Conversational LEAPs in HRI

Prof. Oliver Lemon, Interaction Lab & Alana AI

HERIOT WATT UNIVERSITY

amazon alexa

EDINBURGH CENTRE FOR ROBOTICS

EPSRC Pioneering research and skills

THE FRAMEWORK PROGRAMME FOR RESEARCH AND INNOVATION HORIZON 2020

A L A N A

unicef

SPRING

Types of LEAP from conversational feedback

- Dialogue policy/ Reinforcement Learning – what to say next?
  - Goal-oriented dialogue, MDPs and POMDPs: Rieser & Lemon 2011
  - **Shalyminov et al, SCAI EMNLP 2018**

- User model updates – what interests this user?
  - Alana system: Papaioannou et al. proc. Alexa Prize 2017; Conv AI @ NeurIPS 2017
  - Alana system: Curry et al, proc Alexa Prize 2018

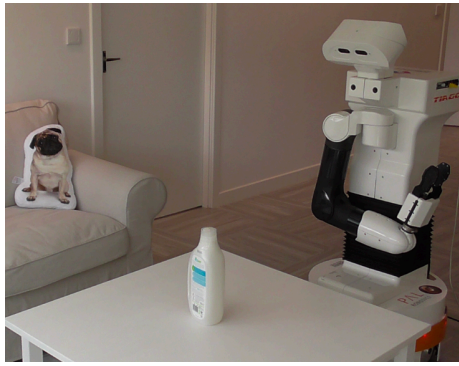

- Incremental Lifelong Learning – learning visual language
  - **Object Attributes - Yu et al ACL Robo-NLP 2017**
  - **Object Classes - Part & Lemon IROS 2019**

- Multimodal deep learning – latent representations via guessing games
  - **Zero-shot tasks: Suglia et al 2020**

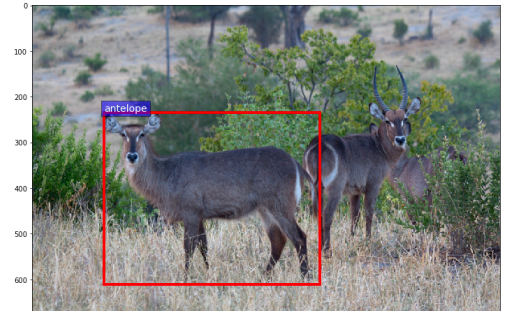

- Transfer Learning for dialogue tasks: Shalyminov et al 2020

Robotarium

- Opening 2022 @ Heriot-Watt University, Edinburgh
- https://www.edinburgh-robotics.org/robotarium
- 6 labs for HRI experiments
- “Living lab” smart home for assistive robots
- We are hiring ! & open for collaboration ☺

EDINBURGH CENTRE FOR ROBOTICS

A L A N A

- Our platform for Conversational AI research (academia and industry)
- Alexa Prize 2-time finalist and spin-out company from HWU
- Learning via feedback: task-based and open-domain social conversation
- Persistent user models – long-term memory
- >6M interactions, 2 years uptime in USA (available via Alexa “let’s chat”)
- Projects: UNICEF, RNIB, PAL (H2020 SPRING), Pepper (H2020 MuMMER)
- ROS interface + deployed on Android, Google Home, Pepper, Furhat, ARI
- Visual dialogue capability (see HRI 2021 demo video)
- www.alanaAI.com @helloalana

What is a “conversational robot” ?

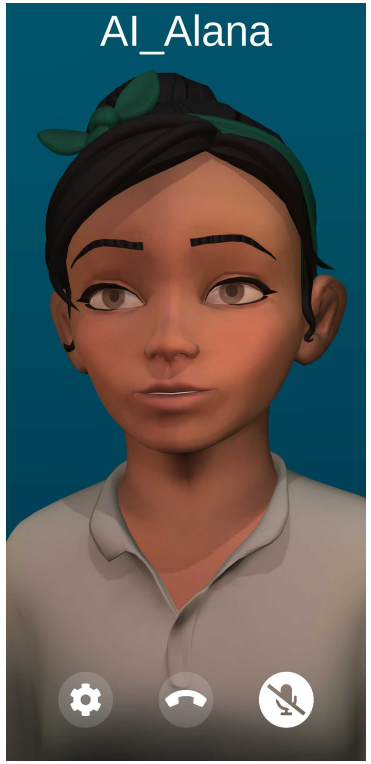

“Conversation is the interface of everyday life”

What do we do in conversations?

- collaborative tasks and negotiation
- visual conversation
- social chit-chat
- personalized entertainment
- wide-coverage Question-Answering
- user modelling and adaptation
- teaching and learning
- …….. “one app to rule them all”

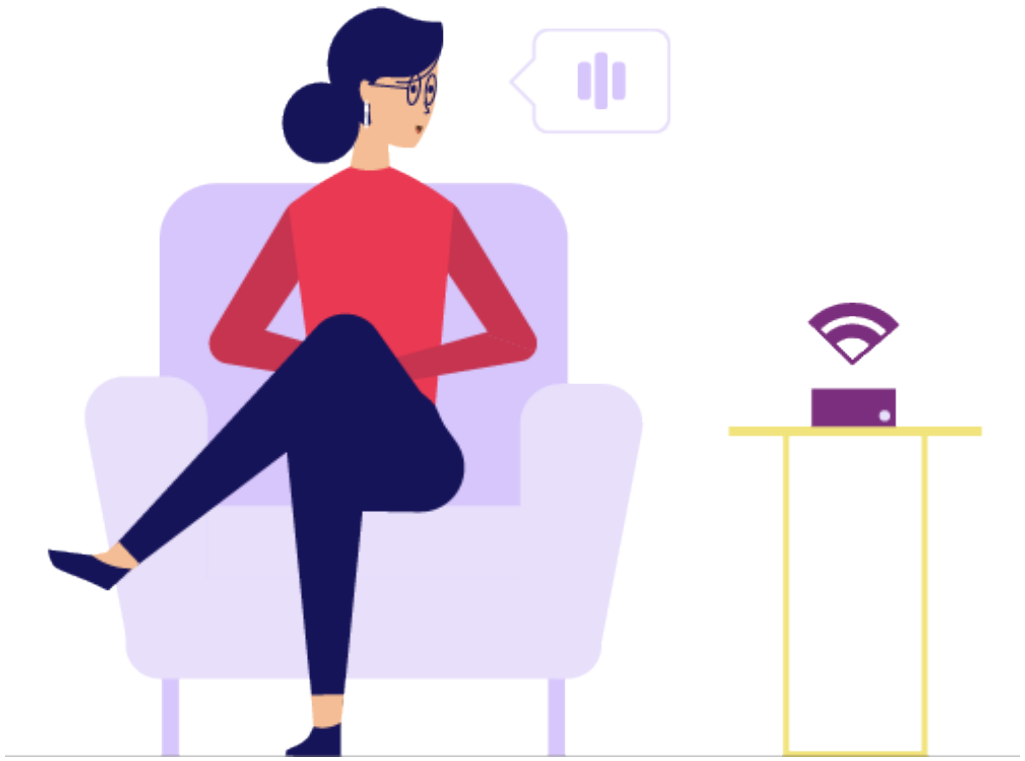

Example: personalization and context (2017)

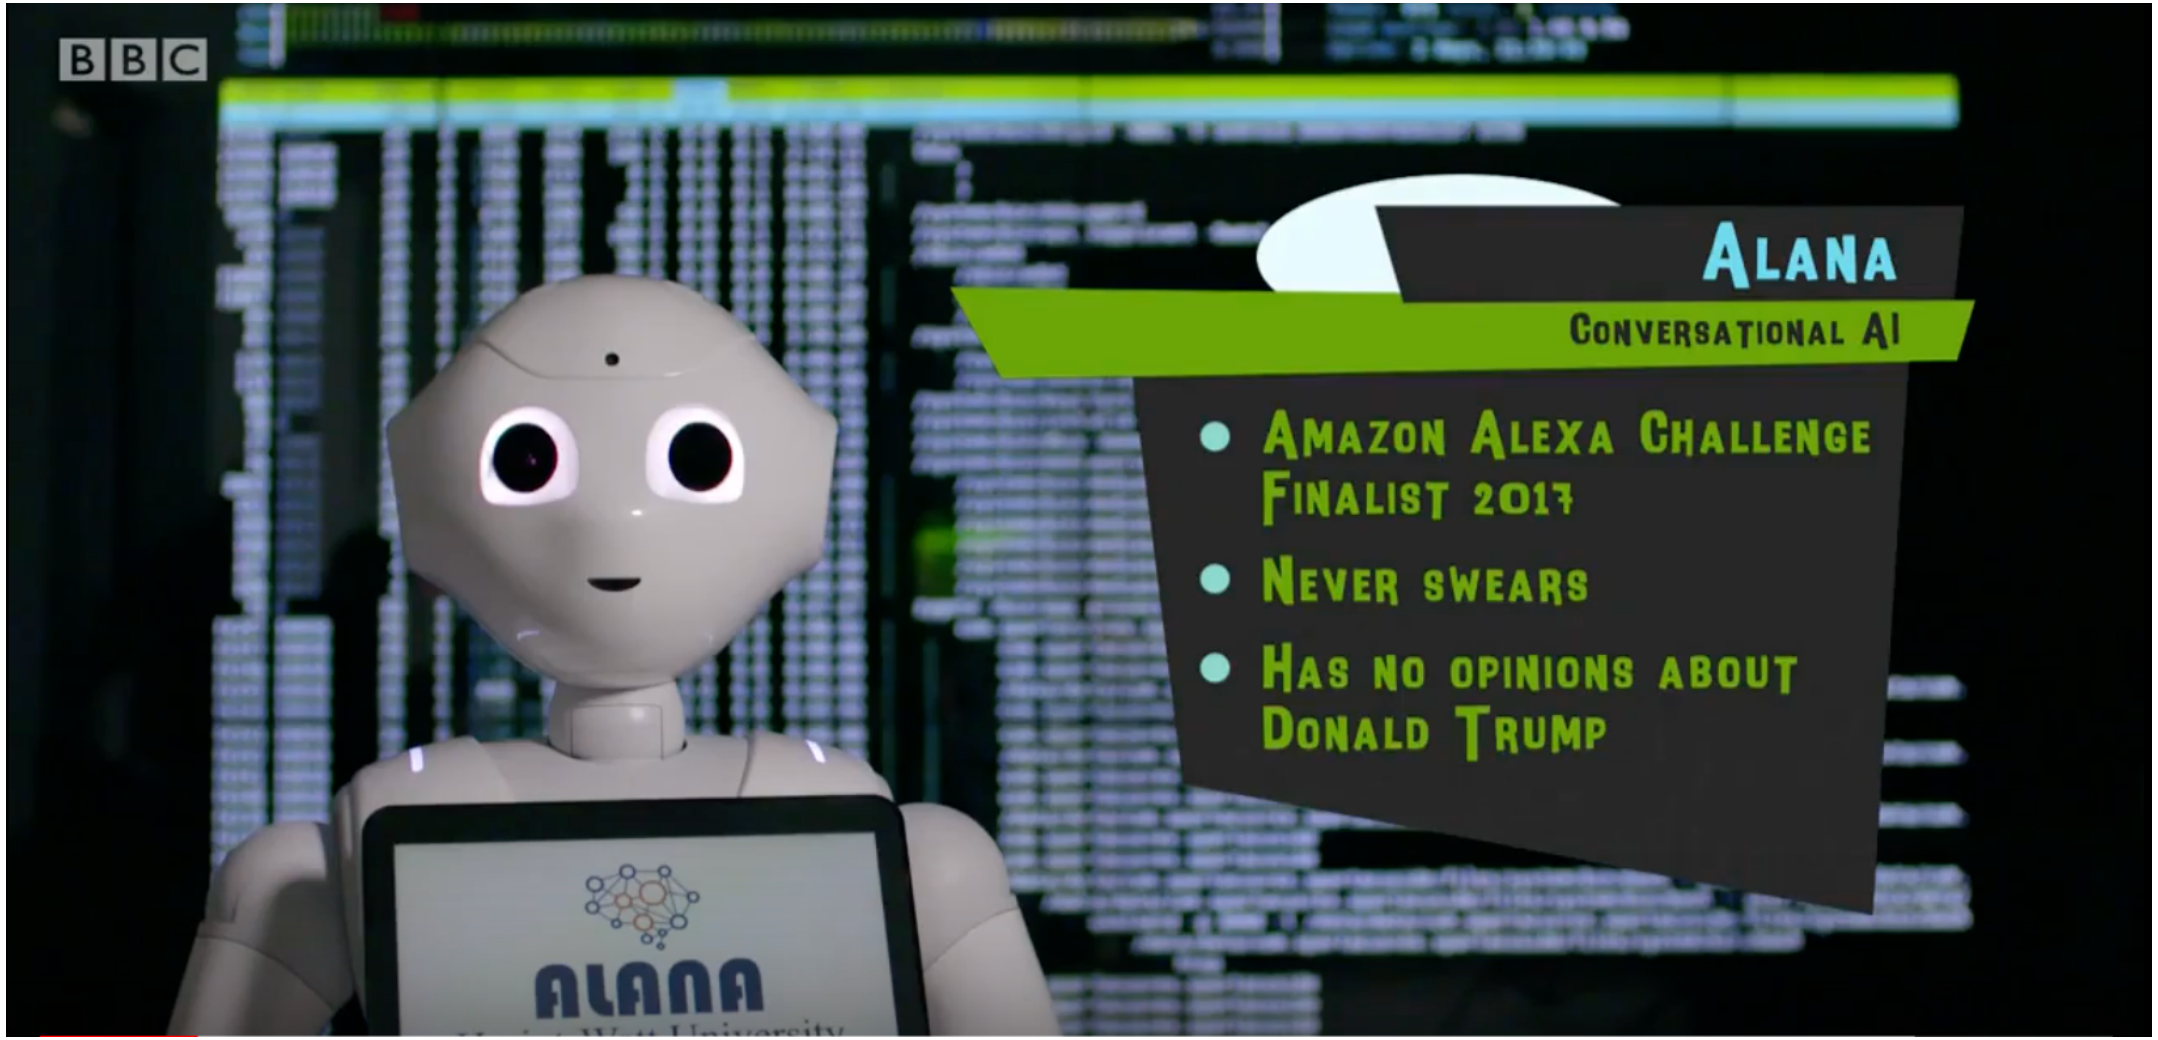

BBC: The Joy of AI : http://www.bbc.co.uk/programmes/p06jt7j4

State-of-the art in Conversational HRI (2021)?

“A bit deaf, stupid, forgetful, rude, and a know-it-all”

- A bit deaf - speech recognition errors in-the-wild
- A bit stupid
  - NLU problems; lacking commonsense; not very aware of surroundings
- A bit rude – not very attentive to users (MM signals)
- But - also a “know-it-all” / idiot savant
- Can’t usually handle > 1 person, not socially intelligent
- Not personalized
  - Not learning about user preferences, or diverse use of language
- → Not good at collaborating with humans ….. Why?

The problems of “context”

- “Context often means that you really had to be there” (Frankie Boyle)
- We are lacking rich **multimodal and social** context models
- Which can be rapidly **updated** by observation and **interaction**

SPRING

What about Transformers / GPT-3?

- Limited notion of context: trained on text, not multimodal, unreliable
- Not (yet) connected to user models, plans, tasks, goals, social scenes …
- E.g. Collaborative planning requires reasoning over abstractions
- We need to use deep learning as part of a larger cognitive architecture
- → Hybrid approach: controllable, inspectable, explainable, trustworthy

[Image from **Hady Elsahar, Naver Labs**]

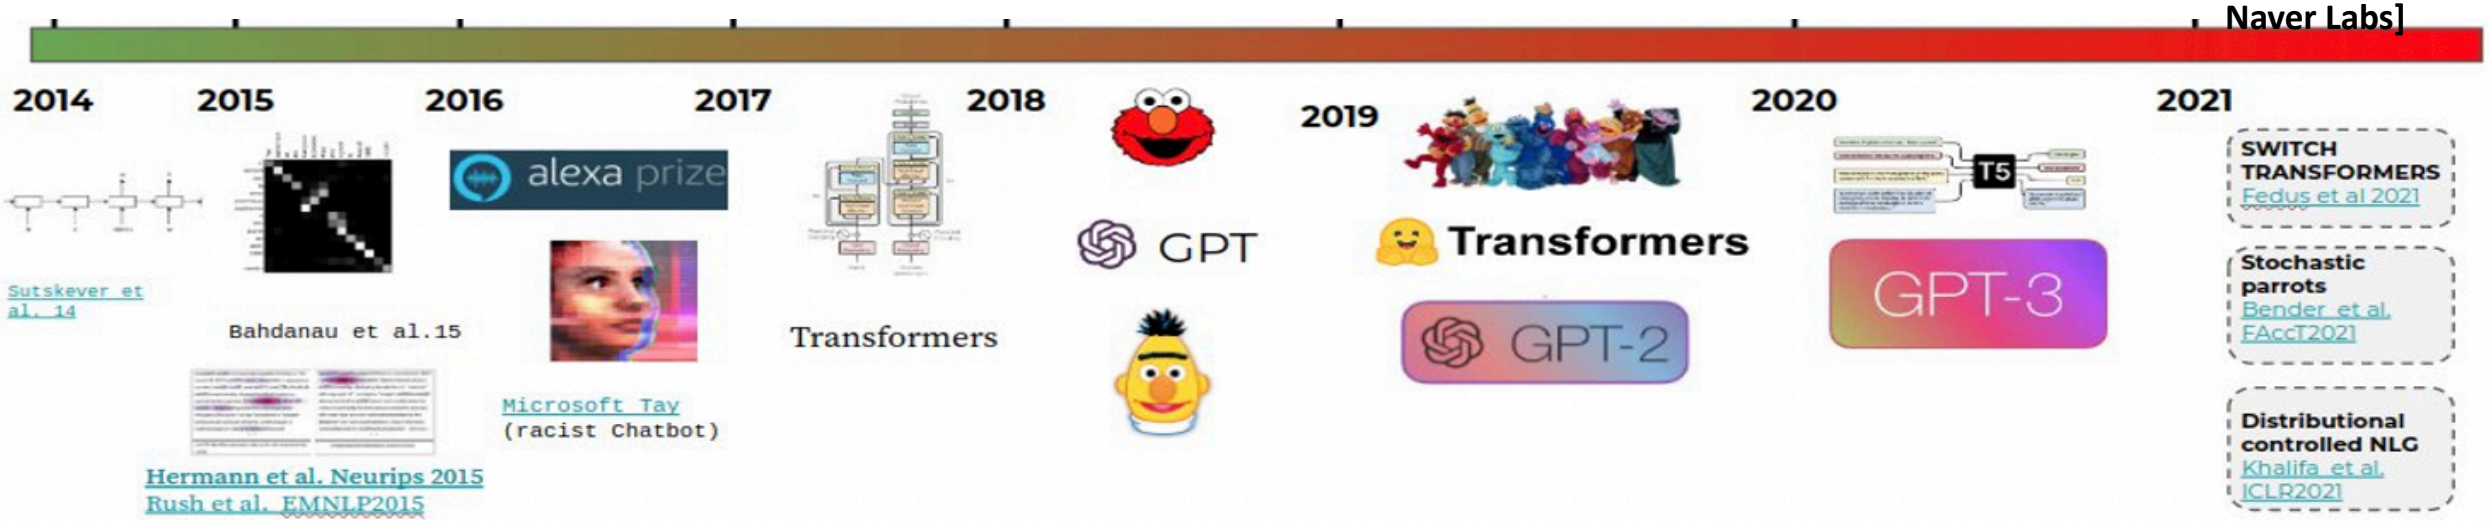

Modular Conversational HRI architecture

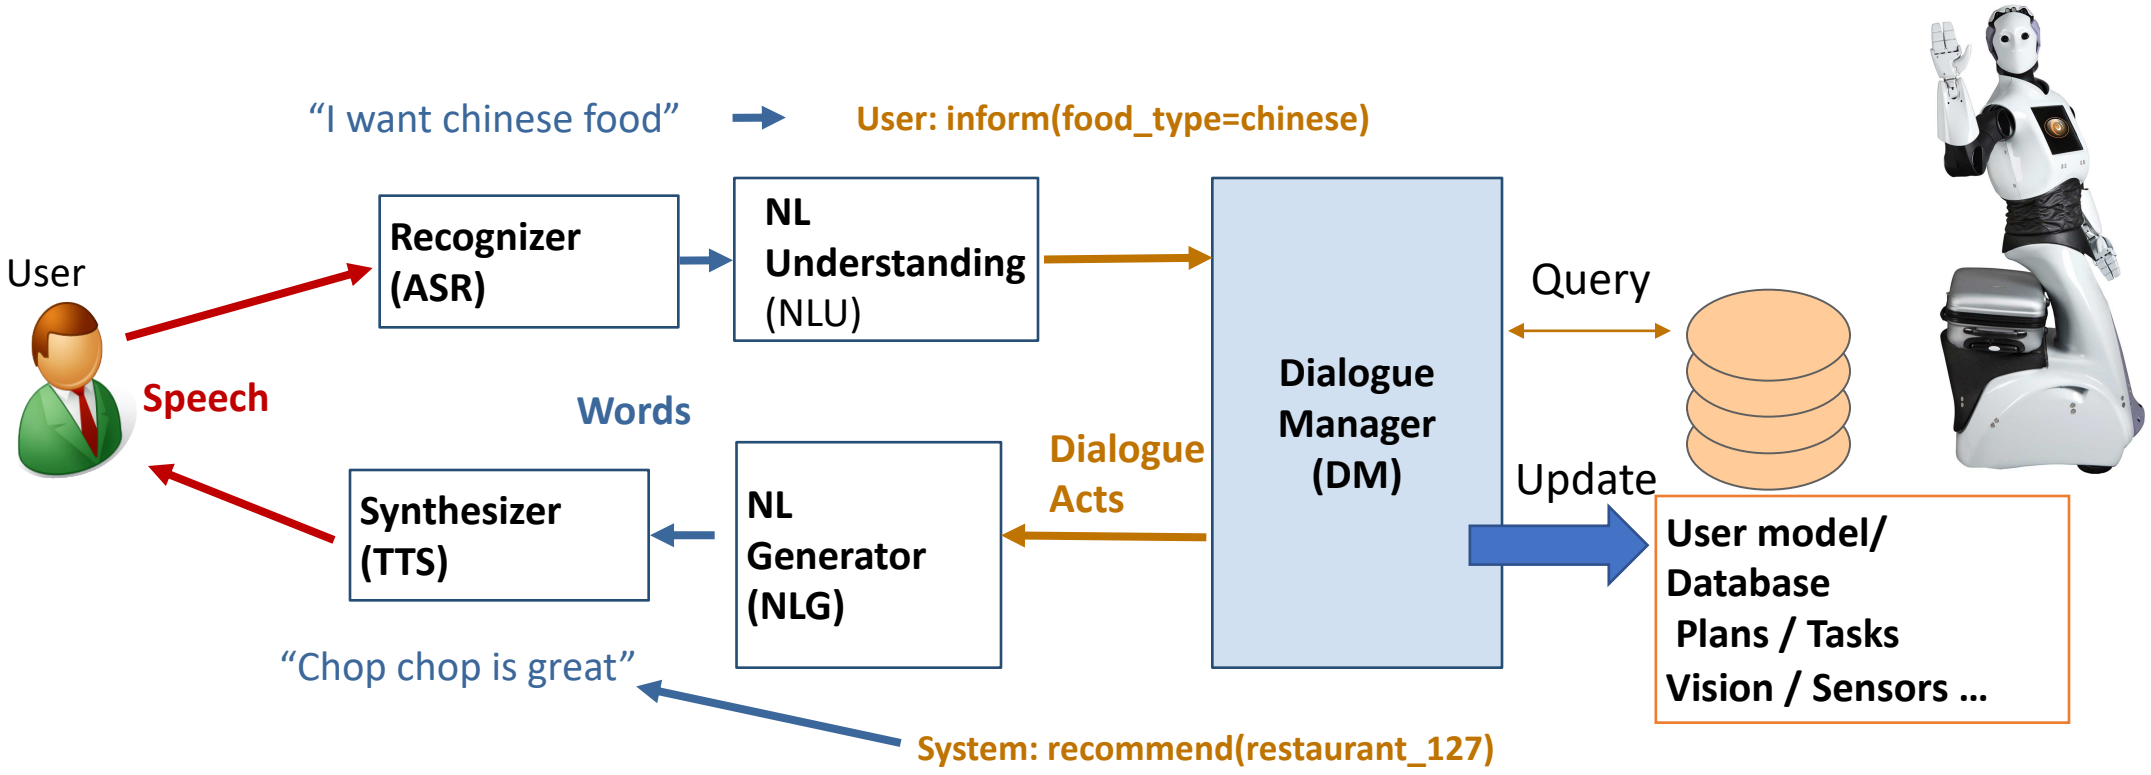

Conversation can update user models, tasks/plans, and visual language understanding

→ ability for NLU + DM to **re-train/update** the user model, visual classifiers etc

- processing of **feedback** and **repairs : 1-shot and few-shot** learning;

Conversational signals and (machine) learning

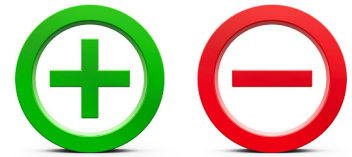

- Conversational feedback signals can be both **explicit** and **implicit**
  - Challenge is to understand them in context

- Explicit **repair == content + reward / supervision signal**
  - “No I meant fourteen not forty” ; “fourTEEN please”
  - “Actually I’m more interested in Volkswagen vans”
  - “No no that’s wrong”; “what?”; “huh?”

- Implicit feedback: pauses, gaze, facial expression, disfluencies (`uhm’, `err’)
  - incremental, continuous, online feedback

Types of Learning from conversational feedback

- Dialogue policy/ Reinforcement Learning – what to say next?
  - Goal-oriented dialogue, MDPs and POMDPs: Rieser & Lemon 2011
  - **Shalyminov et al, SCAI EMNLP 2018**

- User model updates – what interests this user?
  - Alana system: Papaioannou et al. proc. Alexa Prize 2017; Conv AI @ NeurIPS 2017
  - Alana system: Curry et al, proc Alexa Prize 2018

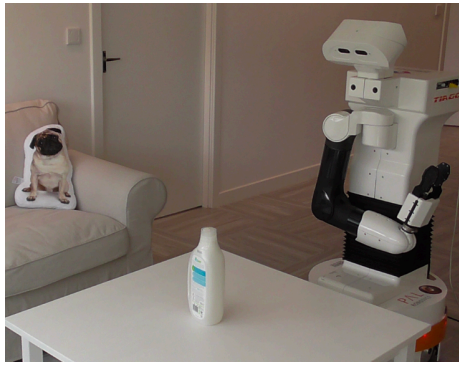

- Incremental Lifelong Learning – learning visual language
  - **Object Attributes - Yu et al ACL Robo-NLP 2017**
  - **Object Classes - Part & Lemon IROS 2019**

- Multimodal deep learning – latent representations via guessing games
  - **Zero-shot tasks: Suglia et al 2020**

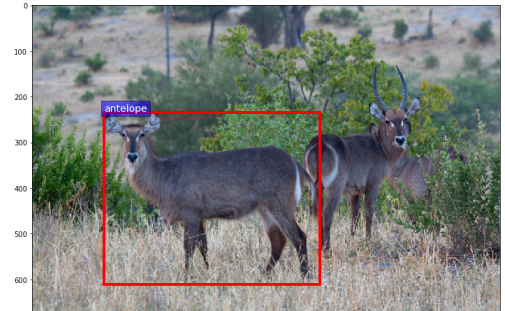

- Transfer Learning for dialogue tasks: Shalyminov et al 2020

User feedback in Alana

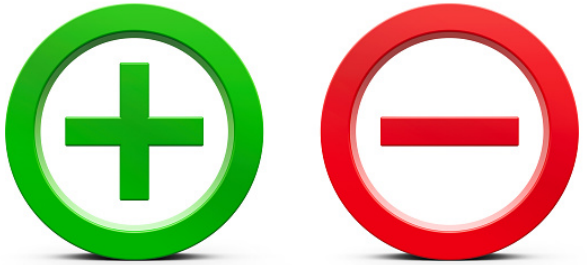

- What is the best thing to say next?

- Policy learning: use feedback signals to train a **strategy** for the Dialogue Manager

- POSITIVE: “thanks / that’s interesting / you’re funny ….”
- NEGATIVE: “no no no that’s wrong / oh stop / that’s stupid / I meant …”

- POSITIVE: long conversations are better (10% were > 10 mins)
- NEGATIVE: short conversations

Alana's Neural Response Ranker

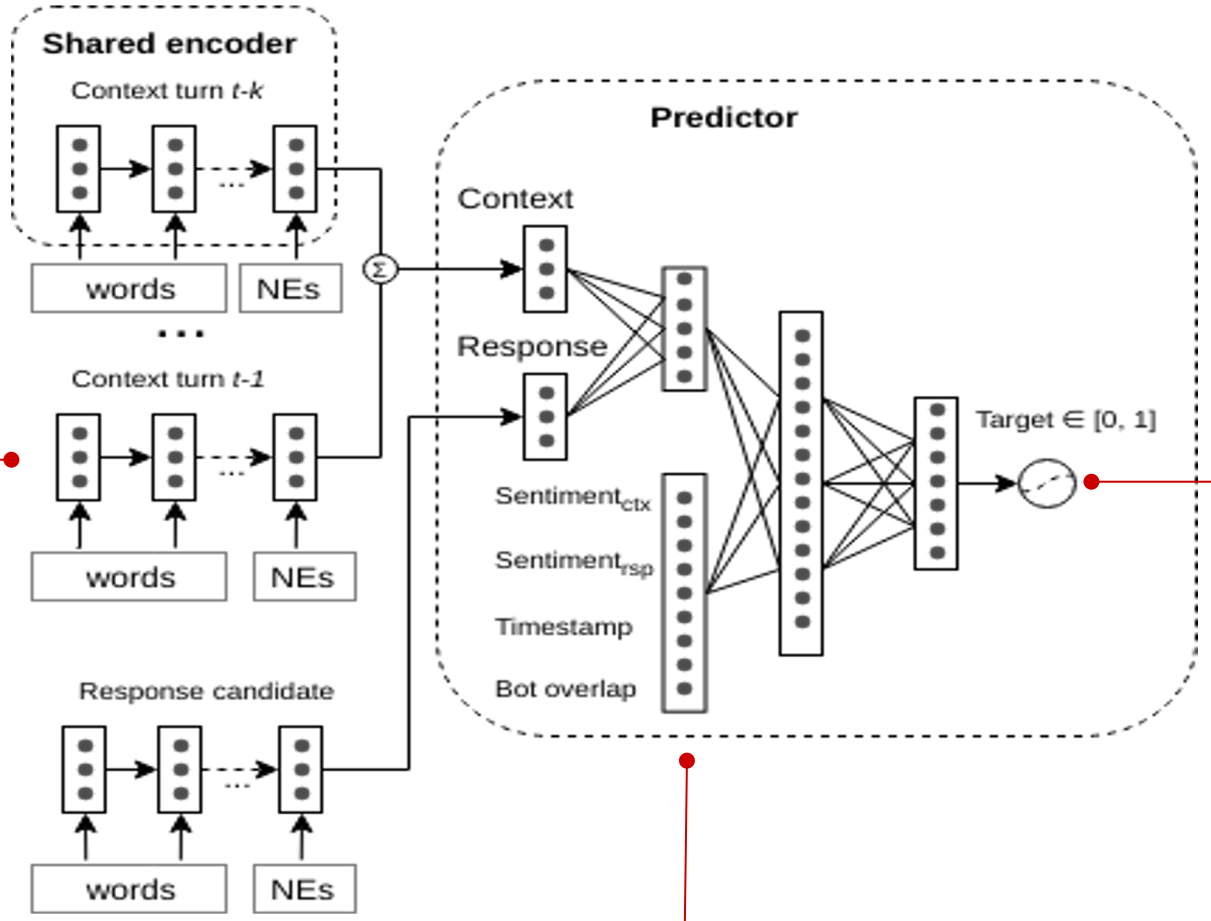

GRU with word, bot, and NE features: encodes context utterances (3 previous user + system turns) and response candidates

Target score is a point-wise prediction for a <*context, response_candidate*> pair

Predictor MLP augments the representation with additional dialogue features and outputs the target value

Igor Shalyminov, Ondřej Dušek and Oliver Lemon. Neural Response Ranking for Social Conversation: A Data-Efficient Approach. **SCAI @ EMNLP**, 2018

HERIOT WATT UNIVERSITY

Multimodal conversational LEAPs

- Lifelong learning: ability to learn new words for new objects and attributes
- Personalisation: in-the-wild: users will have personal/ idiosyncratic language (e.g. object names, colours, relations…) → diversity
- → need to adapt to each person's language use, dialects, etc

Labels, repairs, and retraining

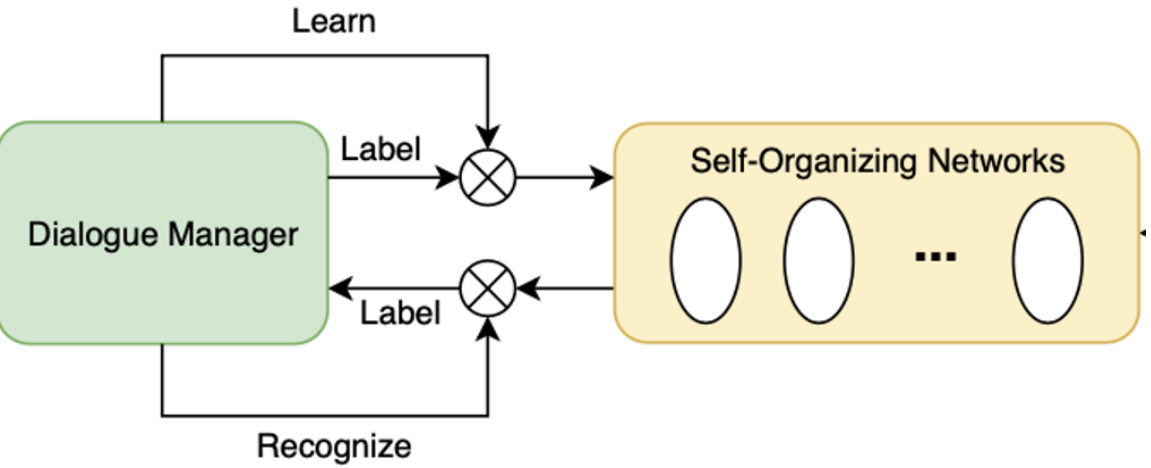

NLU + Dialogue Policy Learning

-- sys: is this a black square?
-- usr: no, the shape is correct but color is wrong.
-- sys: okay. so what's the color?
-- usr: it is red ... sorry ... purple.
-- sys: okay, so a purple square.
-- usr: yes, well done.

VOILA: Incrementally Learning visual attributes

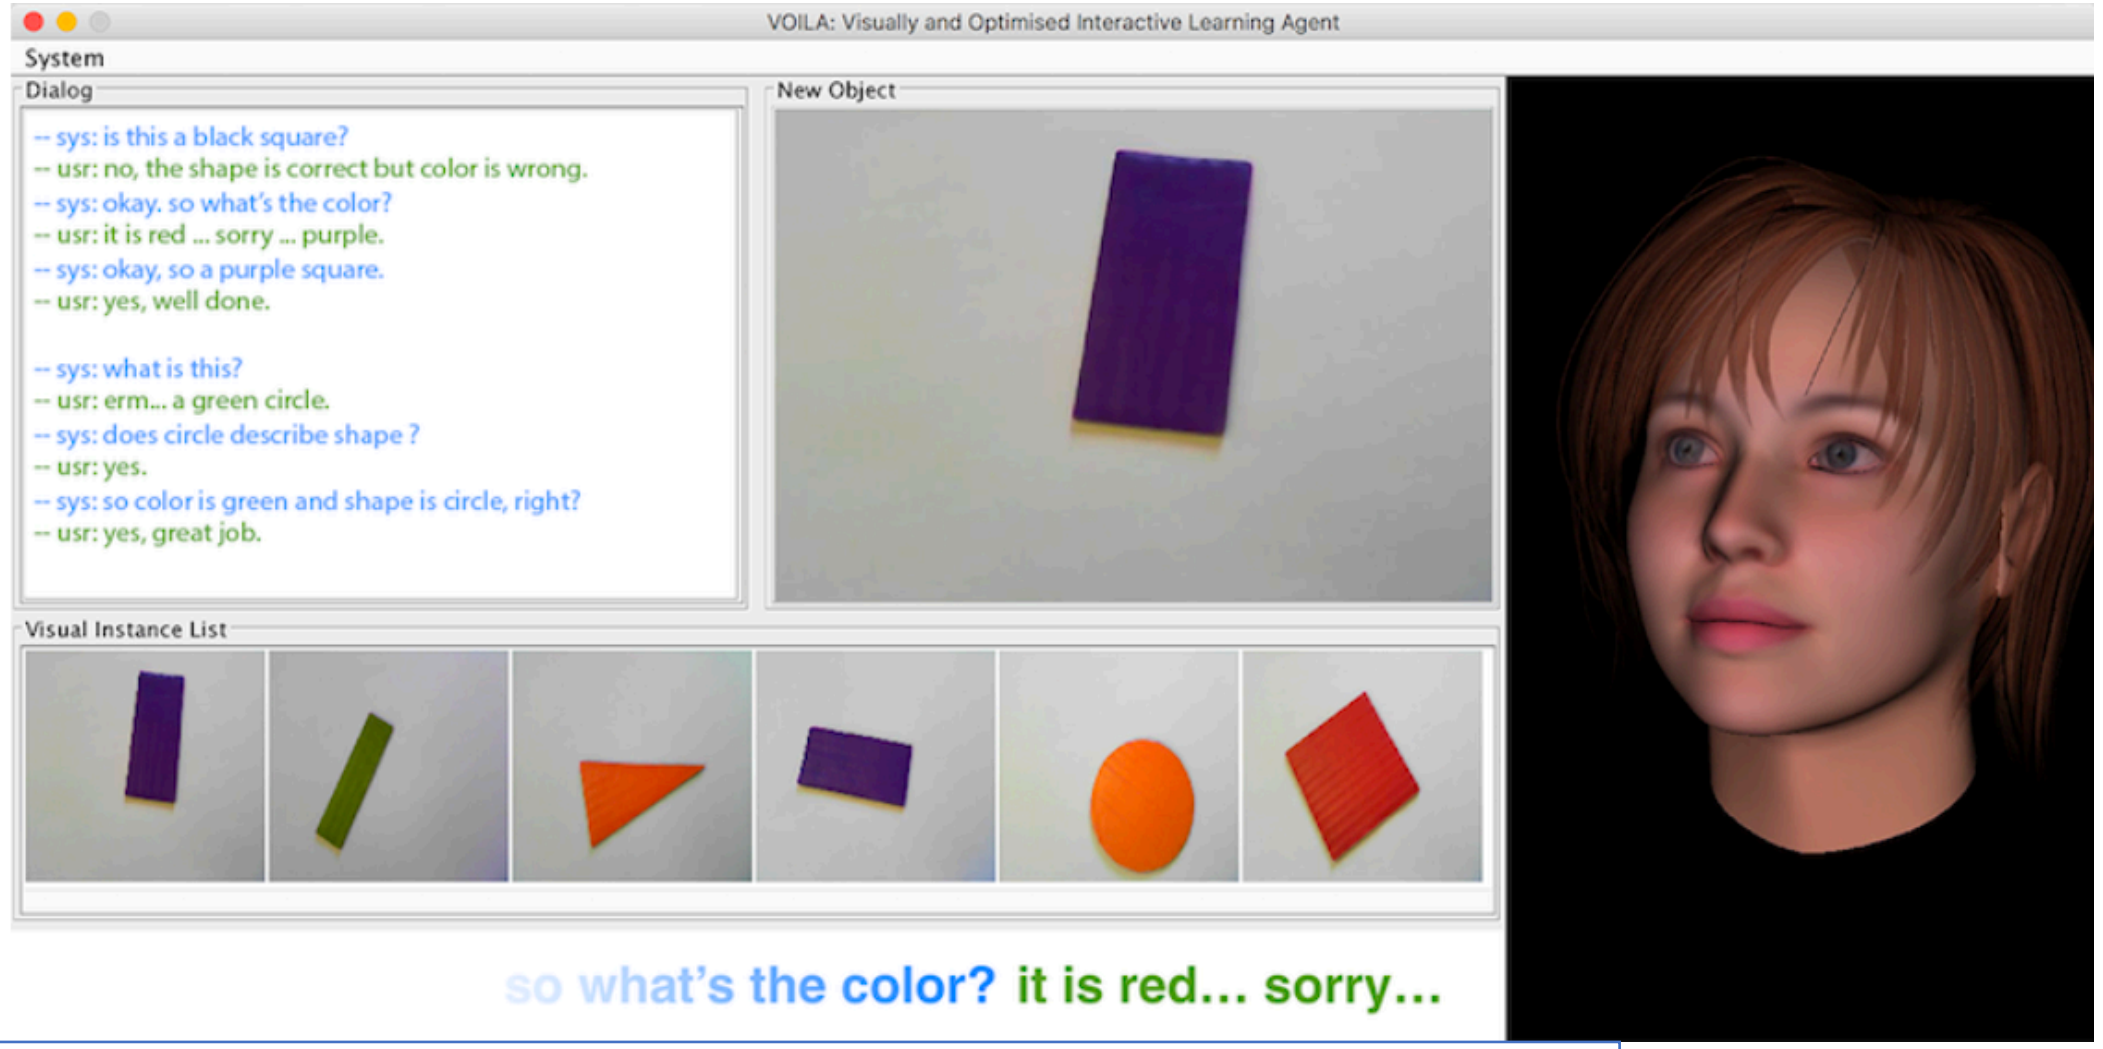

Yanchao Yu, Arash Eshghi, and Oliver Lemon, "Learning how to learn: an adaptive dialogue agent for incrementally learning visually grounded word meanings", **Robo-NLP @ ACL** 2017

Conversational Incremental life-long learning

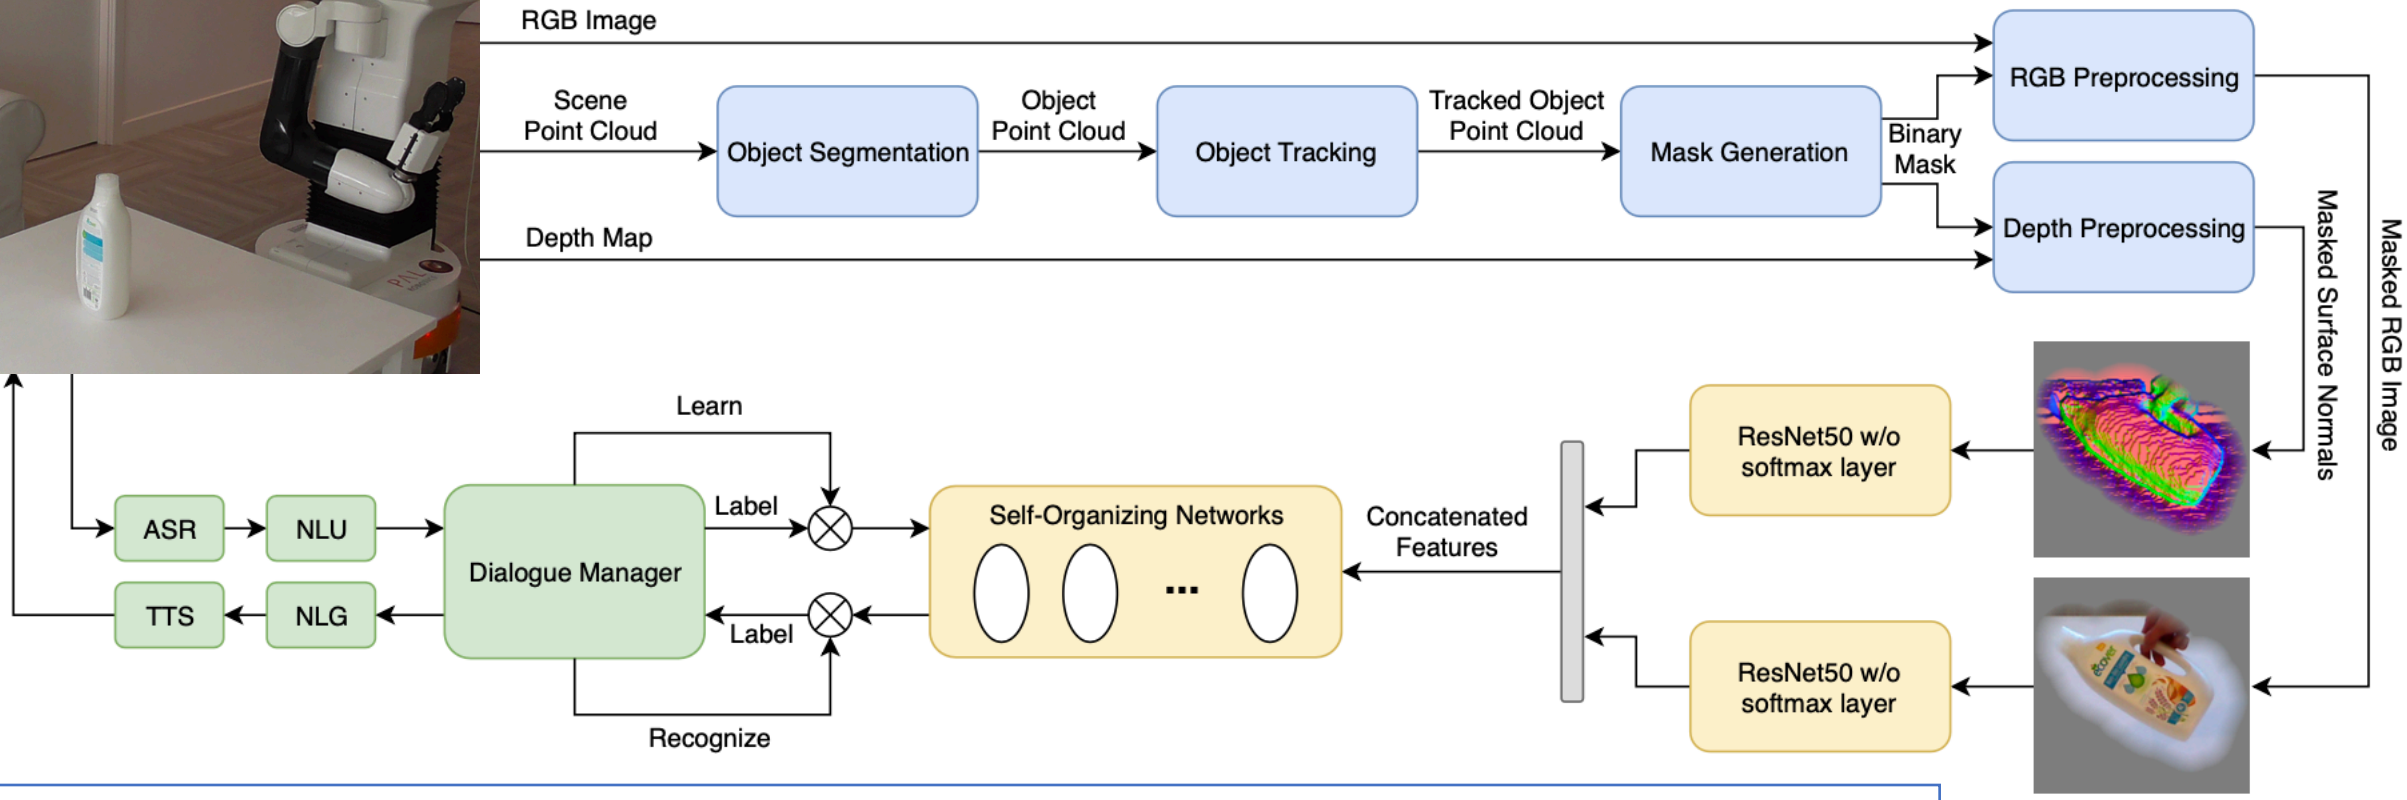

Jose Part and Oliver Lemon, "Towards a Robot Architecture for Situated Lifelong Object Learning", Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (**IROS**), 2019

Grounded language learning with visual games

- “Imagination” module based on Regularized Auto-encoders.
- Learns context-aware latent embeddings
- State of the art (+8%) accuracy in zero-shot scenarios

Alessandro Suglia et al "Imagining Grounded Conceptual Representations from Perceptual Information in Situated Guessing Games", **COLING 2020**

GDSE-CL+**imagination**		GDSE-CL	
is it a person?	yes	is it a person?	no
is it the kid on the bike?	yes	is it a bike?	no
is it the whole person?	yes	is it the green bag?	no
is it wearing an helmet?	yes	is the the bike?	yes
		is it the bike in the back?	no
		is the the bike wheel to the right?	no
		is it the bike to the left of the bike?	yes
		is it the whole bike?	yes
status: success		status: failed	

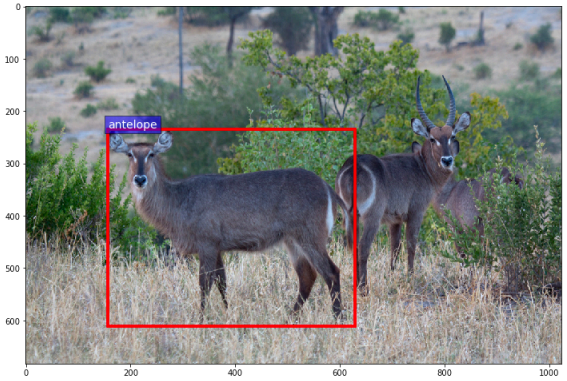

GDSE-CL+**imagination**		GDSE-CL	
is it an animal?	yes	Is it a giraffe?	no
is it the one in the middle?	yes	Is it a car?	no
is it the whole animal?	yes	Is it the wood plank in the background	no
is it the one whose head is visible?	yes	is the animal the animal sitting on?	yes
is it the whole animal?	yes	is the object the animal on the brown animal?	no
is it the entire animal?	yes	is it the black jacket with the brown handle?	no
is the head facing to the left side of the picture?	yes	is the animal's head visible?	no
is the head facing to the left?	yes	is the object brown?	no
		is the handle black?	no
status: success		status: failed	

Types of Learning from conversational feedback

- Dialogue policy/ Reinforcement Learning – what to say next?
  - Goal-oriented dialogue, MDPs and POMDPs: Rieser & Lemon 2011
  - **Shalyminov et al, SCAI EMNLP 2018**

- User model updates – what interests this user?
  - Alana system: Papaioannou et al. proc. Alexa Prize 2017; Conv AI @ NeurIPS 2017
  - Alana system: Curry et al, proc Alexa Prize 2018

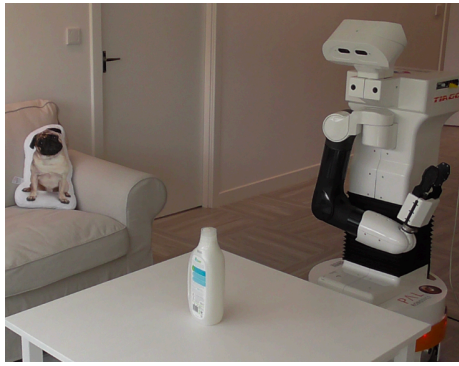

- Incremental Lifelong Learning – learning visual language
  - **Object Attributes - Yu et al ACL Robo-NLP 2017**
  - **Object Classes - Part & Lemon IROS 2019**

- Multimodal deep learning – latent representations via guessing games
  - **Zero-shot tasks: Suglia et al 2020**

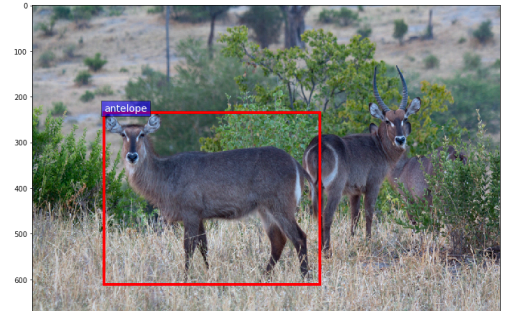

- Transfer Learning for dialogue tasks: Shalyminov et al 2020

SPRING project: Alana + visual + social

SPRING

http://spring-h2020.eu/

ALANA

http://www.alanaAI.com/

Jose Part et al "Towards Visual Dialogue for Human-Robot Interaction", **HRI 2021**

https://youtu.be/eY_BNxwrlPg

RNIB

See differently

Summary / Challenges

- Life-long learning and personalization can benefit from conversational interaction
- Conversational signals can be used to adapt ML models
- But: we need richer MM signals & context models, NLU, DM

Next steps for collaboration?

- Multimodal signals → interactive continuous feedback
- Multi-user HRI → learning social conversation
- Simulations of social / visual / spatial scenes
- Adapting to individuals & diverse language use
- Real evaluations ``in the wild”

EDINBURGH CENTRE FOR ROBOTICS

HORIZON EUROPE

A L A N A

Robotarium

- Let’s collaborate!
- Opening 2022 @ Heriot-Watt University, Edinburgh
- https://www.edinburgh-robotics.org/robotarium
- 6 labs for HRI experiments
- “Living lab” smart home for assistive robots

EDINBURGH CENTRE FOR ROBOTICS

Papers

- Jose Part et al. "Towards Visual Dialogue for Human-Robot Interaction", Human-Robot Interaction (**HRI 2021**), video demonstration [VIDEO]
- Alessandro Suglia et al. , "Imagining Grounded Conceptual Representations from Perceptual Information in Situated Guessing Games", **COLING 2020**
- Daniel Hernandez Garcia, et al. "Explainable Representations of the Social State: A Model for Social Human-Robot Interactions" **AAAI Fall Symposium on AI for HRI**, 2020
- Igor Shalyminov, Sungjin Lee, Arash Eshghi and Oliver Lemon, "Data-Efficient Goal-Oriented Conversation with Dialogue Knowledge Transfer Networks", **EMNLP-IJCNLP** 2019 [arXiv]
- Jose Part and Oliver Lemon, "Towards a Robot Architecture for Situated Lifelong Object Learning", Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (**IROS)**, 2019
- Amanda Curry, et al, "Alana v2: Entertaining and Informative Open-domain Social Dialogue using Ontologies and Entity Linking", **Alexa Prize Proceedings, Amazon RE-INVENT**, Las Vegas, 2018 [pdf]
- Yanchao Yu, Arash Eshghi, and Oliver Lemon, "Learning how to learn: an adaptive dialogue agent for incrementally learning visually grounded word meanings", **Robo-NLP workshop at ACL** 2017 *** BEST PAPER AWARD *** [pdf]
- www.macs.hw.ac.uk/InteractionLab