


The Map Builder

15-494 Cognitive Robotics
David S. Touretzky &
Ethan Tira-Thompson

Carnegie Mellon
Spring 2008

Horizontal Field of View



Rat: 300 deg.




Human: 200 deg.



AIBO: 60 deg.

Seeing A Bigger Picture


- How can we assemble an accurate view of the robot's surroundings from a series of narrow camera frames?
- First, convert each image to symbolic form: shapes.
- Then, match the shapes in one image against the shapes in previous images.



- Construct a “local map” by matching up a series of camera images.

Can't Match in Camera Space

- We can't match up shapes from one image to the next if the shapes are in camera coordinates. Every time the head moves, the coordinates of the shapes in the camera image change.
- Solution: switch to a body-centered reference frame.
- If we keep the body stationary and only move the head, the coordinates of objects won't change (much) in the body reference frame.



Planar World Assumption

- How do we convert from camera-centered coordinates to body-centered coordinates?
- Need to know the camera pose: can get that from the kinematics system.
- Is that enough? Unfortunately, no.
- Add a planar world assumption: objects lie in the plane. The robot is standing on that plane.
- Now we can get object coordinates in the body frame.

Shape Spaces




- **camShS** = camera space
- **groundShS** = camera shapes projected to ground plane
- **localShS** = body-centered (egocentric space);
constructed by matching and importing shapes
from groundShS across multiple images
- **worldShS** = world space (allocentric space);
constructed by matching and importing shapes
from localShS
- The robot is explicitly represented in worldShS

Invoking The Map Builder



- Let's map the tic-tac-toe board:







Frame 1







Frame 2






Frame 3




Frame 4



Frame 5



Final Local Map



Shape Matching Algorithm

- Shape type and color must match exactly.
- Coordinates must be a reasonably close match for points, blobs, and ellipses.
- Lines are special, because endpoints may be invalid:
 - If endpoints are valid, coordinates should match.
 - If invalid in local map but valid in ground space, update the local map to reflect the true endpoint location.
- Coordinates are updated by weighted averaging.

Noise Removal

- Noise in the image can cause spurious shapes. A long line might appear as 2 short lines separated by a gap, or a noisy region might appear as a short line.
- Assign a confidence value to each shape in local map.
- Each time a shape is seen: increase its confidence.
- If a shape *should* be seen but is not, decrease its confidence.
- Delete shapes with negative confidence.

Where to Look?

- Start with the shapes visible in the camera frame.
- Move the head to fixate each shape: get a better look.
- If a line runs off the edge of the camera frame, move the head to try to find the line's endpoints.
 - If the head can't rotate any further, give up on that endpoint.
- If an object is partially cut off by the camera frame, don't add it to the map because we don't know its true shape.
 - Move the head to bring the object into view.

Programming the MapBuilder

- A instance of MapBuilder is included as a member of VisualRoutinesBehavior, called mapbuilder.

```
#include "DualCoding/DualCoding.h"

class LocalMapDemo : public VisualRoutinesBehavior {
public:
    LocalMapDemo() : VisualRoutinesBehavior() {}

    virtual void DoStart() {
        VisualRoutinesBehavior::DoStart();

        MapBuilderRequest req;

        ... program the mapbuilder instructions

        mapbuilder.executeRequest(req);
    }
};
```

2 Ways To Invoke MapBuilder

- If processing a single camera image, the MapBuilder returns immediately:

```
mapbuilder.executeRequest(req);  
NEW_SHAPEVEC(blobs, BlobData, select_type<BlobData>(localShS));
```

- If multiple camera images and/or head motion required, must wait for completion (status event):

```
erouter->addListener(this, EventBase::mapbuilderEGID);  
mapreq_id = mapbuilder.executeRequest(req);  
...  
  
void processEvent (const EventBase &e) {  
    if ( e.getGeneratorID() == EventBase::mapbuilderEGID &&  
        e.getSourceID() == mapreq_id ) {  
        NEW_SHAPEVEC(blobs, BlobData, select_type<BlobData>(localShS));  
        ...  
    }  
}
```

MapBuilderRequest Parameters

- RequestType
 - cameraMap
 - groundMap
 - localMap
 - worldMap
- Shape parameters:
 - objectColors
 - occluderColors
 - maxDist
 - minBlobArea
- Lookout control:
 - motionSettleTime
 - numSamples
 - sampleInterval
 - pursueShapes
 - searchArea
 - doScan, dTheta
 - manualHeadMotion

Programming the MapBuilder

```
const int pink_index = ProjectInterface::getColorIndex("pink");  
const int blue_index = ProjectInterface::getColorIndex("blue");  
const int orange_index = ProjectInterface::getColorIndex("orange");
```

```
MapBuilderRequest req(MapBuilderRequest::localMap);
```

```
req.numSamples = 5; // take mode of 5 images to filter out noise  
req.maxDist = 1200; // maximum shape distance 1200 mm  
req.pursueShapes = true;
```

```
req.objectColors[lineDataType].insert(pink_index);  
req.occluderColors[lineDataType].insert(blue_index);  
req.occluderColors[lineDataType].insert(orange_index);
```

```
req.objectColors[ellipseDataType].insert(blue_index);  
req.objectColors[ellipseDataType].insert(orange_index);
```

```
unsigned int mapreq_id = MapBuilder::executeRequest(req);  
erouter->addListener(this, EventBase::mapBuilderEGID,  
                    mapreq_id, EventBase::statusETID);
```



Programming the MapBuilder

```
void processEvent(const EventBase &event) {  
    cout << "Map builder returned " << event.getDescription()  
        << endl;  
    cout << "Found " << localShS.allShapes().size() << " shapes." << endl;  
}
```



Qualitative Spatial Reasoning

- Reading for today:
How qualitative spatial reasoning can improve strategy game AIs
Ken Forbus, James Mahoney, and Kevin Dill (2002)
- Uses visual routines to “reason about” maps, e.g., compute reachability, calculate paths, etc.
- Possible research topic: applying these ideas to world maps in Tekkotsu.

