative:
  - Generally uses the cognitive map
  - Typically recognized Behaviors
  - E.g., fight, follow, search
- Viewer-relative:
  - Generally through local spatial queries
  - Things that should just sort of, you know, happen

Fighting

Position evaluation based on

- Range-to-target
- Line-of-sight to target
- Distance from current position
- Distance to the player and other allies
- Easy!

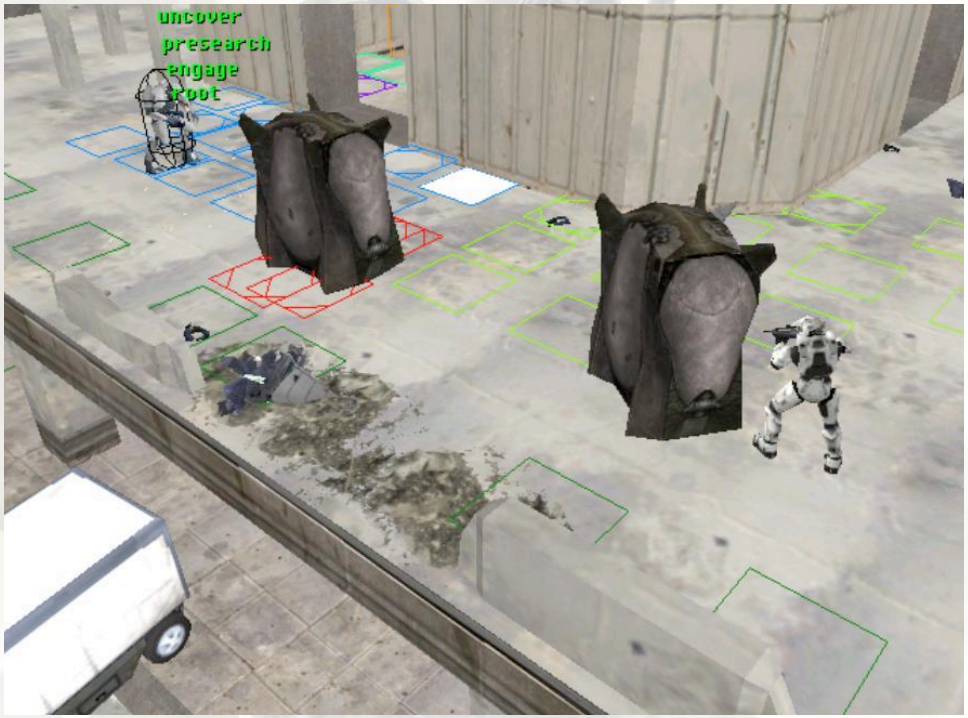

This is the tactical spatial analysis problem.

And there are lots of published solutions out there. See in particular *Van Der Sterren, Killzone's AI: Dynamic Procedural Combat Tactics, GDC 2005*

Following

Easy to do mediocrely

Hard and complicated to do well

- Stay close
- Not too close
- Try and stay in front (so that player can see and appreciate) but don't get in the way and don't block the player's line of fire
- What does "in front" even mean?
- Don't follow when not appropriate

In the ideal case, need player-telepathy

- Look for explanation for the player's movement, then determine whether that explanation warrants MY adjusting my position as well.

Search

The most interesting of the spatial behaviors

As complicated as you want to get:

- Fake it completely
  - Play a “look around and shrug” animation
- Pretend you don’t know where the player is while exclaiming “Where’d he go?!”
- Simple scripted search routines
- Basic stateless hidden location-uncovering
- … based on spatial structure and spatial semantics …
- … based on spatial structure and semantics and player model

The more complicated the search model, the more complicated the perception and knowledge models and the **maps** needed to support it.

Viewer-relative Reference Frames

The most interesting use:
frames of motion

E.g., AIs running around on the back of the giant scarab tank

Reference Frames

The hard part:

- Moving sectors
- Adapting A*
  - A* in local space except across ref-frame boundaries
  - Final path cached in local space(s)

- A new point representation:

(x,y,z,f)

Reference Frames

Once we start using it one place, we have to use it everywhere!

- Sectors
- Firing-positions
- Scripting points
- Target locations
- Last-seen-location
- Burst targets
- Etc.

Results in a generalized “understanding” of reference frames

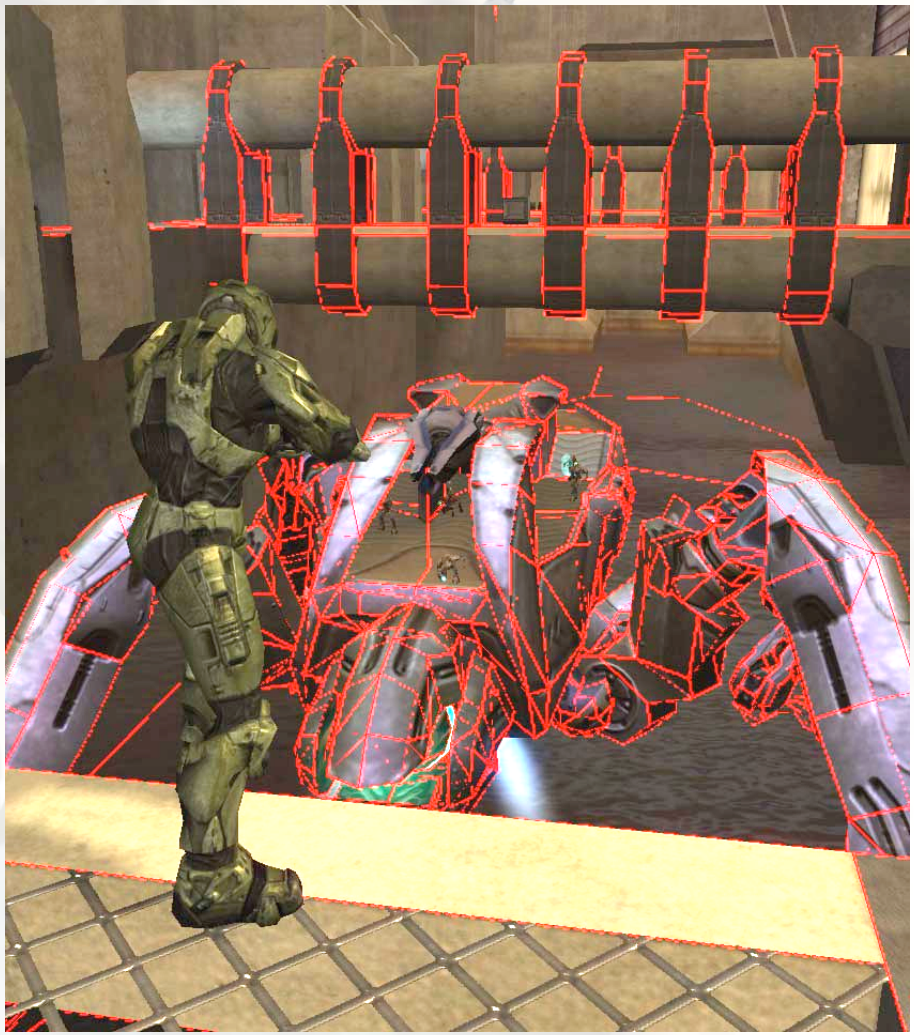

Viewer-relative Behaviors

The grab-bag:

- When stopped, don't face into walls ["react to mistakes"]
- Don't pick a spot that blocks a friend's line of fire
- Don't block the player's line of fire **ever**
- Don't even **cross** the player's line of fire
- Crouch down when someone behind me is shooting
- Move with my allies, rather than treating them as obstacles
- Get off non-pathfindable surfaces

These are hard, because they're not exclusive behaviors

- Things to "keep in mind".
- Which means that high-level behaviors always need to be robust to their effects.

Unsolved Mysteries

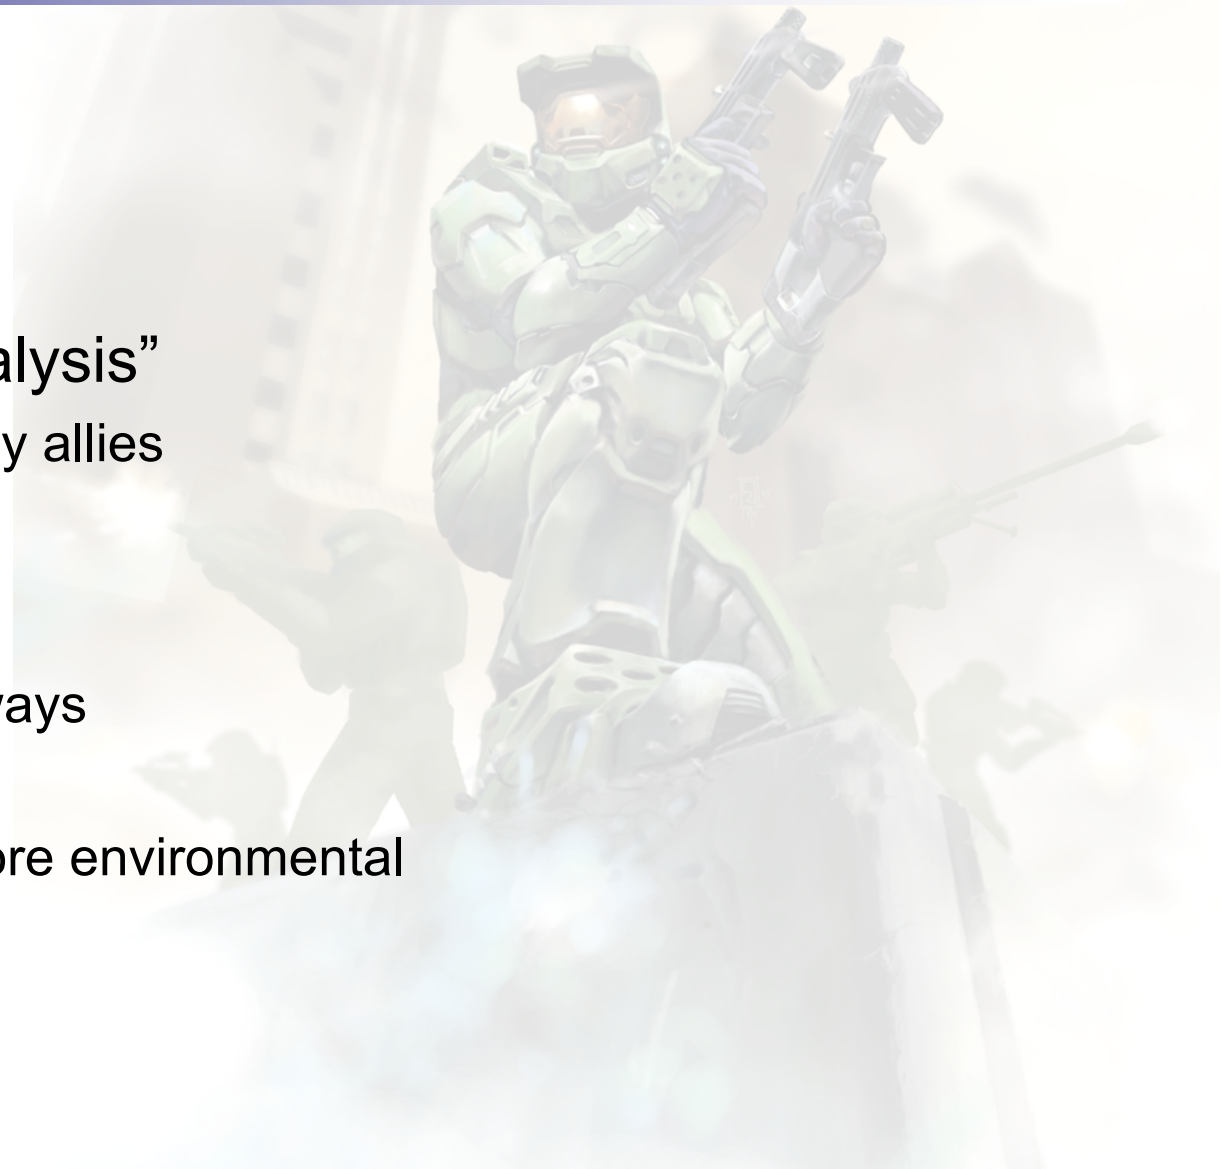

- Group movement
  - Queuing
  - Formations
- “Configuration analysis”
  - My relation with my allies
- Anticipation
- Spatial Semantics
  - Rooms and doorways
  - Inside / outside
  - Understanding more environmental spatial features