LONG-TERM SIMULATION OF DYNAMIC, INTERACTIVE WORLDS WITH MORSE

Marc Hanheide, University of Lincoln
(in collaboration with Lars Kunze, Birmingham, and many other colleagues)

STRANDS

Spatio-Temporal Representation and Activities for Cognitive Control in Long-Term Scenarios

STRANDS will produce intelligent mobile robots that are able to run for months in dynamic human environments. We will provide robots with the longevity and behavioural robustness necessary to make them truly useful assistants in a wide range of domains. Such long-lived robots will be able to learn from a wider range of experiences than has previously been possible, creating a whole new generation of autonomous systems able to extract and exploit the structure in their worlds.

STRANDS

Spatio-Temporal Representations and Activities For Cognitive Control in Long-Term Scenarios

LINDA AT NIGHT

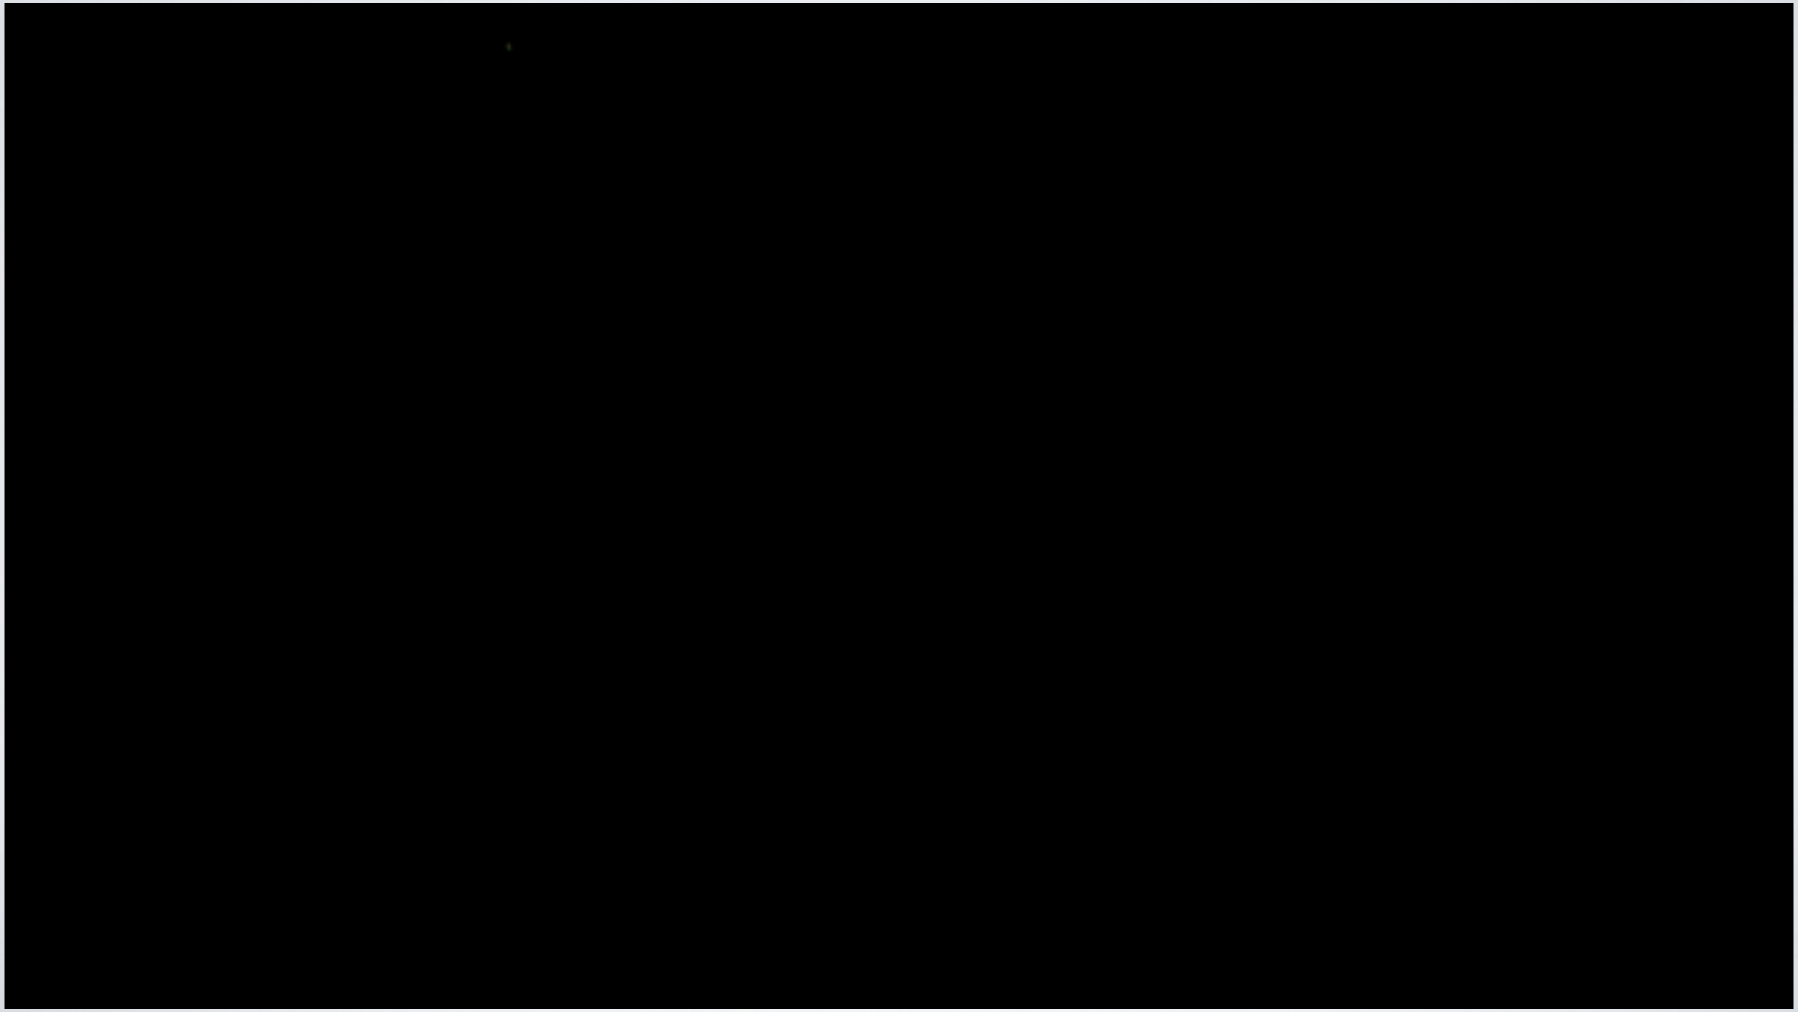

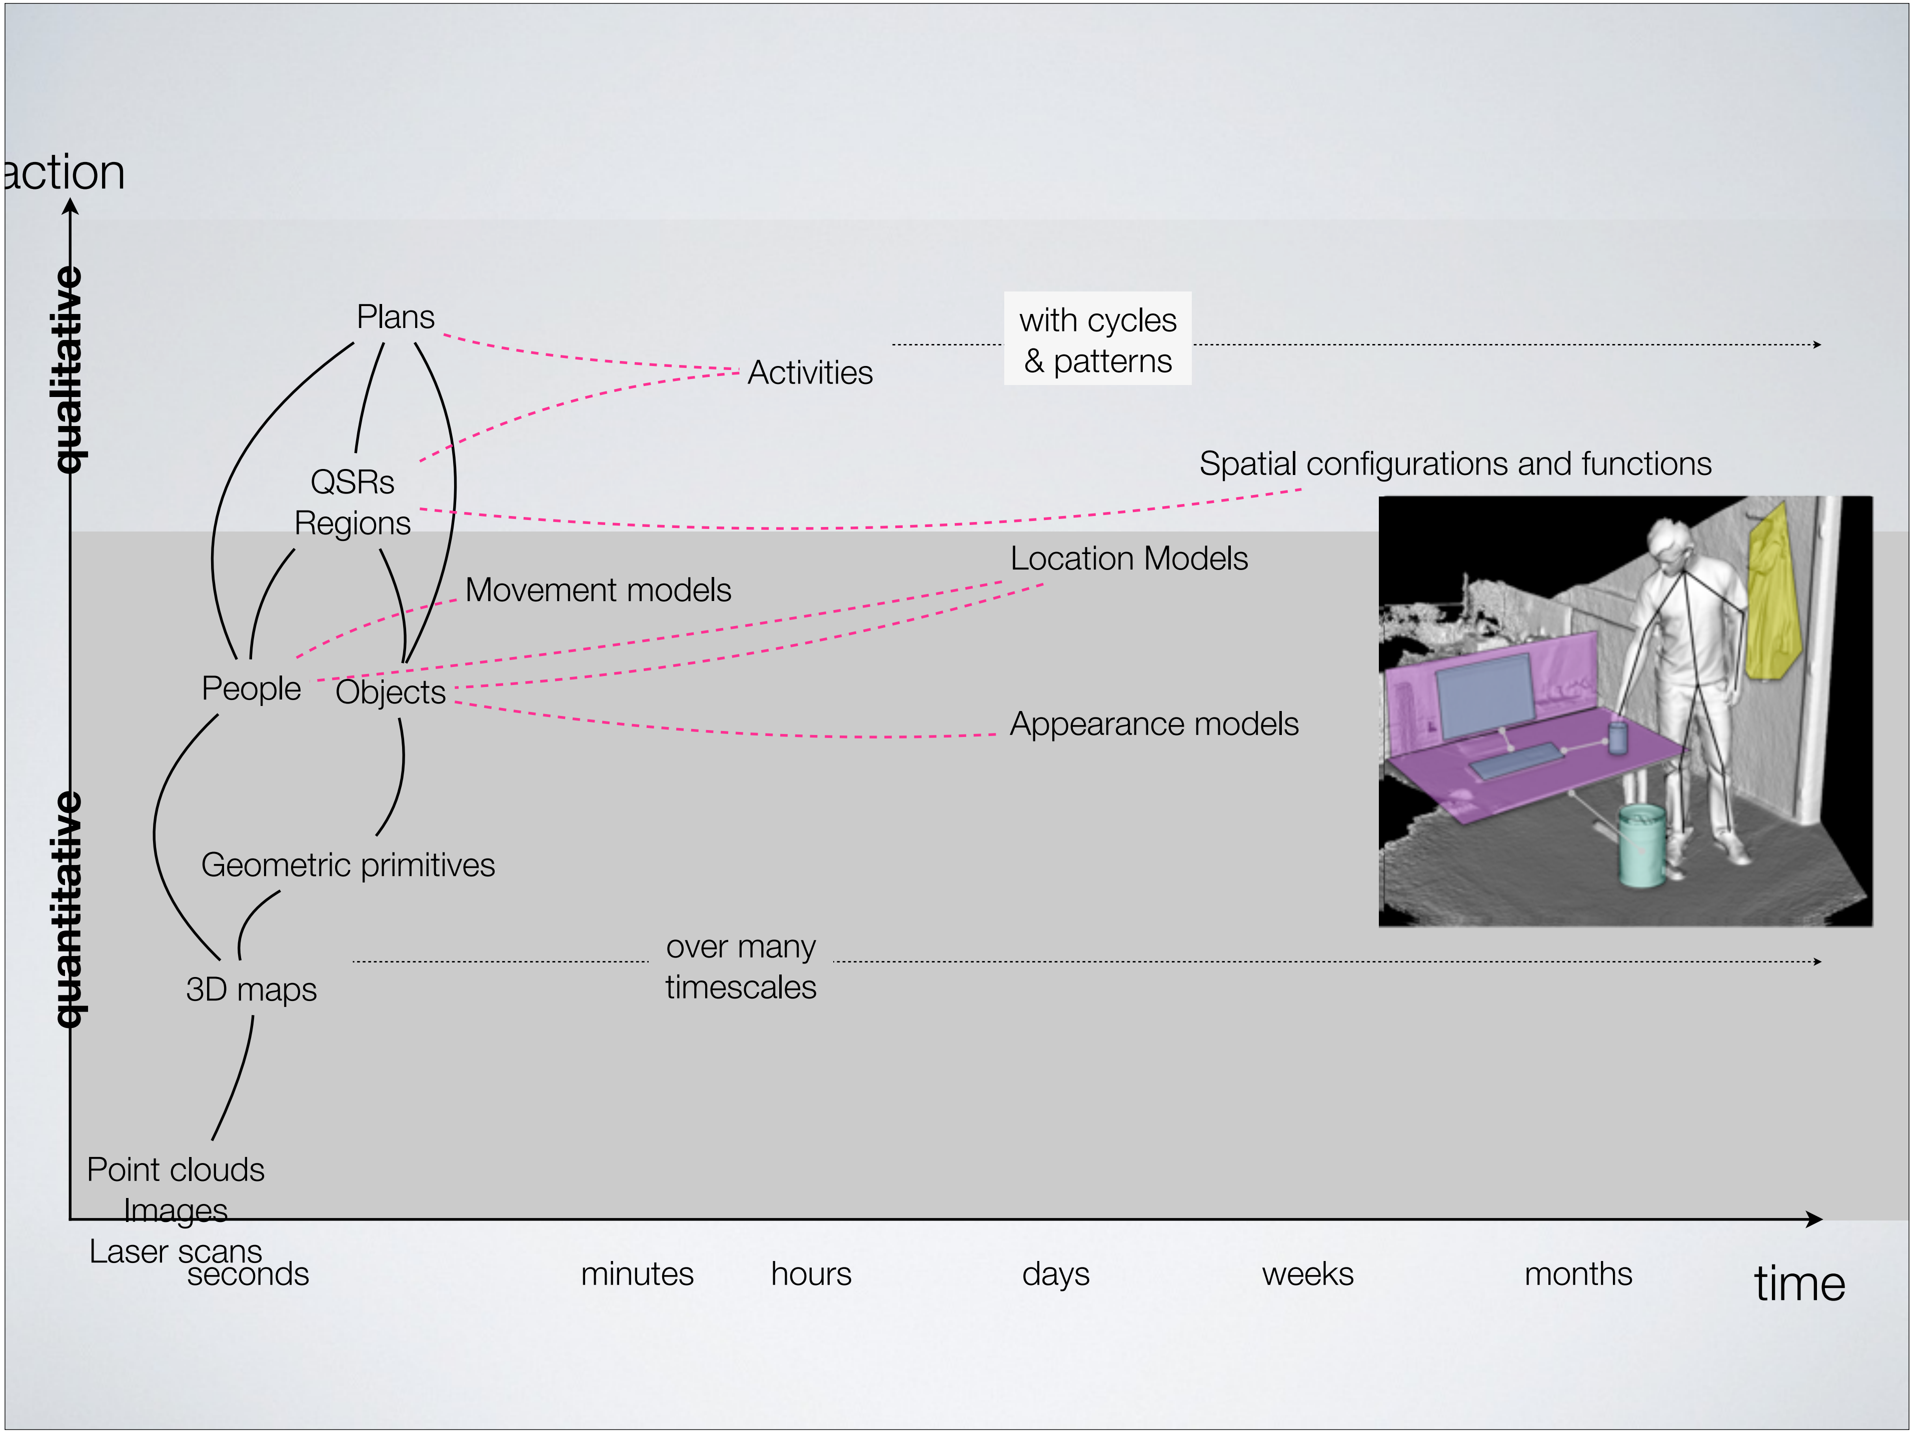
action
qualitative
quantitative
Plans
Activities
with cycles
& patterns
QSRs
Regions
Spatial configurations and functions
Location Models
Movement models
People
Objects
Appearance models
Geometric primitives
3D maps
over many
timescales
Point clouds
Images
Laser scans
seconds
minutes
hours
days
weeks
months
time

O1: A unified understanding of space *over time*

O2: Semantic segmentation of space

HR(S)I

O3: Understanding human activities

O4: Cognitive control of a robot's activities from spatio-temporal information

O5: Interpreting long-term experience from sparse observations

Robustness

O6: Integration and validation of a long-lived cognitive robot for dynamic, real-world tasks

THE STRANDS ROBOT MARATHON

Fork me on GitHub

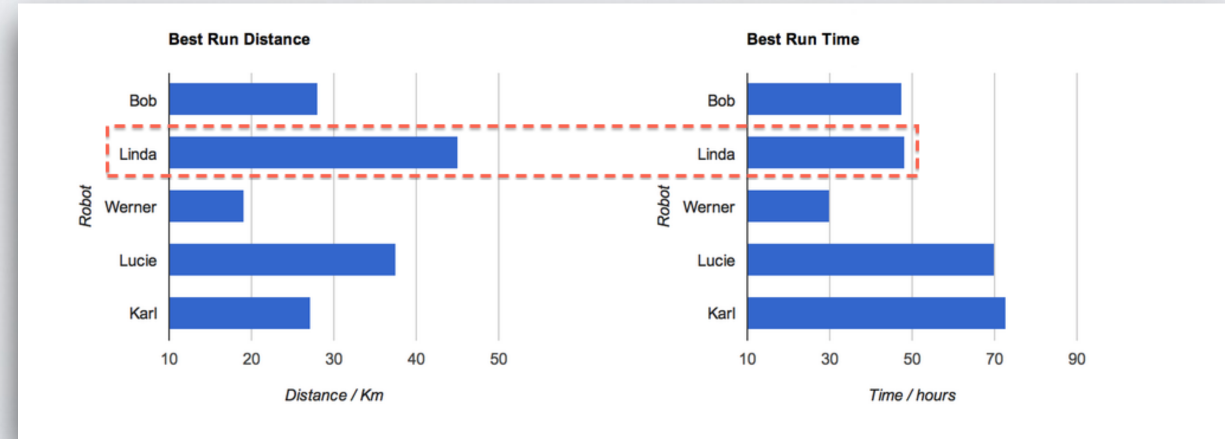

UNIVERSITY OF LINCOLN

https://github.com/strands-project/strands_systems/wiki/Hydro-Installation-Instructions

LINCOLN ROBOTICS

ROBOTS ARE DANGEROUS

- We need a safe robot software implementation
- We need robots that navigate safely among humans and real world
- We cannot do much testing before deployment

ROBOTS ARE ENDANGERED

UNIVERSITY OF
LINCOLN

LINCOLN
ROBOTICS

MORSE'S ROLE IN STRANDS

- cannot forgo any real-world studies, but
  - we need *robust* software implementations,
  - we need *tested* software implementations and deployment,
  - we need *ground truth performance* for comparison and algorithm quantification/bootstrapping

Continuous Integration and System Testing

Generating Dynamic Worlds

Testing long-term Human-Robot Spatial Interaction

MORSE'S ROLE IN STRANDS

- cannot forgo any real-world studies, but
  - we need *robust* software implementations,
  - we need *tested* software implementations and deployment,
  - we need *ground truth performance* for comparison and algorithm quantification/bootstrapping

Continuous Integration and System Testing

Generating Dynamic Worlds

Testing long-term Human-Robot Spatial Interaction

UNIVERSITY OF LINCOLN

LINCOLN ROBOTICS

TOWARDS CONTINUOUS LONG-TERM SYSTEM TESTING

- similar aims to Florian’s for Flobi
- but with the specific ambition to eventually run continuously
- go beyond compile-time and unit testing
- system architecture testing

FSMT/JENKINS @ STRANDS

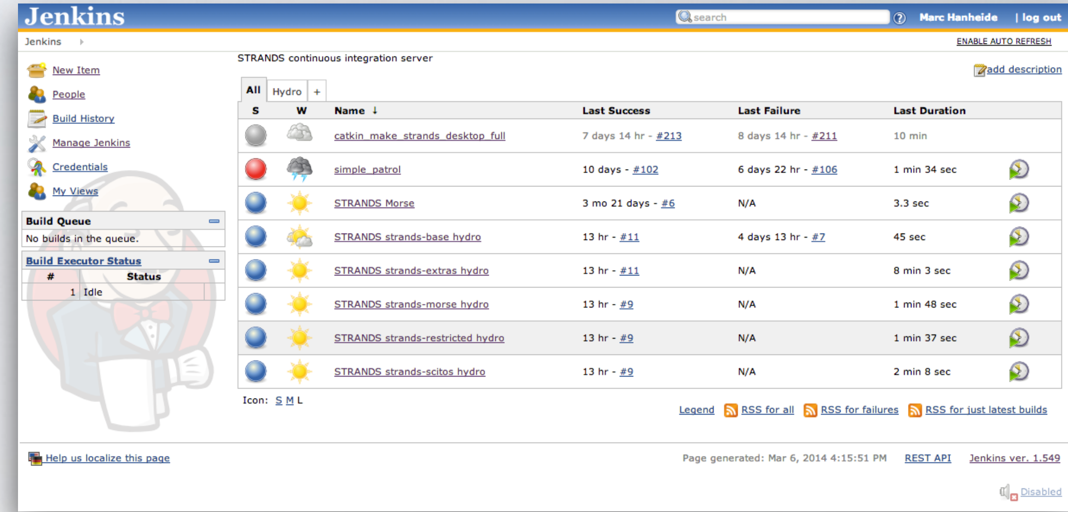

- based on FSMT
- Patrolling scenario (like in Robot Marathon)
- monitors ROS topics to check accomplishments of tasks (e.g. reaching goals), based on SMACH executor
- uses mongodb to store logs and selected topics
- runs in VirtualGL environment to record MORSE video for debugging

UNIVERSITY OF LINCOLN

LINCOLN ROBOTICS

MORSE'S ROLE IN STRANDS

- cannot forgo any real-world studies, but
  - we need *robust* software implementations,
  - we need *tested* software implementations and deployment,
  - we need *ground truth performance* for comparison and algorithm quantification/bootstrapping

Continuous Integration and System Testing

Generating Dynamic Worlds

Testing long-term Human-Robot Spatial Interaction

DYNAMICS IN REAL

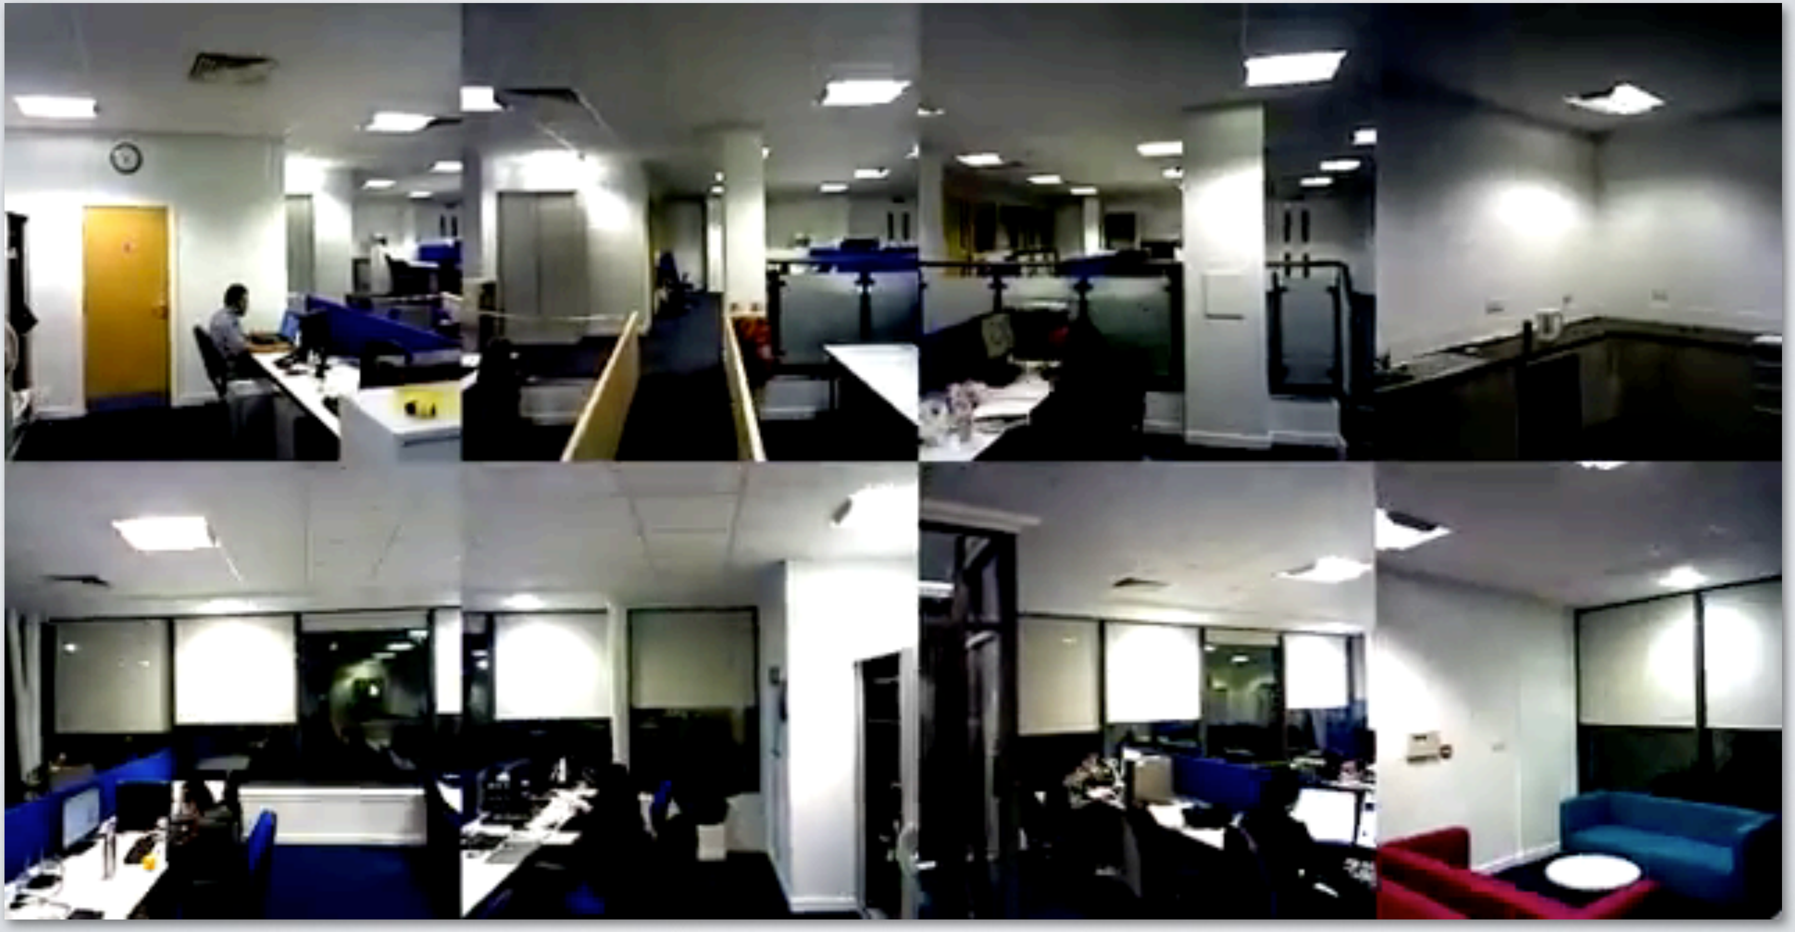

DYNAMIC WORLDS

- STRANDS sets out to develop "intelligent mobile robots that are able to run for months in ***dynamic*** human environments"
- generating controlled *random* worlds
- specific build extensions to sample worlds from probability distributions

UNIVERSITY OF LINCOLN

LINCOLN ROBOTICS

QUALITATIVE SPATIAL RELATIONS

$$
\begin{aligned}
&\mathit{scene}(\mathit{Monitor}, \mathit{Keyboard}, \mathit{Laptop}, \mathit{Cup}, \mathit{Bottle}) \Leftrightarrow \\
&\quad \mathit{in\text{-}front\text{-}of}(\mathit{Keyboard}, \mathit{Monitor}) \wedge \\
&\quad \mathit{left\text{-}of}(\mathit{Laptop}, \mathit{Keyboard}) \wedge \\
&\quad \mathit{right\text{-}of}(\mathit{Cup}, \mathit{Keyboard}) \wedge \\
&\quad \mathit{behind\text{-}of}(\mathit{Bottle}, \mathit{Cup}) \wedge \\
&\quad \mathit{close\text{-}to}(\mathit{Bottle}, \mathit{Cup}).
\end{aligned}
$$

- search objects for efficiently
- utilise 3D structures and probabilistic models of likely QSRs

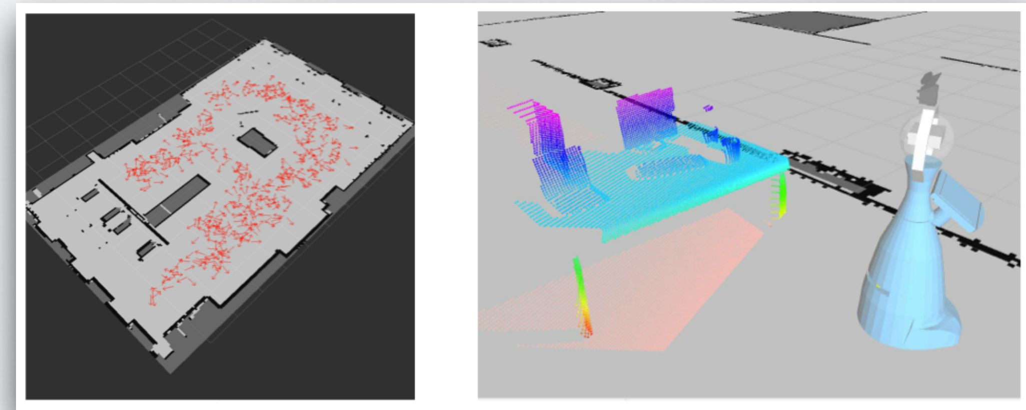

Kunze et al, "Using Qualitative Spatial Relations for Indirect Object Search", ICRA 2014

UNIVERSITY OF LINCOLN

LINCOLN ROBOTICS

QUALITATIVE SPATIAL RELATIONS

- Evaluation in MORSE, using sampling of scenes :
  - random positioning
  - supporting faces
  - QSR-based
- uses semantic camera
- performance comparable to real-world experiments

Search method	Found objects	Average time (sec)	Average poses
random	6/10	68.5	4.8
supporting planes	10/10	33.6	2.3
QSR	10/10	15.6	1.1

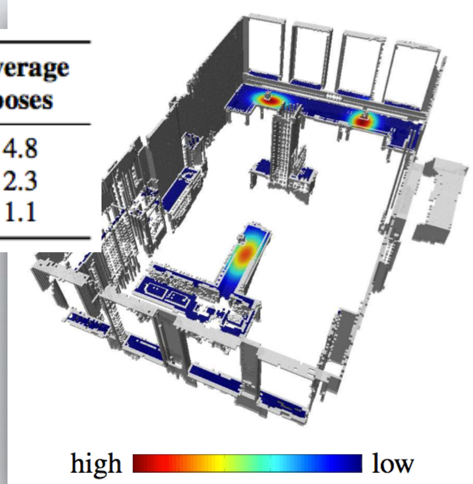

Object type	Location changes	Examples
static	rare and only quantitative	PC, monitor, printer
dynamic	frequent and qualitative	keyboard, mouse, cup

SCENES CAN BE DESCRIBED USING QUALITATIVE SPATIAL RELATIONS (QSRS)

QSR models abstract geometric information away and describe a scene qualitatively:

$$
\begin{aligned}
& scene(Monitor, Keyboard, Laptop, Cup, Bottle) \Leftrightarrow \\
& \quad \textit{in-front-of}(Keyboard, Monitor) \wedge \\
& \quad \textit{left-of}(Laptop, Keyboard) \wedge \\
& \quad \textit{right-of}(Cup, Keyboard) \wedge \\
& \quad \textit{behind-of}(Bottle, Cup) \wedge \\
& \quad \textit{close-to}(Bottle, Cup).
\end{aligned}
$$

Idea: Use qualitative descriptions (QSRs) to generate scenes of the same structure

```
# Create a root note
table = RootNode('table')

# Create several object nodes
pc1 = ObjectNode('pc1')
mon1 = ObjectNode('monitor1')
key1 = ObjectNode('keyboard1')
laptop1 = ObjectNode('laptop1')
cup1 = ObjectNode('cup1')
bottle = ObjectNode('bottle')

# Random rotations of objects
mon1.set_yaw(random.gauss(0.0,math.pi/16))
key1.set_yaw(random.gauss(0.0,math.pi/16))
cup1.set_yaw(random.uniform(0,2*math.pi))

# Add object nodes to the root node
table.add(pc1,'north_west')
table.add(mon1,'north')

# Add object nodes relative to other object nodes
mon1.add(key1,'front')
key1.add(cup1, 'right','close')
key1.add(laptop1,'left')
cup1.add(bottle, 'back','close')

table.place_objects()
```

Table

Objects

Rotations

Absolute positions

Relative positions

Fork me on GitHub

UNIVERSITY OF LINCOLN

https://github.com/strands-project/strands_morse

LINCOLN ROBOTICS

WE BOOTSTRAPPED A STATISTICS OF QSRS BY ANNOTATING IMAGES OF OFFICE DESKS (TO GENERATE SCENES WITH DIFFERENT STRUCTURES)



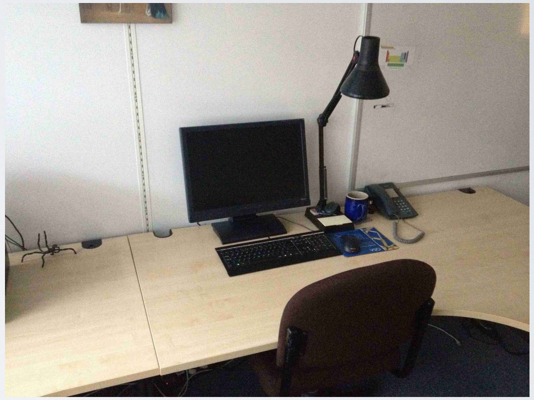
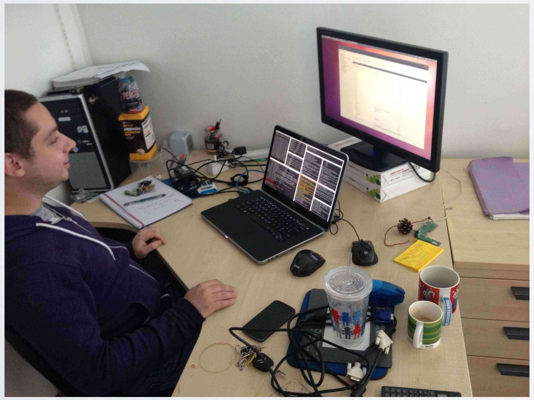

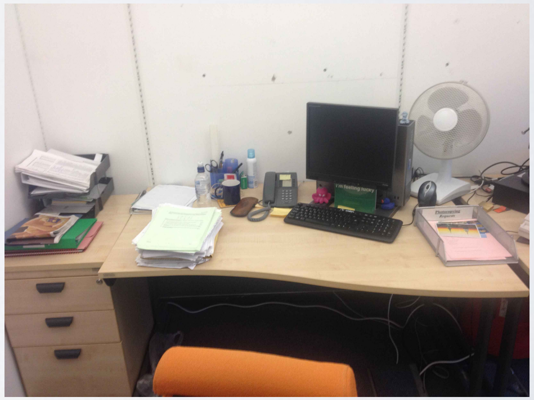

Q: WHERE IS A CUP WRT A MONITOR?

SAMPLING

Given a set of qualitative scene descriptions:

right-of(Cup, Monitor)

left-of(Cup, Monitor) and **in-front-of**(Cup,Monitor)

right-of(Cup, Monitor) and **close-to**(Cup,Monitor)

…

We generate object positions using the
Ternary Point Calculus [Moratz et al 2003]
by sampling a **relative angle** and **distance**

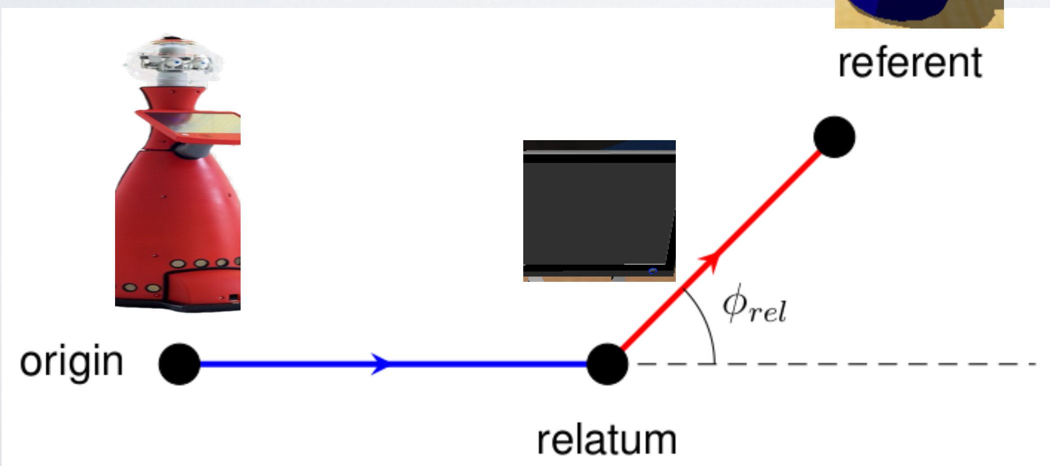

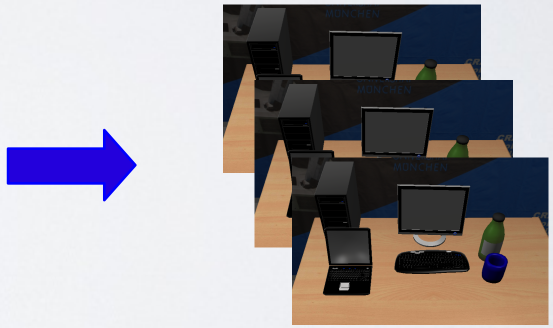

ORIENTATION

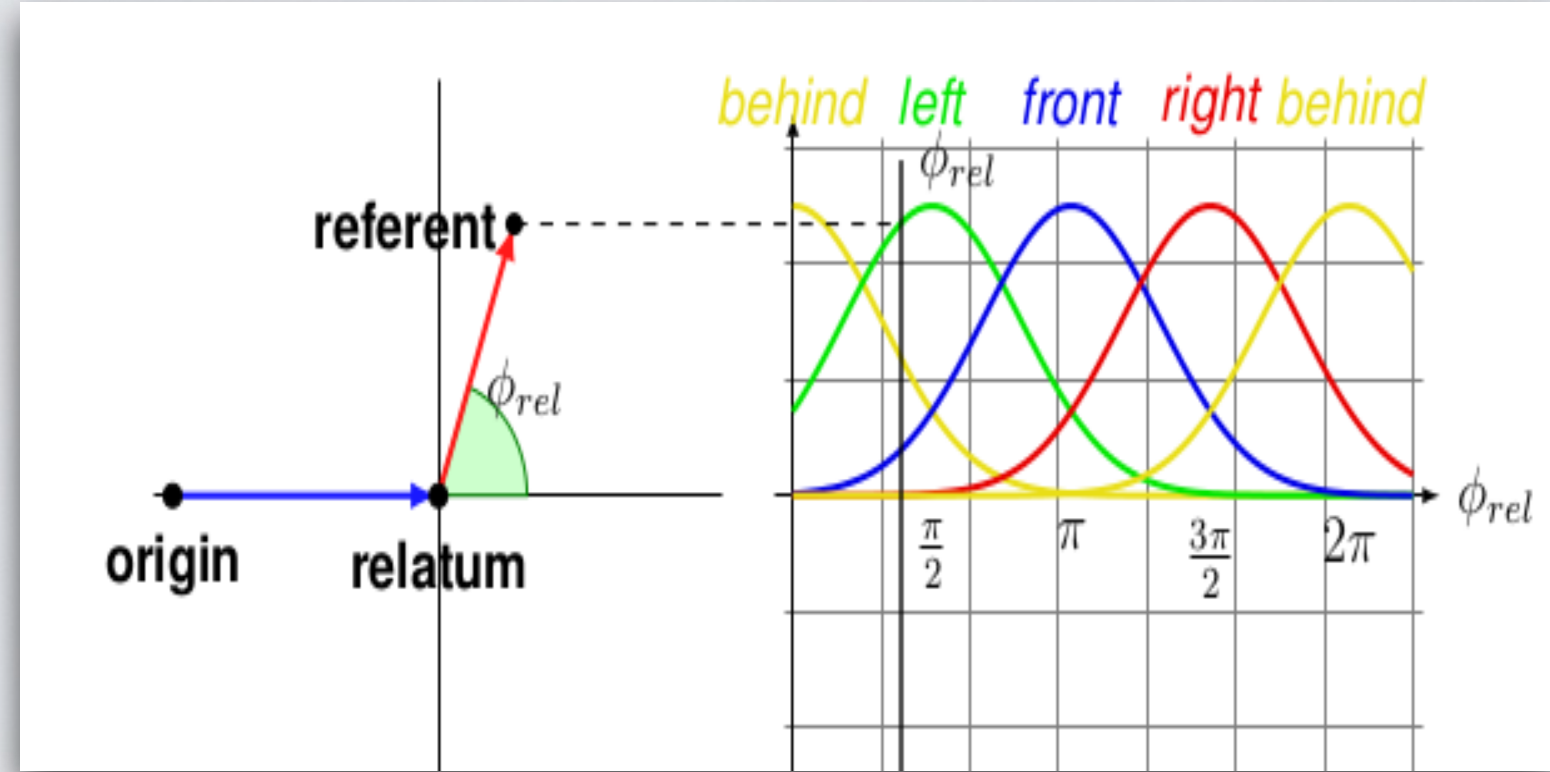

SIMILARLY, WE ALSO SAMPLE A DISTANCE RELATION (E.G. CLOSE)

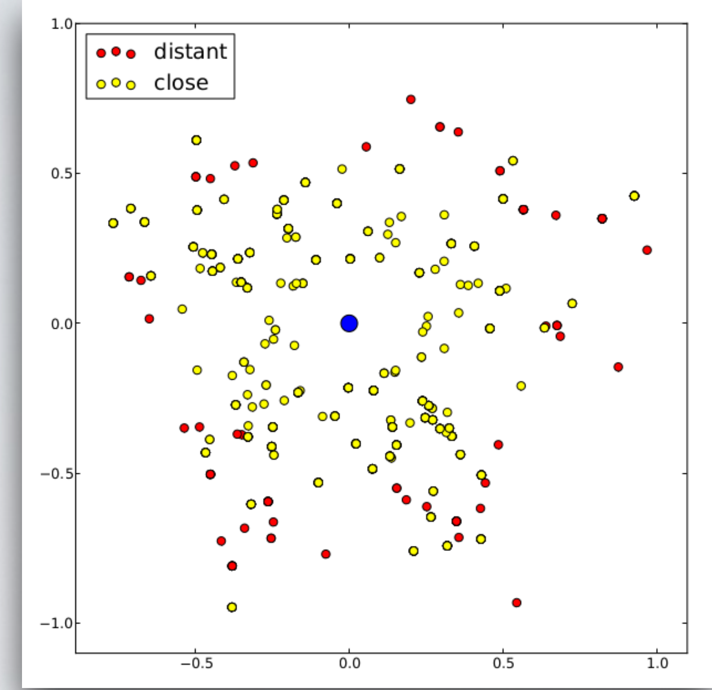

$$r_{rel} = \frac{distance(referent, relatum)}{distance(origin, relatum)}$$

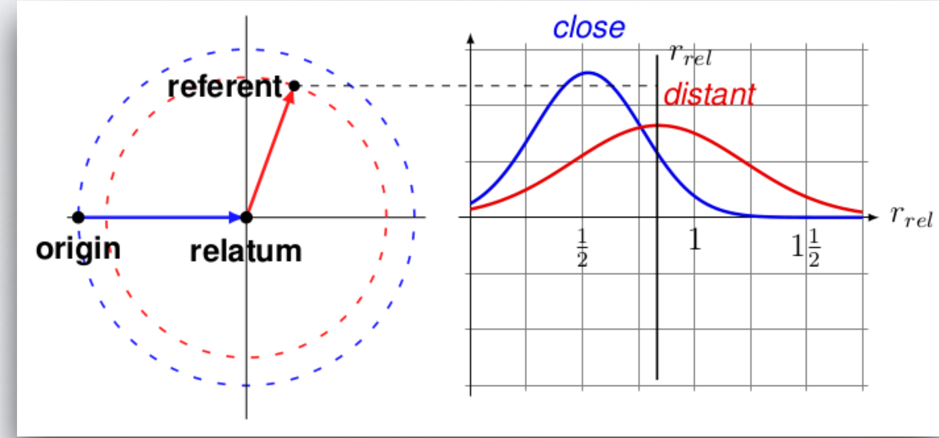

What about the context? For example: other objects, object size

THEN CHECK PHYSICAL CONSTRAINTS!

- avoid collisions
  - simply resample

SOME RESULTS...

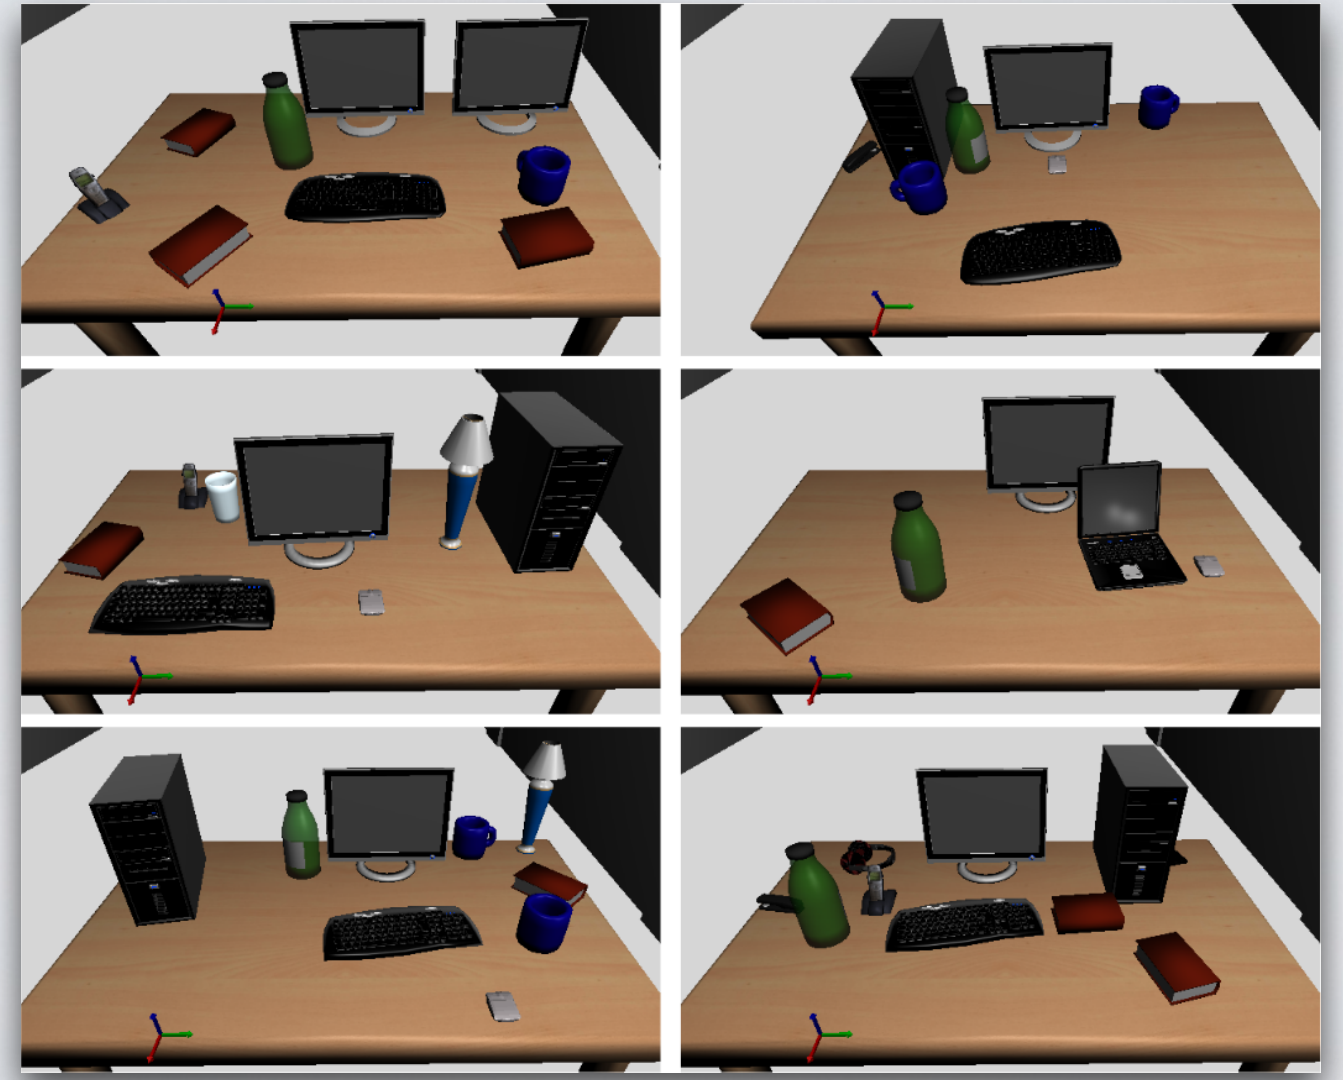

WHAT WE CAN DO WITH THIS BEYOND QSR EXPERIMENTS?

- algorithm testing for anomaly detection in “periodic worlds”
- exploring the boundaries of “dynamicness"

WHAT WE CAN DO WITH THIS BEYOND QSR EXPERIMENTS?

- algorithm testing for anomaly detection in "periodic worlds"

namicness"

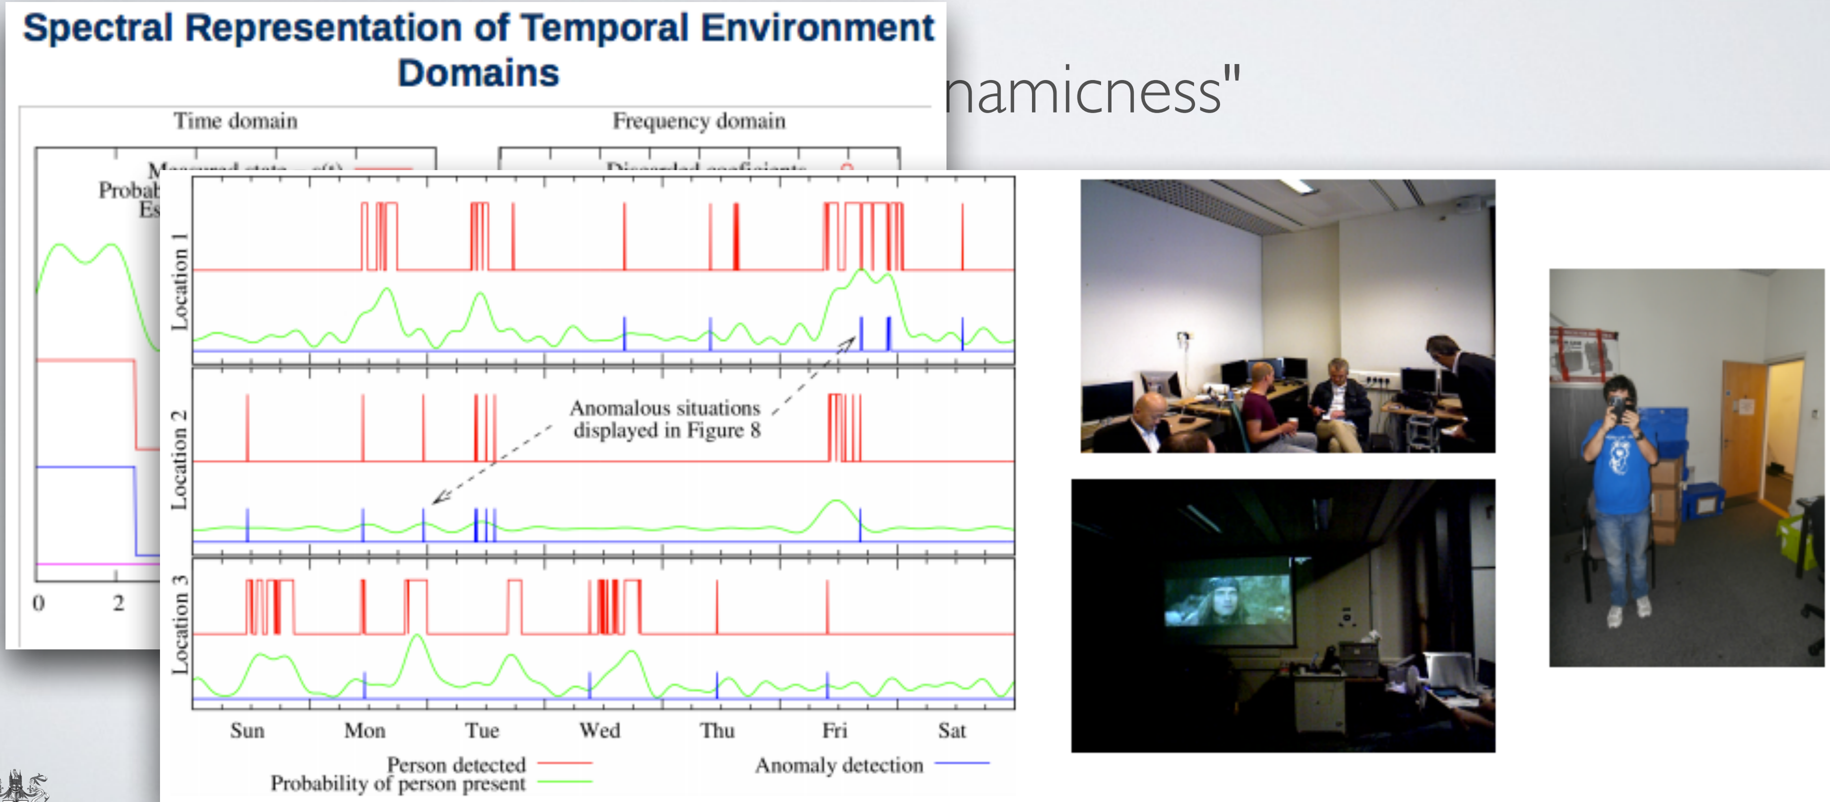

WHAT WE CAN DO WITH THIS BEYOND QSR EXPERIMENTS?

- algorithm testing for anomaly detection in “periodic worlds”
- exploring the boundaries of “dynamicness"

What is the localisation error, if… ?

MORSE'S ROLE IN STRANDS

- cannot forgo any real-world studies, but
  - we need *robust* software implementations,
  - we need *tested* software implementations and deployment,
  - we need *ground truth performance* for comparison and algorithm quantification/bootstrapping

Continuous Integration and System Testing

Generating Dynamic Worlds

Testing long-term Human-Robot Spatial Interaction

UNIVERSITY OF LINCOLN

LINCOLN ROBOTICS

Human-Robot Spatial Interaction (HRSI)

Prompting by Moving

Modelling HRSI

Predicting and Recognising HRSI Behaviour/Intentions

NARROW SPACES

Coordination or negotiation in
Human-Robot close encounters

Mutual Understanding of goals and intention

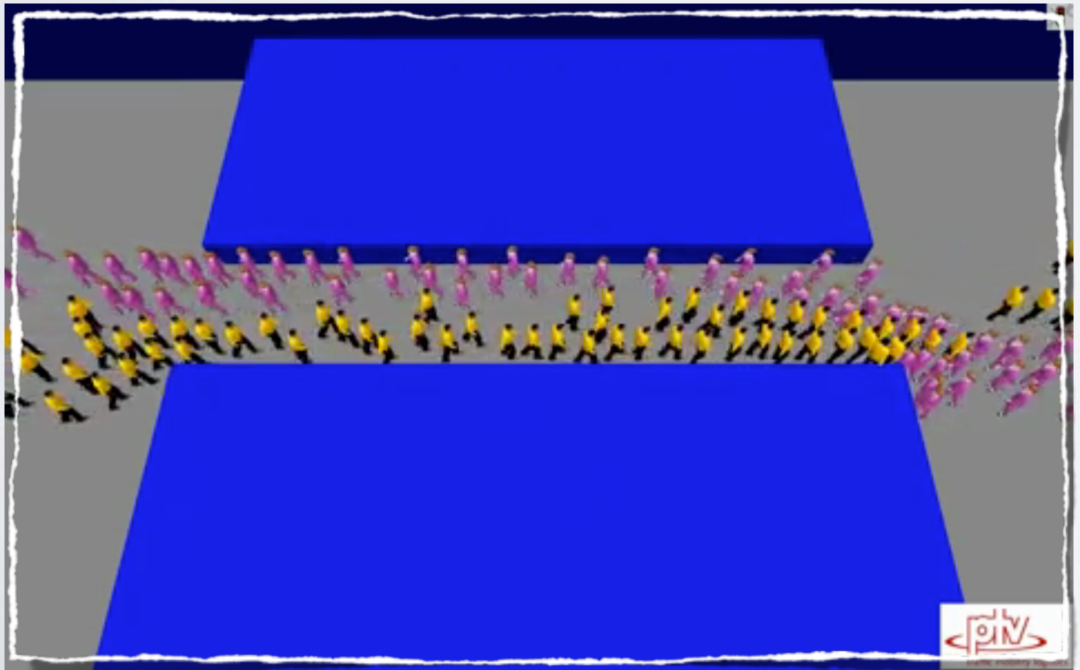

http://www.youtube.com/watch?v=QkVOjbRq-lo

Representing HRSI Situations and Behaviour

How to effectively capture and reason about joint movement of human and robot?

Statistical models for time series

A Qualitative Trajectory Calculus

Order matters, different relative movements need to be represented

like QSR for motion!

SIMPLIFIED QTC$_C$

Fork me on GitHub

q_1) movement of k with respect to l

- $-$: k is moving towards l
- 0 : k is stable with respect to l
- $+$: k is moving away from l

q_2) movement of l with respect to k

as above, swapping k and l

q_4) movement of k with respect to $\overline{k\,l}$

- $-$: k is moving to the left side of $\overline{k\,l}$
- 0 : k is moving along $\overline{k\,l}$
- $+$: k is moving to the right side of $\overline{k\,l}$

q_5) movement of l with respect to $\overline{l\,k}$

as above, swapping k and l

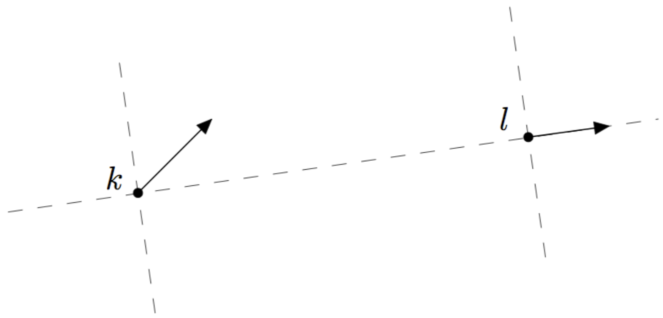

Figure 1: Example of moving points k and l. The respective QTC$_C$ relation is $(- + - 0)$.

- for 2D space
- 4-tupel per state, e.g. (-,+,-,0)
- 3^4=81 possible states

https://github.com/LCAS/qtc-toolbox

QTC$_C$ - BY EXAMPLE

QTC$_C$ represents the relative motion of two points in a time interval with respect to the reference line that connects them on a 2D plane.

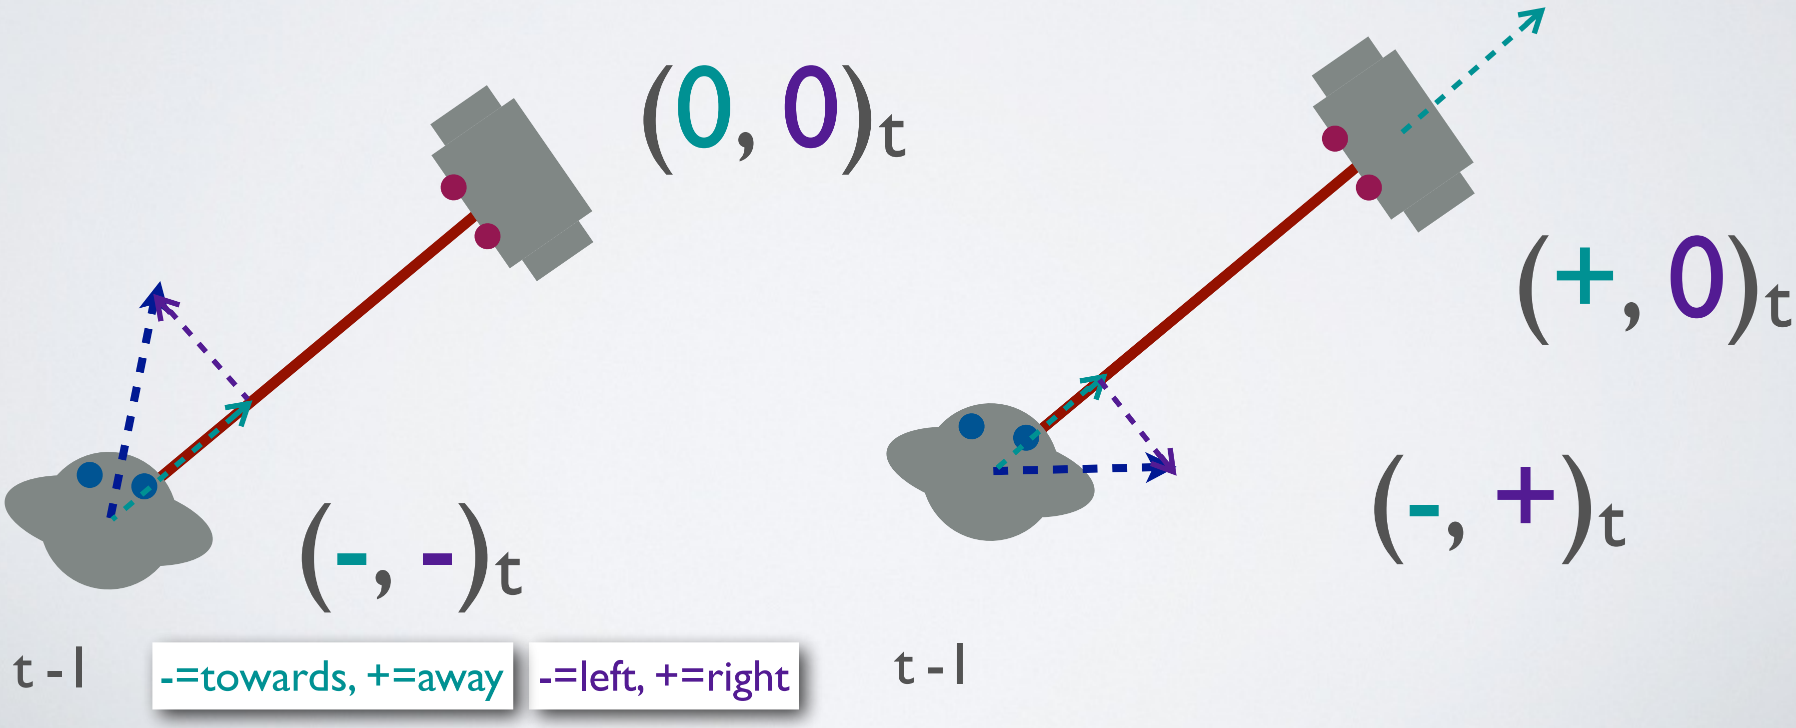

QTC$_C$ - BY EXAMPLE

$(-, 0, -, 0)_t$

H R H R

$(-, +, +, 0)_t$

H R H R

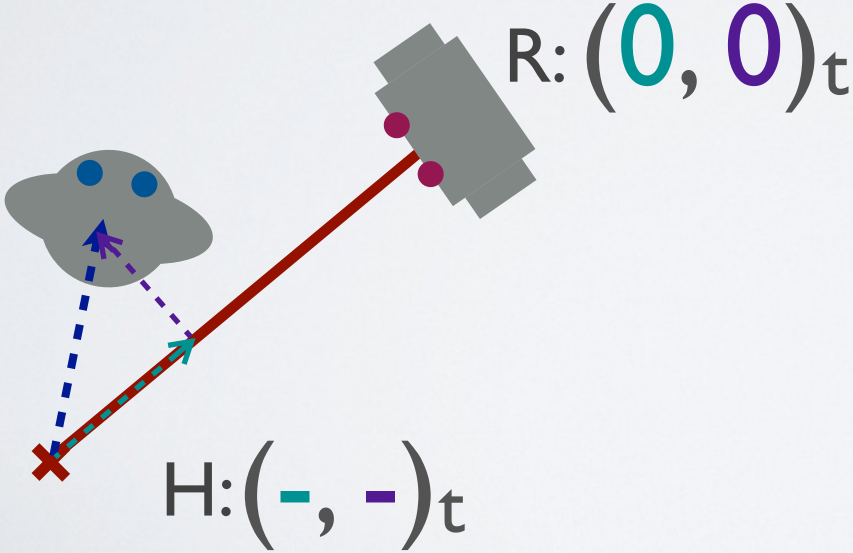

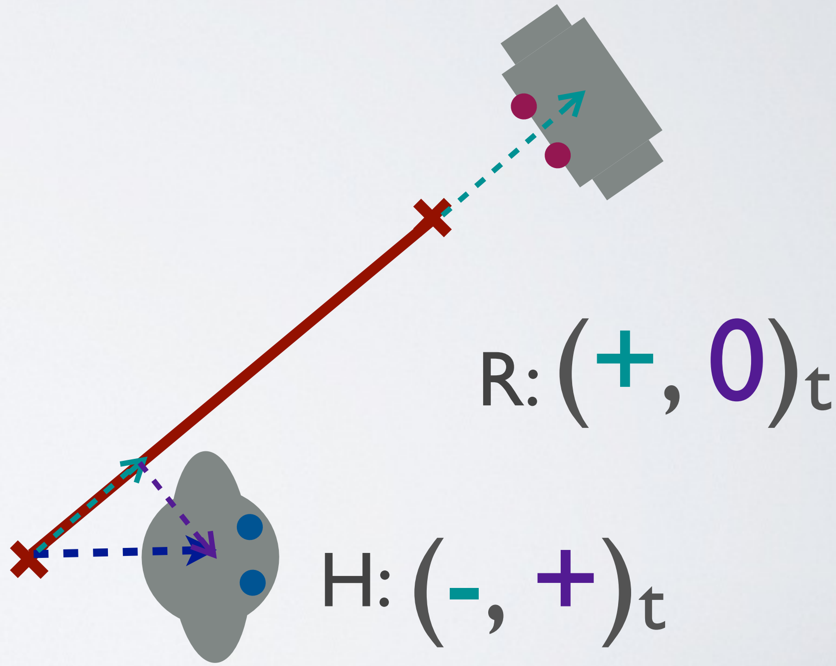

QTC$_C$ STATE

$(-, 0, -, 0)_t$

H R H R

state no. H R H R

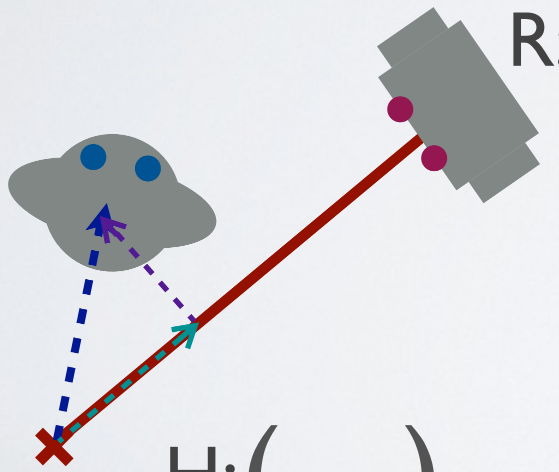

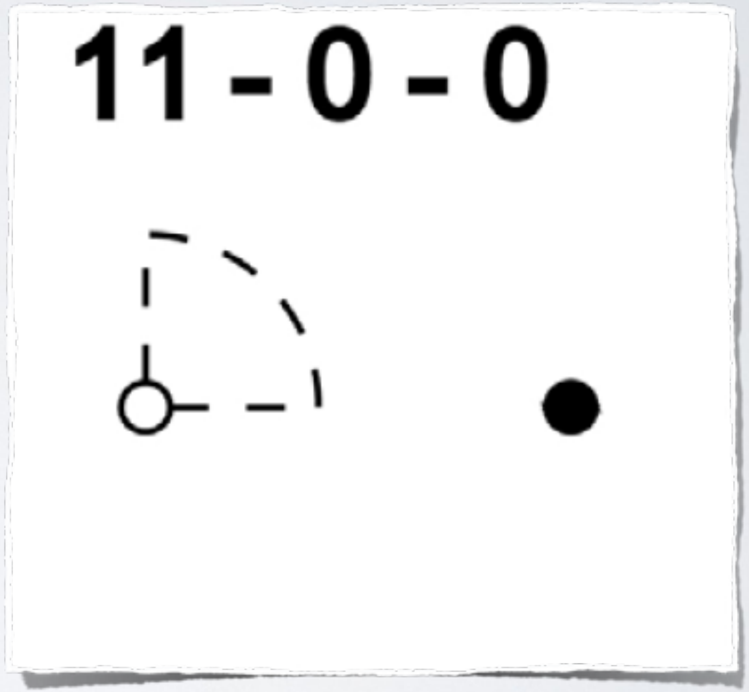

QTC$_C$ STATE

$(-, +, +, 0)_t$

H R H R

state no. H R H R

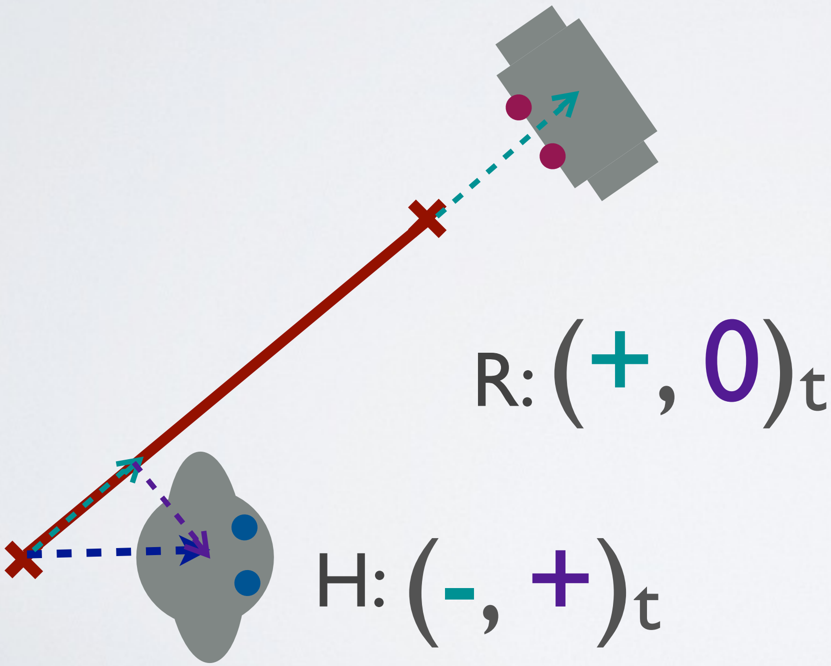

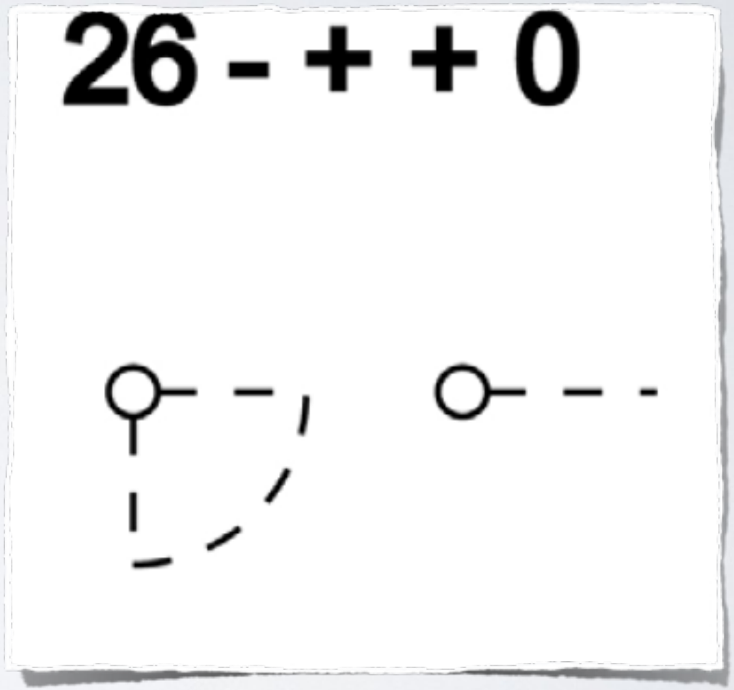

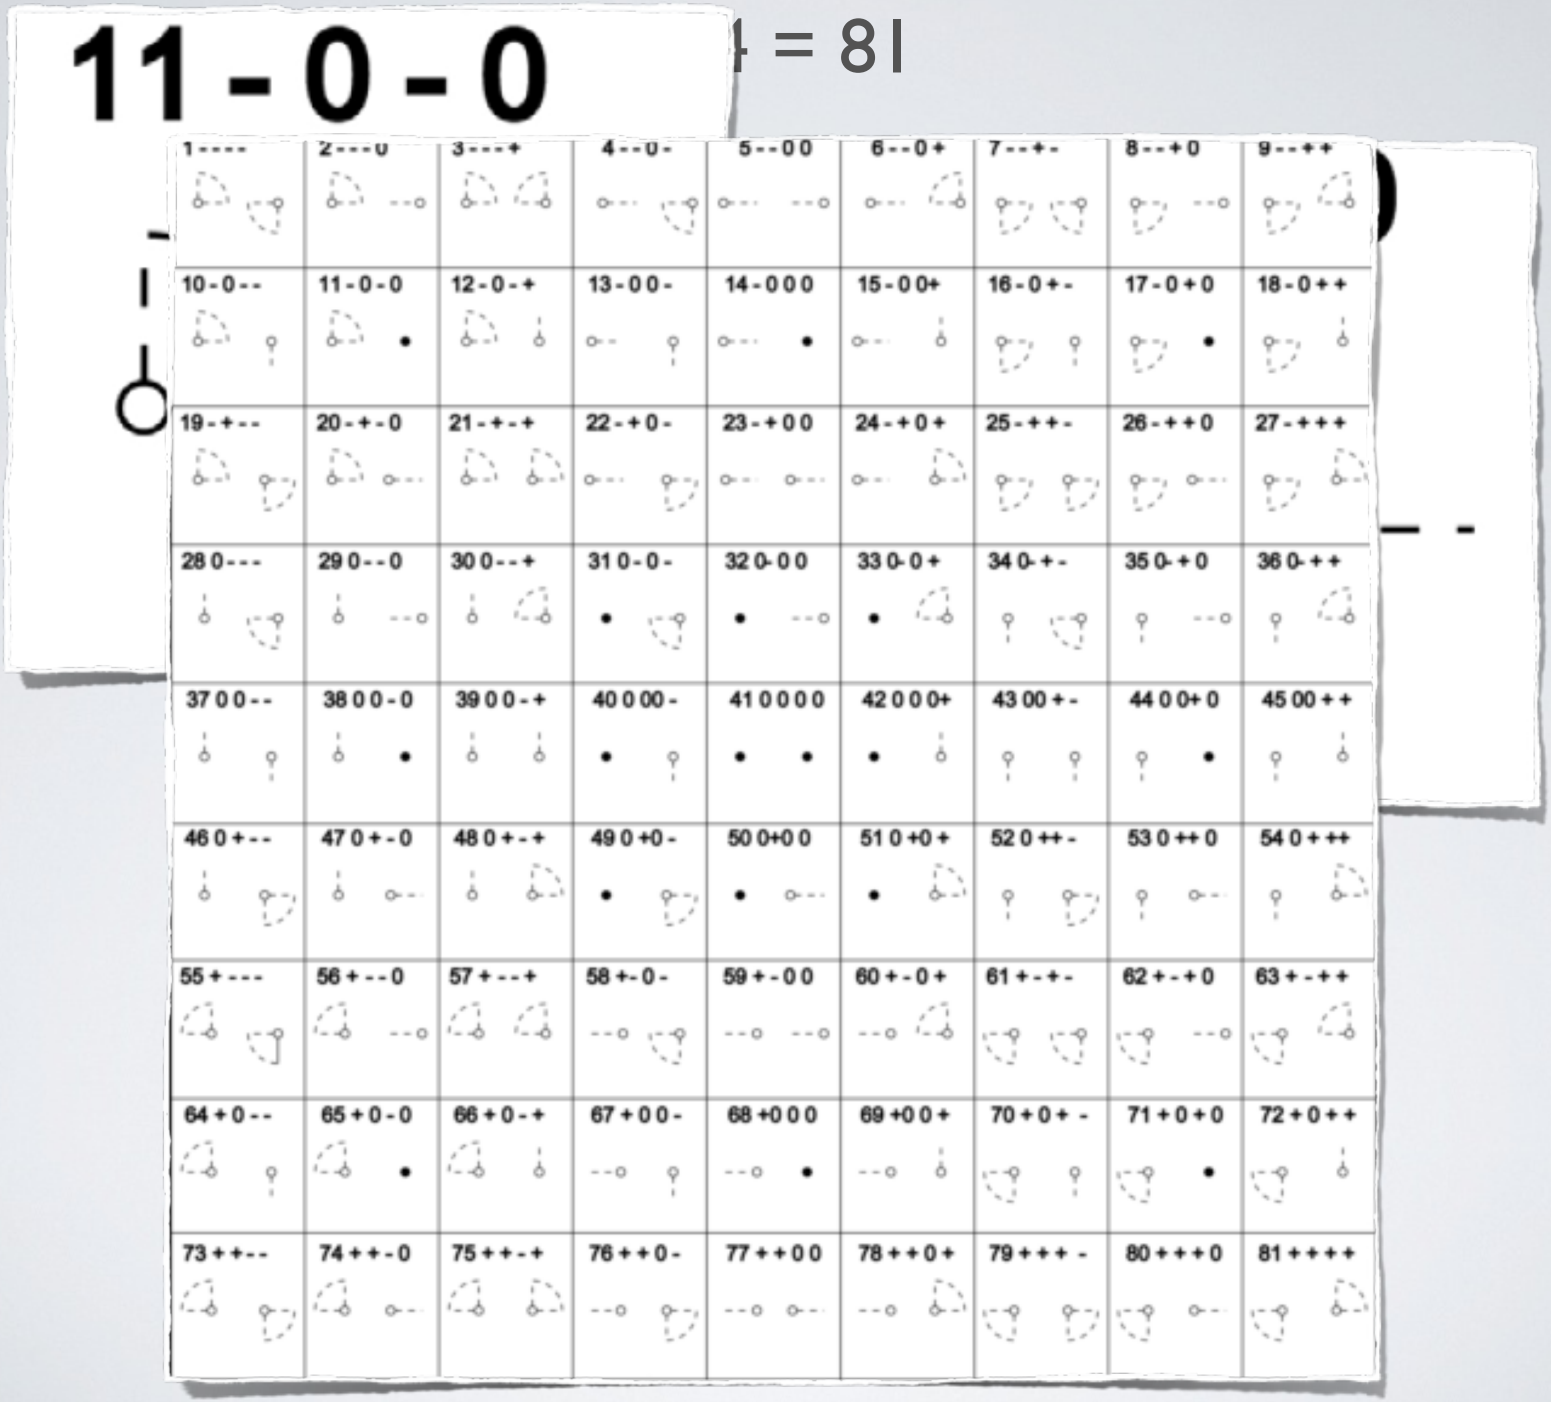
11 - 0 - 0
4 = 81
1 - - - -
2 - - - 0
3 - - - +
4 - - 0 -
5 - - 0 0
6 - - 0 +
7 - - + -
8 - - + 0
9 - - + +
10 - 0 - -
11 - 0 - 0
12 - 0 - +
13 - 0 0 -
14 - 0 0 0
15 - 0 0+
16 - 0 + -
17 - 0 + 0
18 - 0 + +
19 - + - -
20 - + - 0
21 - + - +
22 - + 0 -
23 - + 0 0
24 - + 0 +
25 - + + -
26 - + + 0
27 - + + +
28 0 - - -
29 0 - - 0
30 0 - - +
31 0 - 0 -
32 0- 0 0
33 0- 0 +
34 0- + -
35 0- + 0
36 0- + +
37 0 0 - -
38 0 0 - 0
39 0 0 - +
40 0 0 0 -
41 0 0 0 0
42 0 0 0+
43 0 0 + -
44 0 0+ 0
45 0 0 + +
46 0 + - -
47 0 + - 0
48 0 + - +
49 0 +0 -
50 0+0 0
51 0 +0 +
52 0 ++ -
53 0 ++ 0
54 0 + ++
55 + - - -
56 + - - 0
57 + - - +
58 +- 0 -
59 + - 0 0
60 + - 0 +
61 + - + -
62 + - + 0
63 + - + +
64 + 0 - -
65 + 0 - 0
66 + 0 - +
67 + 0 0 -
68 +0 0 0
69 +0 0 +
70 + 0 + -
71 + 0 + 0
72 + 0 + +
73 + + - -
74 + + - 0
75 + + - +
76 + + 0 -
77 + + 0 0
78 + + 0 +
79 + + + -
80 + + + 0
81 + + + +

REPRESENTING HRSI BEHAVIOUR AS QTC-MM

- Take the CND and create a (hidden) Markov Model topology from it
- Add start and stop states
- Discretise Motion of Human and Robot into QTC states
- Train with Baum-Welch Algorithm for set of specific behaviours

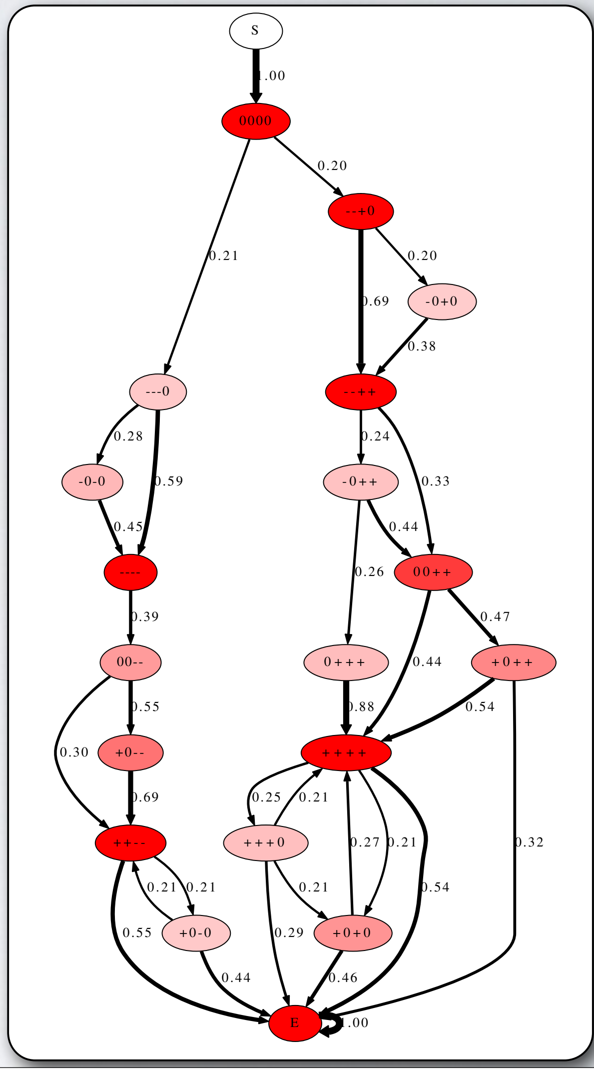

HRSI IN MORSE

- Human model is nice, but
  - …currently our advanced Kinect-based person detector is not directly applicable in MORSE
  - limited similarity to real-world ASUS (resolution, synchronisation)
- we can use modified Human model (with collisions for legs enabled) to track legs in a Bayesian Framework (https://github.com/LCAS/bayestracking)

T. Baumgartner, D. Mitzel, B. Leibe. "Tracking People and Their Objects" in *IEEE Conference on Computer Vision and Pattern Recognition*, 2013.

FROM INDIVIDUALS TO CROWDS

- so far we looked at individual humans,
- but we deal with this:
- and, for testing in simulation, these need to move realistically

LAYERED FLOW FIELDS

- individual goals
- physical
- navigational
- occupancy
- generate trajectories

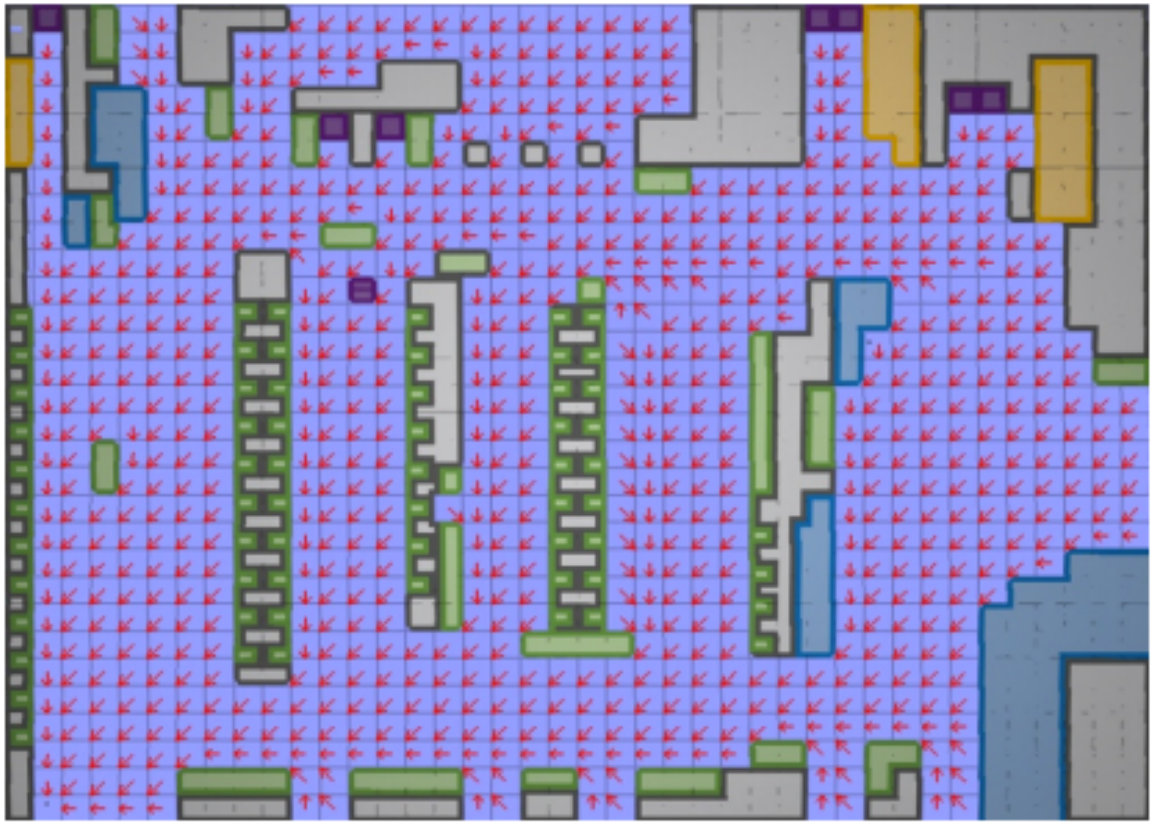

Fig. 2. Navigation flow-field produced by Dijkstra's algorithm, stored in the Navigation layer with goal the bottom-left corner.

UNIVERSITY OF LINCOLN

LINCOLN ROBOTICS

AT GATWICK SOUTH

Olivier Szymanezyk, Patrick Dickinson, Tom Duckett. Towards Agent-Based Crowd Simulation in Airports Using Games Technology. In proceeding of Agent and Multi-Agent Systems: Technologies and Applications - 5th KES International Conference, KES-AMSTA 2011

MULTI-HUMAN SIMULATION STILL CHALLENGING

- Performance! Simple human model needed?
- Integration into Morse as "active objects"
- Extend to changing environments
- Extensive testing for "realism" needed

Good News!
New Lincoln MSc student to work on this starting in April.

UNIVERSITY OF LINCOLN

LINCOLN ROBOTICS