Machine Learning para los Retos de la Robótica Social

Luis Merino Cabañas Service Robotics Laboratory Universidad Pablo de Olavide, Sevilla

http://robotics.upo.es

Next step: robots in our daily lives







People Perception

Haru (春): Tabletop Social Robot









The Effects of Robot Cognitive Reliability and Social Positioning on Child-Robot Team Dynamics

V. Charisi, L. Merino, M. Escobar, F. Caballero, R. Gomez and E. Gómez ICRA 2021

Haru's People Perception Pipeline













Multi-Modal People Fusion

- A method to integrate global features into a unified view of the persons in the scene
- Combination of multi-modal global features employing vision and sound
- Extensible data fusion architecture based on a cascaded of filters and pipelines
- Computation of contextual features, like proxemics and Fformations

Ragel, Ricardo, et al. "Multi-modal Data Fusion for People Perception in the Social Robot Haru." *International Conference on Social Robotics*. Cham: Springer Nature Switzerland, 2022.



Multi-Modal People Fusion







F-formations classification



0.00 0.05 0.10 0.15 Características

People detection with knee-high 2D LiDAR



- Motivation: Low cost + Ubiquity of 2D LiDAR
- Input data: Range data vector (1D, distance in meters for each angle)
- Challenges: Few datasets, complex problem, non-deep processing
- UPO developed dataset: FROG Real Alcázar de Sevilla











People detection with knee-high 2D LIDAR

- UPO end-to-end models:
 - People Proposal Network
 - Laser Feature Extractor



People detection with knee-high 2D LiDAR





		0.5 m			0.3 m		1
	AP	Peak F_1	EER	AP	Peak F_1	EER	Time
ROS leg_detector	20.0	35.2	34.9	10.5	24.3	24.1	1.77ms
PeTra	47.9	66.6	66.4	47.5	66.1	65.9	28.17ms
PeTra*	59.1	67.9	67.3	58.7	67.4	66.8	_
LFE-Peaks (ours)	62.4	70.7	70.7	61.4	69.0	69.0	1.76ms
LFE-PPN (ours)	65.6	68.4	68.2	59.2	65.8	65.4	1.49ms
DROW3 $(T = 1)$	72.3	72.0	71.8	71.8	71.6	71.4	13.08ms
DR-SPAAM $(T = 1)$	72.2	72.1	71.8	71.6	71.6	71.3	13.95ms
DR-SPAAM $(T = 5)$	74.2	73.6	73.4	73.7	73.2	73.0	13.99ms

Fernando Amodeo, Noé Perez-Higueras, Luis Merino, Fernando Caballero, FROG: A new people detection dataset for knee-high 2D range finders, 2023 <u>https://github.com/robotics-upo/2DLaserPeopleBenchmark</u> <u>https://arxiv.org/abs/2306.08531</u>

Context Estimation

Never Home Alone













LIGHIHUUSE DISRUPTIVE INNOVATION GROUP

NHoA

- Goal: Automatic scene description by social robots
- Motivation:
 - Description of scenes for visually impaired
 - Feeding robot knowledge bases for semantic perception



Existing work



- Datasets:
 - Visual Relationship Dataset (VRD)
 - Visual Genome (VG) and its many derivatives
 - Models: VRD-RANS, VR-LP, VTE, IMP, Motifs, ...
- Problems:
 - Biased, noisy, overly general datasets
 - Unrealistic images



Qualitative tests







Person 1 is sitting on chairs 1, 2, 3, 4, 5 and 6. Person 2 is sitting on chairs 1 and 2.





Persons 3, 4, 5 are sitting at table 1. Person 2 is next to Person 3 and 4. Person 3 is next to Person 2. Person 4 is next to Person 6. Person 5 is next to Person 6. Person 6 is next to Person 2 and 4.

Main proposals



- Specific Idea: Design and train a scene graph generation ML model taking into account domain specific information
- Prior domain knowledge: Ontology
- Datasets:
 - Filter and preprocess an existing dataset using ontology \checkmark
- Process:
 - Adapt and reimplement an existing model architecture \checkmark
 - Filter impossible outputs according to ontology \checkmark

Fernando Amodeo, Fernando Caballero, Natalia Díaz-Rodríguez, Luis Merino (2022) OG-SGG: Ontology-Guided Scene Graph Generation. A Case Study in Transfer Learning for Telepresence Robotics. ArXiv https://github.com/robotics-upo/og-sgg



Fernando Amodeo, Fernando Caballero, Natalia Díaz-Rodríguez, Luis Merino (2022) OG-SGG: Ontology-Guided Scene Graph Generation. A Case Study in Transfer Learning for Telepresence Robotics. ArXiv https://github.com/robotics-upo/og-sgg

Scene Graph Generator





Test dataset: TERESA

- Scenario: Telepresence robot for elderly care
- Simple ontology
- Domain/Range restrictions, Inverse/Functionality/Transitivity/etc
- 25 images
- Object/predicate annotations (with help from ontology)







Quantitative tests



	Metrics for Predicate Detection (PredDet)												
Dataset	Post	R@K (k = 1)		R@K (k = 8)		mR@K (k = 1)		mR@K (k = 8)					
Dataset		20	50	100	20	50	100	20	50	100	20	50	100
VG-SGG unmodified dataset	X	27.0	34.7	41.9	23.8	34.8	51.1	19.1	30.6	36.4	29.3	42.7	57.0
	1	28.4	36.0	43.2	29.5	42.2	57.9	32.4	44.6	51.1	33.5	49.7	63.0
VG-SGG with filtering	X	40.0	43.4	47.7	42.0	49.0	60.0	42.2	48.8	50.9	43.5	51.5	57.4
	1	44.7	47.9	53.2	46.5	53.4	66.3	44.0	51.2	53.6	44.7	53.9	60.5
VG-SGG with filtering/augmentation	X	39.8	43.9	48.6	40.7	49.8	61.5	42.5	47.1	50.6	43.4	51.2	57.9
	1	44.9	49.3	54.0	46.5	54.0	68.3	44.6	49.9	53.3	45.2	53.0	61.7
VG-indoor unmodified dataset	X	26.1	33.7	40.2	26.0	41.4	56.1	10.5	20.6	25.2	19.8	35.5	53.6
	1	30.7	38.4	45.2	32.7	45.6	59.6	22.0	39.5	44.7	23.3	39.5	58.1
VG-indoor with filtering	X	41.2	41.3	46.7	42.2	48.1	59.2	43.7	50.6	45.5	44.7	56.0	65.8
	1	43.6	45.3	51.8	44.2	51.5	62.8	45.0	52.9	59.6	47.1	57.7	69.3
VG-indoor with filtering/augmentation	X	42.1	42.6	46.9	42.2	48.5	58.1	45.5	53.0	56.1	46.2	56.8	63.8
	1	44.4	46.1	51.9	43.3	51.5	62.0	46.8	54.8	61.0	46.6	58.3	66.6

Notice that the test set is completely different from the training set!

Qualitative tests





Person 1 is sitting on chairs 1, 2, 3, 4, 5 and 6. Person 2 is sitting on chairs 1 and 2.

sittir

sitting o

sitting

Persons 3, 4, 5 are sitting at table 1.

Person 2 is next to Person 3 and 4. Person 3 is next to Person 2.

Person 4 is next to Person 6.

Person 5 is next to Person 6.

Person 6 is next to Person 2 and 4.

next t

person₂ person

sitting

next

window₁



Person 1 and 2 are in front of a door. Person 1 is holding a piece of paper. Person 1 is next to chair 1. Person 2 is holding a piece of paper.



Persons 1 and 2 are in front of a window. Persons 2, 3, 4, 5 and 6 are in front of a window. Person 2 is holding a cup.



Human-aware navigation

Human-aware path planning



- Encoding social norms into path/motion planners
- Social navigation tasks are easier to demonstrate than formalise
- Given demonstrations by humans: Learning by Demonstration





Learning navigation cost functions from demonstrations



- Inverse Reinforcement Learning (IRL) approach
- Given successful demonstrations, estimate cost function thee demonstrator is minimising





- Initial formulations using discrete Markov Decision Processes (MDPs) as planners
- MDPs defined by the tuple <S,A,T,R>:
 - S: state space
 - A: action space
 - T: transition function, T(x',a,x)=p(x'|x,a)
 - R: reward (cost) function, R(x,a)
- Solving a MDP is finding a policy $\pi^*(x)$ that maximizes (minimizes) cumulative reward (cost)



- IRL as an "incomplete" MDP, <S,A,T>
- Given a set of demonstrations by an expert, determine the cost (reward) function used by the demonstrator
 - Cost function as a linear combination of features:

$$c(x) = \sum_{j} \omega_{j} f_{j}(x) = \omega^{T} f(x)$$





- IRL as an "incomplete" MDP, <S,A,T>
- Given a set of demonstrations by an expert, determine the cost (reward) function used by the demonstrator
 - Cost function as a linear combination of **features**:

$$c(x) = \sum_{j} \omega_{j} f_{j}(x) = \omega^{T} f(x)$$

• Determine weights so that the planner returns the same features as the demonstrations in expectation (Abbeel and Ng, 2004)

$$\mathbb{E}(f(\zeta)) = \frac{1}{D} \sum_{i=1}^{D} f(\zeta_i)$$

D

D

 $f(\zeta_i)$

Maximum Entropy Inverse Reinforcement Learning

- Challenges:
 - Expert suboptimal
- Maximum Entropy Assumption (Ziebart et al., 2008)
 - Probability distribution on paths that does not exhibit any additional preference except feature expectation

$$p(\zeta|\omega) = \frac{1}{Z(\omega)} e^{-\omega^T f(\zeta)}$$

Issues with discrete MDPbased IRL

- Discretization
- Scalability with the state space
 - It renders difficult to include new information about the problem
 - For instance, reasoning about static obstacles, goals, etc
- However, good motion planners available (RRT* for instance)
 - They reason about obstacles, they can reason about dynamics and kinematics of the robot, etc

Learning from demonstrations with RRT*

Learning from demonstrations SRL with RRT*

Maximum Entropy IRL for continuous spaces (Kretzschmar et al., 2014)

$$p(\zeta|\omega) = \frac{1}{Z(\omega)} e^{-\omega^T f(\zeta)}$$

• Determine the weights that maximize the (log-)likelihood of the demonstrations

$$\mathcal{L}(\mathcal{D}|\omega) = -D\log(Z(\omega)) + \sum_{i=1}^{D} (-\omega^{T} f(\zeta_{i}))$$

• The gradient is given by:
$$\frac{\partial \mathcal{L}(\mathcal{D}|\omega)}{\partial \omega} = D(\mathbb{E}(f(\zeta)) - \frac{1}{D} \sum_{i=1}^{D} f(\zeta_{i}))$$

Learning from demonstrations SRL with RRT*

$$\mathbb{E}(f(\zeta)) = \frac{1}{Z(\omega)} \int f(\zeta) e^{-\omega^T f(\zeta)} d\zeta$$

- Approximate it by the features of the best RRT* trajectory (minimum cost)
- RRT* is asymptotically optimal:
 - Run several iterations to account for variations

Learning from demonstrations with RRT*

Learning from demonstrations with RRT*

UNIVERSIDAD

Service Robotics Lab

ABLO

Noé Pérez-Higueras, Fernando Caballero, and Luis Merino. Teaching robot navigation behaviours to optimal RRT planners. In *International Journal of Social Robotics*, 2018.

https://github.com/robotics-upo/upo_nav_irl

Applying IRL to human-aware motion planning

- Demonstrations: human navigation datasets
 - ETHZ BIWI Walking Pedestrians dataset

Applying IRL to human-aware motion planning

Features for social navigation

- Using Deep Networks for path planning directly from sensor input
- Path planning as a classification problem
 - Use a Fully Convolutional Network (FCN) to directly predict the path given an scenario

FCN input:

200x200 px
res: 0.05 m/px
Robot-centered

Labels for training: - 200x200 px - res: 0.05 m/px - robot-centered

Labels

Validation of the FCN

Validation of the FCN

FCN-RRT*

RLT

3

RTIRL

Distance metric u

2

Sets

1

0.6

0.5

Distance metric (m) 0.3 0.3 2.0

0.

N. Pérez-Higueras, F. Caballero and L. Merino, Learning Human-Aware Path Planning with Fully Convolutional Networks. In *Proceedings of the IEEE International Conference on Robotics and Automation*, ICRA, 2018

https://github.com/robotics-upo/upo_fcn_learning

Robot Expressivity

Haru's Expressivity

- amused_3
- cheerful_3
- curious_3
- grumpy_3
- guilty_3
- sympathetic_3
- worried_3
- amused_7
- cheerful_7
- guilty_7
- sympathetic_7

- worried_7
- amused_10
- cheerful_10
- curious_10
- ecstatic_10
- guilty_10
- listening_10
- sympathetic_10
- listening_0
- listening_1
- listening_2

- thinking_0
- thinking_1
- thinking_2
- waking_up_1
- waking_up_2
- grumpy_7

Expressivity through motion

- Given animations designed to convey expressions like agreement or disagreement, happiness, sadness or shyness.
- We consider the question if such animations can be combined so to generate automatically expressions like, for instance, a shy disagreement.
- How the motion trajectories should be interpolated:
 - The manifold of motions representing a particular expression is complex.
 - Learning approaches can be applied to extract the relevant information, but only a handful of animations are available.

Encoding the demonstrated trajectories

- Each animation includes a demonstrated trajectory:
 - Temporal evolution of the 5 degrees of freedom
 - Expression
 - Intensity
- TP-GMM method (Calinon, 2015) to encode these examples

Encoding the demonstrated trajectories

Encoding the demonstrated trajectories

Wheel of Emotion (Plutchik, 1980)

Agree 7

Shy 7

Agree – Shy TPGMM

Generating Interpolated Motion

Agree-Shy 7

Gonzalo Mier, Fernando Caballero, Keisuke Nakamura, Luis Merino, and Randy Gomez.

Generation of expressive motions for a tabletop robot interpolating from hand-made animations.

In Proceedings of the IEEE International Conference on Robot and Human Interactive Communication, RO-MAN, 2019.

Conclusions and Outlook

- Social robotics require to consider the presence of people as a central part of a robot architecture
- Multi-modal perception of people is key
- Pertinent, expressive and legible robot behavior are required
- Data-driven approaches: social behaviors easy to demonstrate than to formalise

Machine Learning para los Retos de la Robótica Social

Luis Merino Cabañas Service Robotics Laboratory

http://robotics.upo.es