B-CRATOS

Wireless Brain-Connect inteRfAce TO machineS

Machine Learning for BCIs

Paolo Viviani, PhD
Senior Researcher, Advanced Computing - LINKS Foundation
paolo.viviani@linksfoundation.com

ICTP - ADVANCED SCHOOL ON APPLIED MACHINE LEARNING

Speaker introduction

- Theoretical physics Msc – University of Torino
- Computer Science PhD – University of Torino
  - Distributed training of DL models before it was cool
- A couple of years in CAE industry doing numerical/ML stuff
- Now I'm back to research **at Advanced Computing, Photonics and Electromagnetics research domain of LINKS Foundation**, in Torino
  - HPC (and some ML)
  - Quantum Computing

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044
CRATOS
FONDAZIONE LINKS PASSION FOR INNOVATION
linksfoundation.com
COPYRIGHT ©2021 LINKS

Background – B-Cratos project

Applied Data Science: focus on application → context of research in a larger project

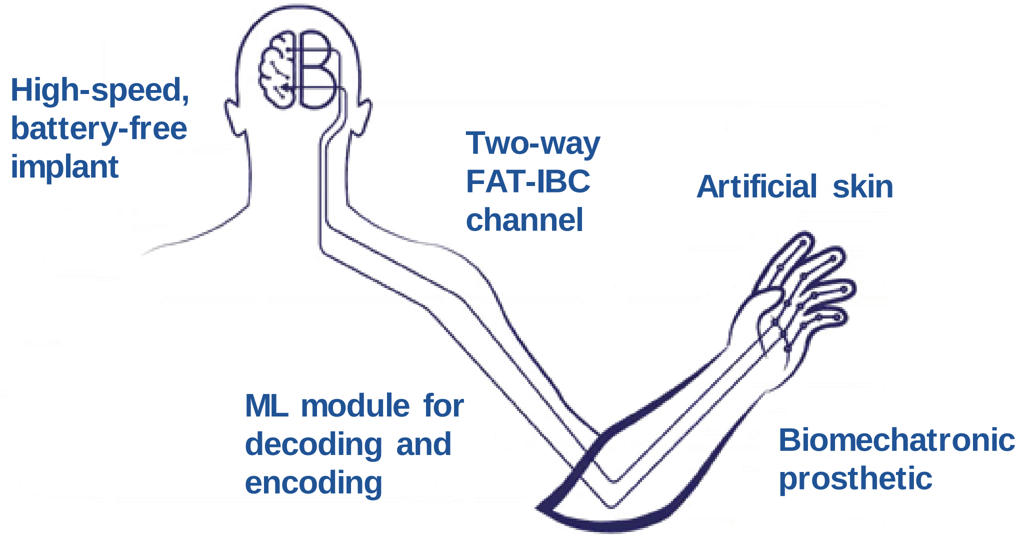

March 2021 – February 2025

FET-OPEN 4.7M EU funding

Partners from 5 countries

Coordinated by University of Uppsala

The goal is to develop a closed-loop BCI and to validate it with NHPs (Non-Human Primates), with the following technical objectives:

1. Proof-of-concept, high-speed, wireless brain implant capable of two-way communication without battery
2. General-purpose, high-speed intra-body communications technology (Fat-IBC)
3. **HPC based ML models deployed on embedded board for low-power inference and control**
4. Improvements of biomechatronic hand prothesis
5. High-resolution sensorized skin

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

What is a BCI?

The New York Times

Despite Setback, Neuralink's First Brain-Implant Patient Stays Upbeat

2 days ago • Christina Jewett

Bloomberg

Podcast: What Neuralink's First Patient Is Thinking

3 days ago • Reyhan Harmanci

WIRED

Neuralink's First User Is 'Constantly Multitasking' With His Brain Implant

2 days ago • Emily Mullin

Funded by the European Union

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

What is a BCI?

- Brain-Computer Interface
  - A device that enables communication and control without movement (BNCI Roadmap, Horizon 2020)
  - i.e., **readout of brain activity**
    → action/control signals

- “Closing the loop”: providing feedback to the brain
  - E.g., tactile stimulation

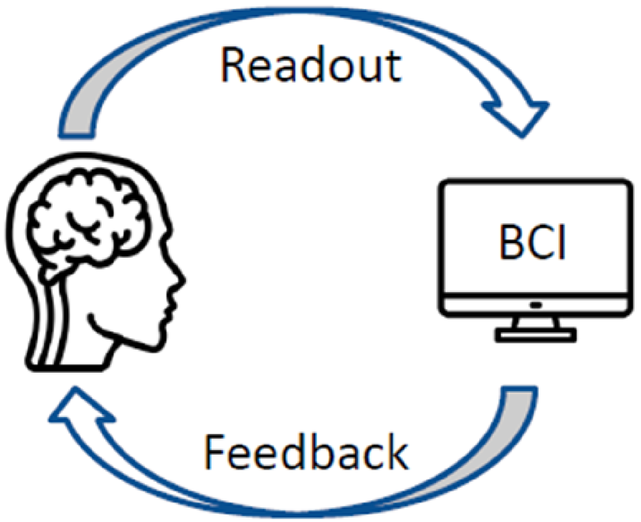

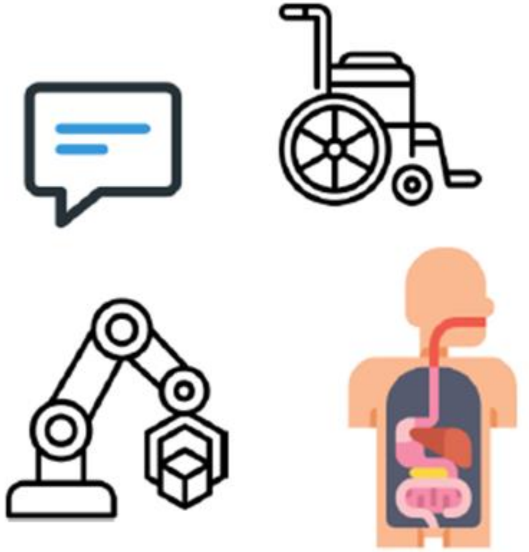

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

Project Overview

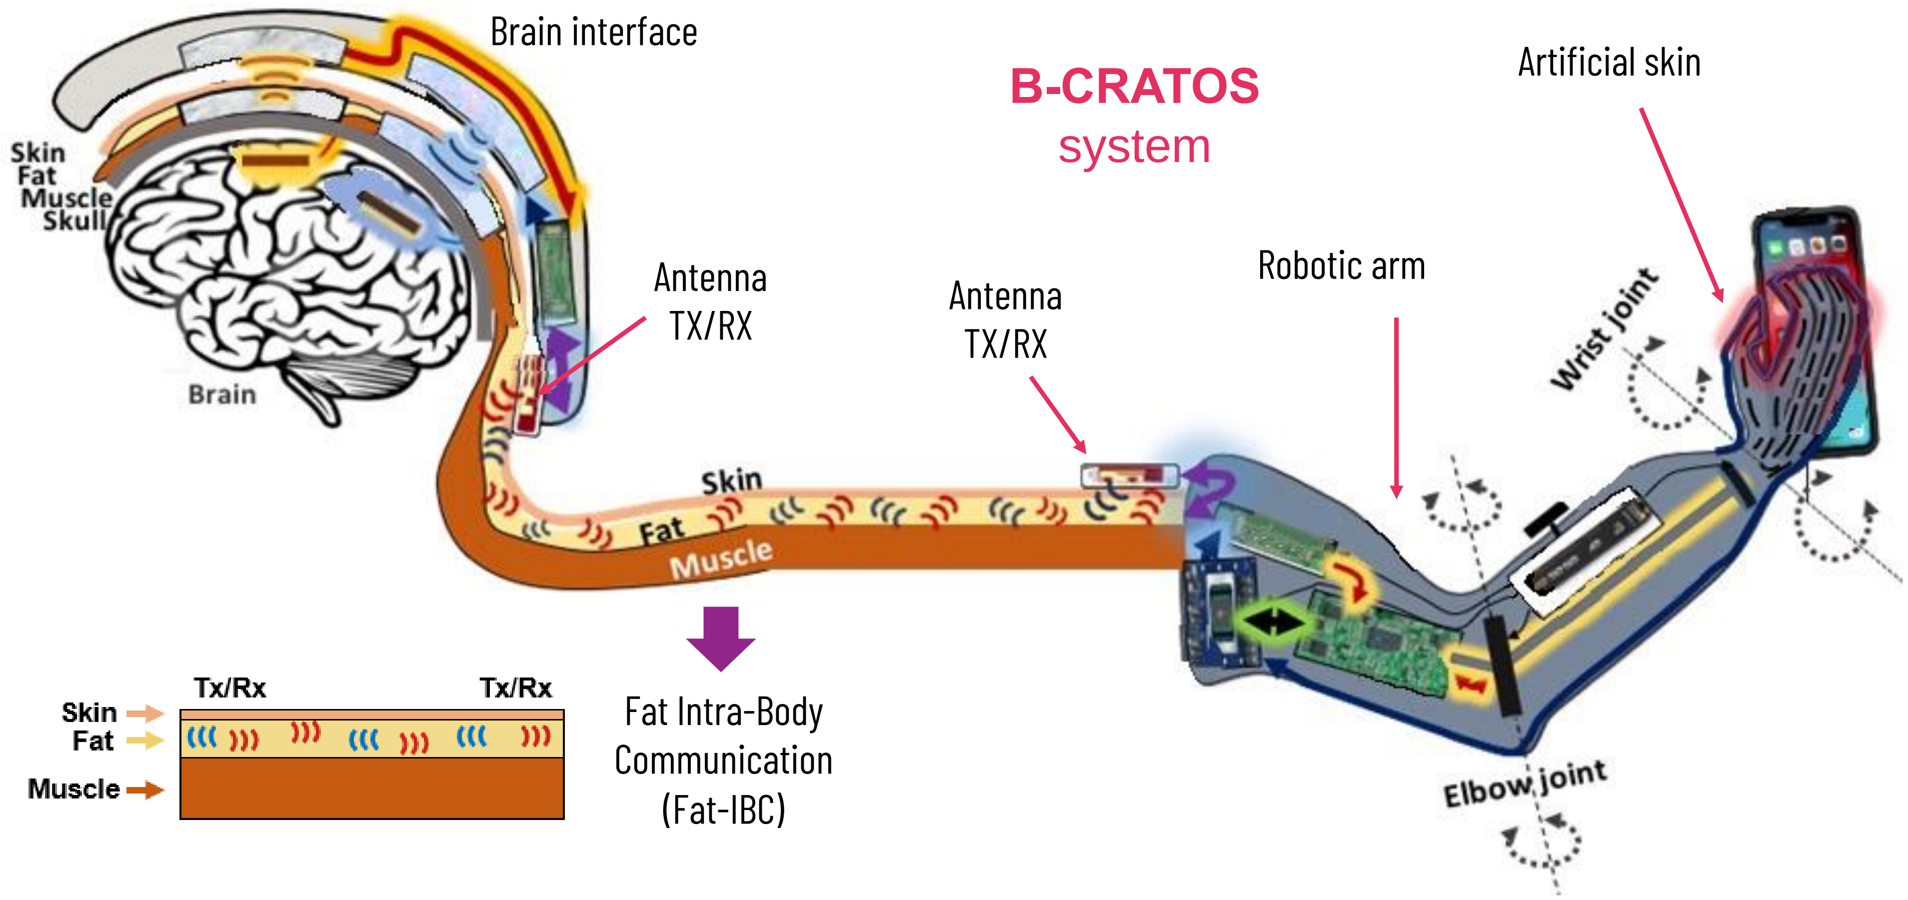

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

COPYRIGHT ©2021 LINKS

Why are we discussing BCIs in an ML school?

- To translate brain signals into commands for a computer/prosthesis, we need some kind of model
- The typical approach in **neuroscience** is to **use very simple models** (linear classifiers, Kalman filters, etc.) because neuroscientists want interpretable models
  - to understand how the brain works
  - because they work for their purposes
- **ML can play a role → Neural decoding**

- **"The man was able to train his mind to use the device"**
  - https://www.technologyreview.com/2022/03/22/1047664/locked-in-patient-bci-communicate-in-sentences/
- The goal of BCIs is to improve the life of patients, more sophisticated models can improve the experience and go in the direction of reversing this paradigm

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

COPYRIGHT ©2021 LINKS

Neural Decoding for Neuroprosthetics

Overview

- **Goal: translate the signal from brain implants into commands for a prothesis**
- Main steps
  - Brain electrodes pick up analogue signals from large neuron populations
  - Post-processing is applied to convert the signal to "spike-trains": multi-channel time series of binary data
  - Decoding algorithm is applied to convert the signal to command for the prothesis

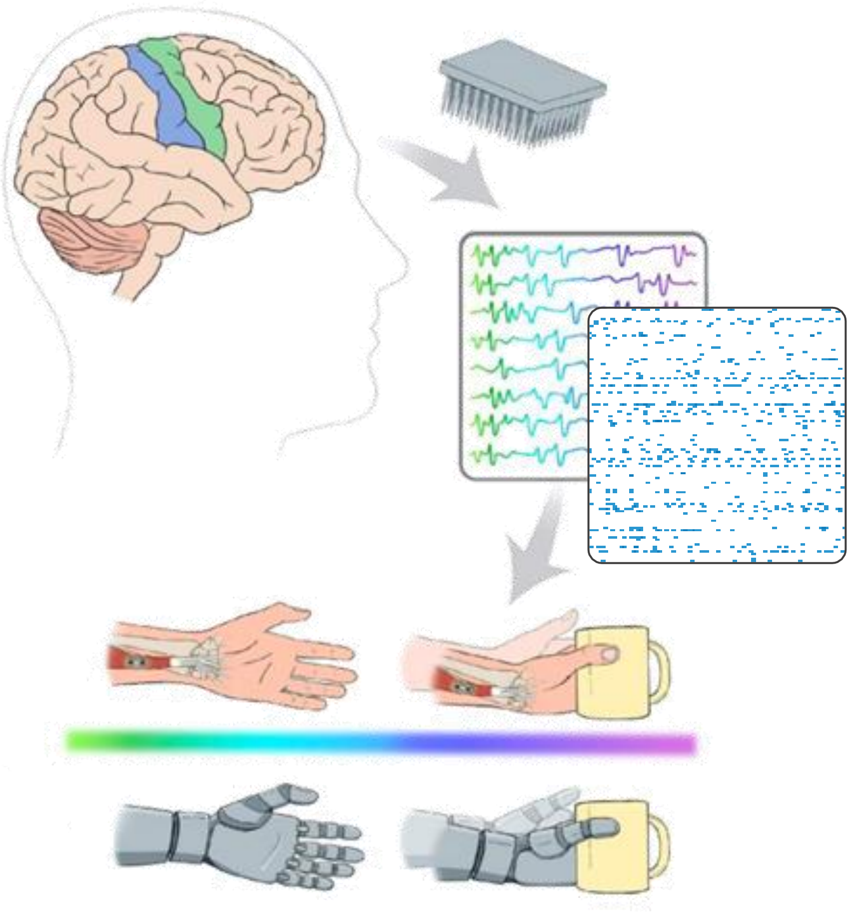

From C. Pandarinath and S. J. Bensmaia, "The science and engineering behind sensitized brain-controlled bionic hands," Physiological Reviews, p. physrev.00034.2020, Sep. 2021 doi: 10.1152/physrev.00034.2020. **Edited.**

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044
linksfoundation.com
COPYRIGHT ©2021 LINKS

Neural Decoding

State of the art

Neural signal

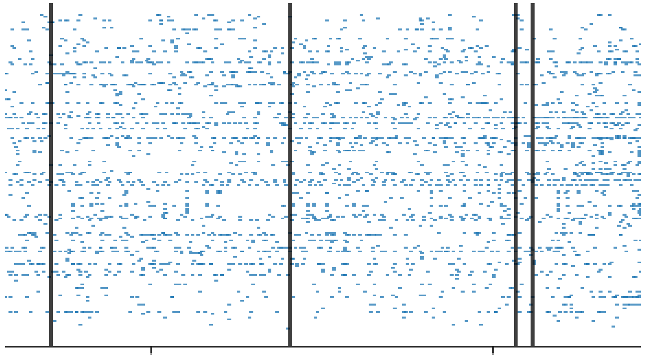

Decoder

Intended outcome

- Off-line decoding is well understood, **real-time decoding is challenging**
- State-of-the-art models are simple: **Patient “learns to drive” the model**
  - Also reflects neuroscientist’s need for explainability
- Neural signal changes over time and requires daily re-training of models
- B-Cratos aims to **improve the overall accuracy of models and to re-use information across sessions**
  - Lower effort for patient, easier adoption, more natural usage
  - **Modern machine learning techniques can support this goal**

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

Enough talking, let's get to the ML

No online data yet, but reference dataset was available

- Focused on a reference dataset provided by a partner (German Primate Center)
- Monkey trained for a grasping task
- Neural recording from multiple cortex regions
- Monkey grasping objects presented in sequence
- Objects grouped by shape and size, presented randomly
- 2 monkeys, 6 recording sessions (i.e., different days)
- **Learning Tasks**
  - Grasping phase detection
  - Object classification

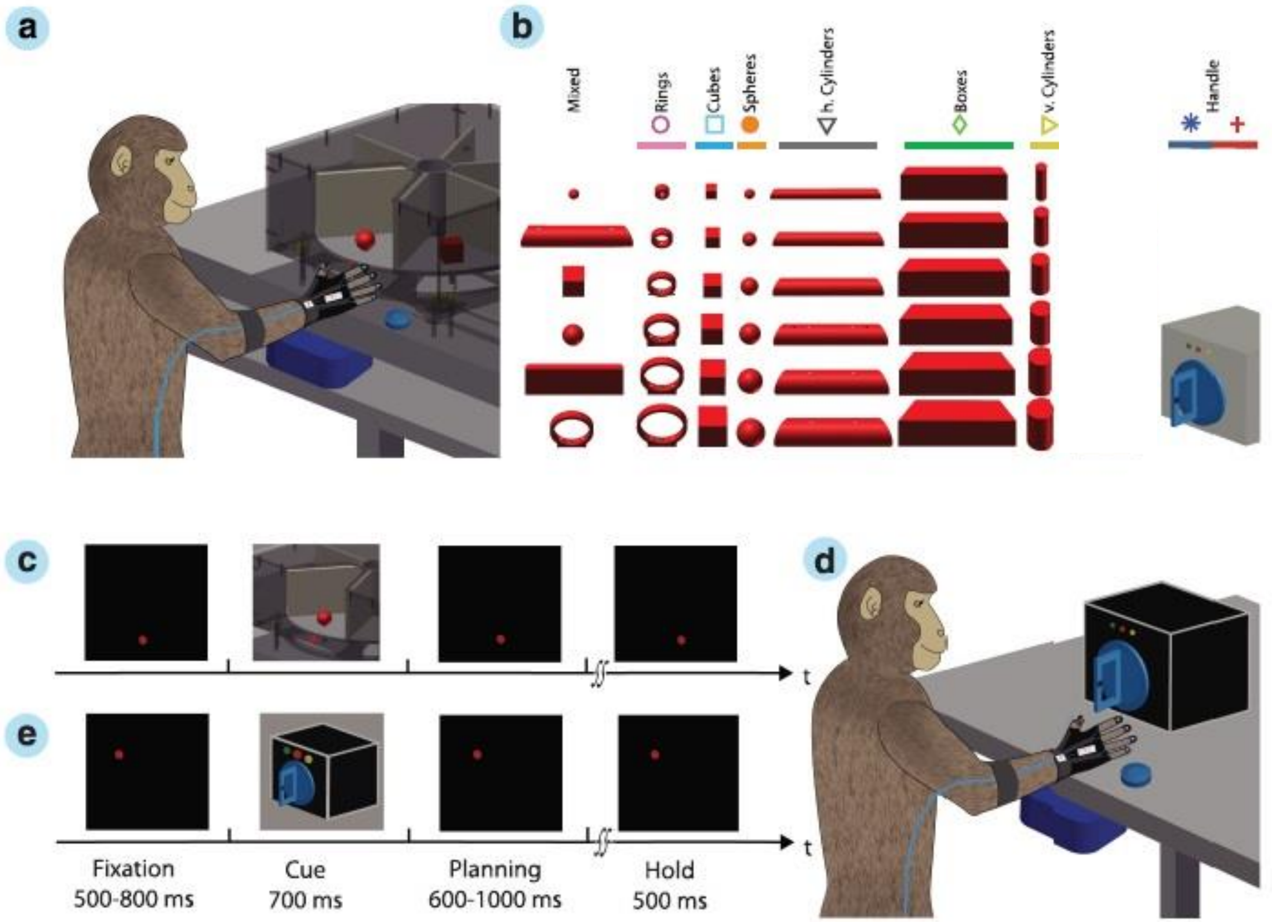

S. Schaffelhofer, A. Agudelo-Toro, and H. Scherberger, "Decoding a Wide Range of Hand Configurations from Macaque Motor, Premotor, and Parietal Cortices," *Journal of Neuroscience*, vol. 35, no. 3, pp. 1068–1081, Jan. 2015, doi: 10.1523/JNEUROSCI.3594-14.2015.

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044
linksfoundation.com
COPYRIGHT ©2021 LINKS

Dataset structure

Recording sessions, trials, metadata

NHP identifier	Dataset identifier	# Channels	# Trials
M	MRec40	552	745
	MRec41	568	757
	MRec42	554	653
Z	ZRec32	391	687
	ZRec35	388	724
	ZRec50	369	610

Recording sessions

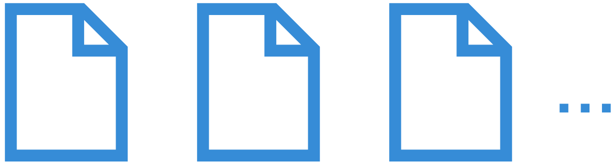

Trials

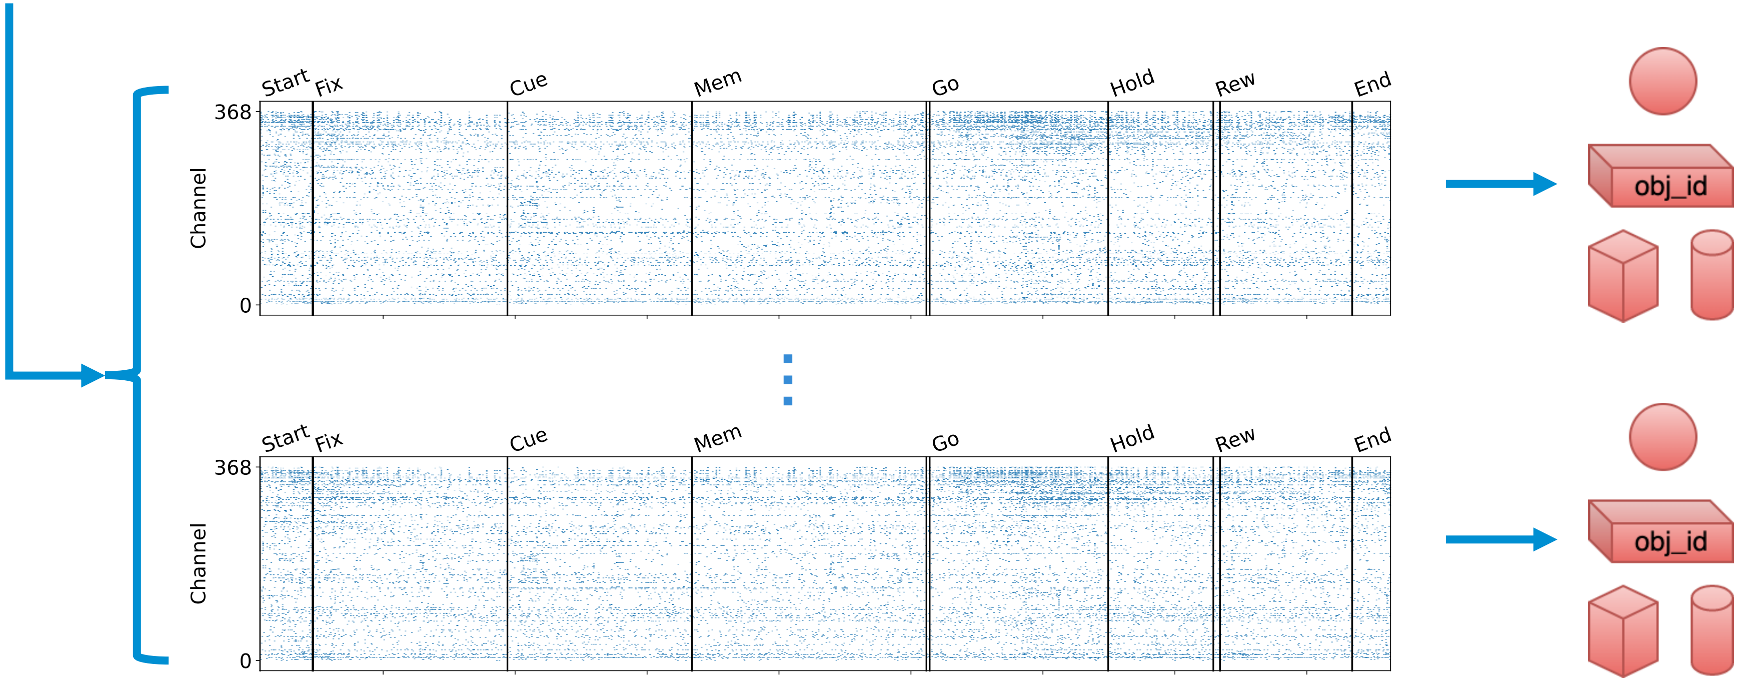

50 objects

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

Pre-processing

Time-discretization

- Each trial discretised in bins of 40ms
- Results in dense matrix
  - *channels x time_bins*
  - *Each bin contains the number of spikes*
- Time stamps associated to a bin
- Each matrix is Associated to object id

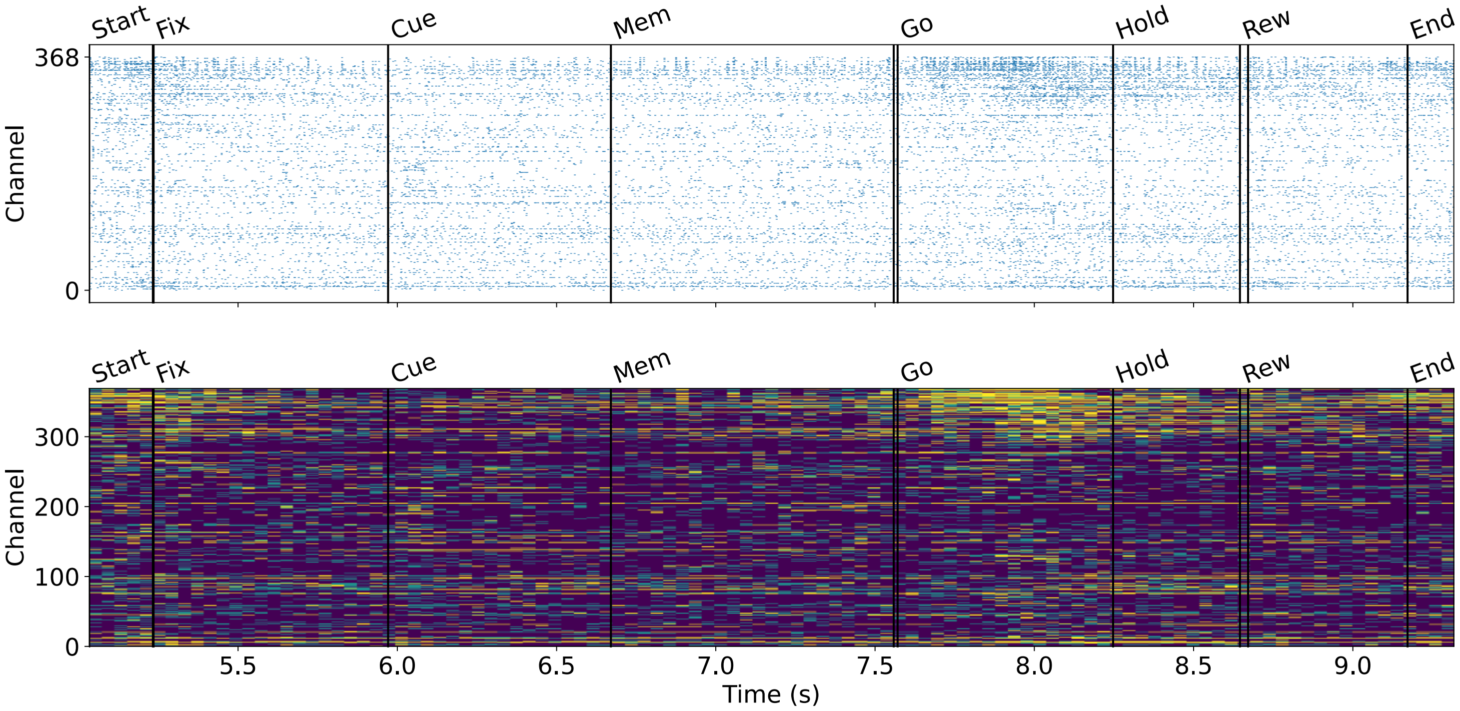

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

COPYRIGHT ©2021 LINKS

Pre-processing

Dataset split

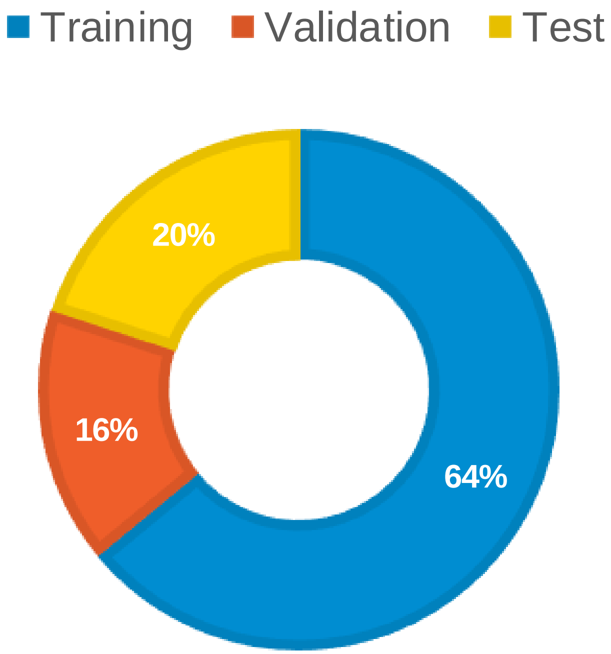

1. Classes are **de-duplicated** (i.e., same object in different groups)
2. Dataset is **split by trial**
    1. 80% training + validation, 20% test)
    2. stratified by class
3. Split is stratified by class
4. Under-represented class are removed (i.e., less than 3 trials per session)
5. Total of 39 classes left

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044
linksfoundation.com
COPYRIGHT ©2021 LINKS

Pre-processing

Sequence extraction

Goal: simulate real-time decoding

1. Sliding windows extract sub-sequnces of 12 bins
2. Grasp phase classification includes all data labelled with 0 (no grasp) and 1 (grasp)
3. Object classification includes only the sub-sequences in the grasp phase, labelled with object ID

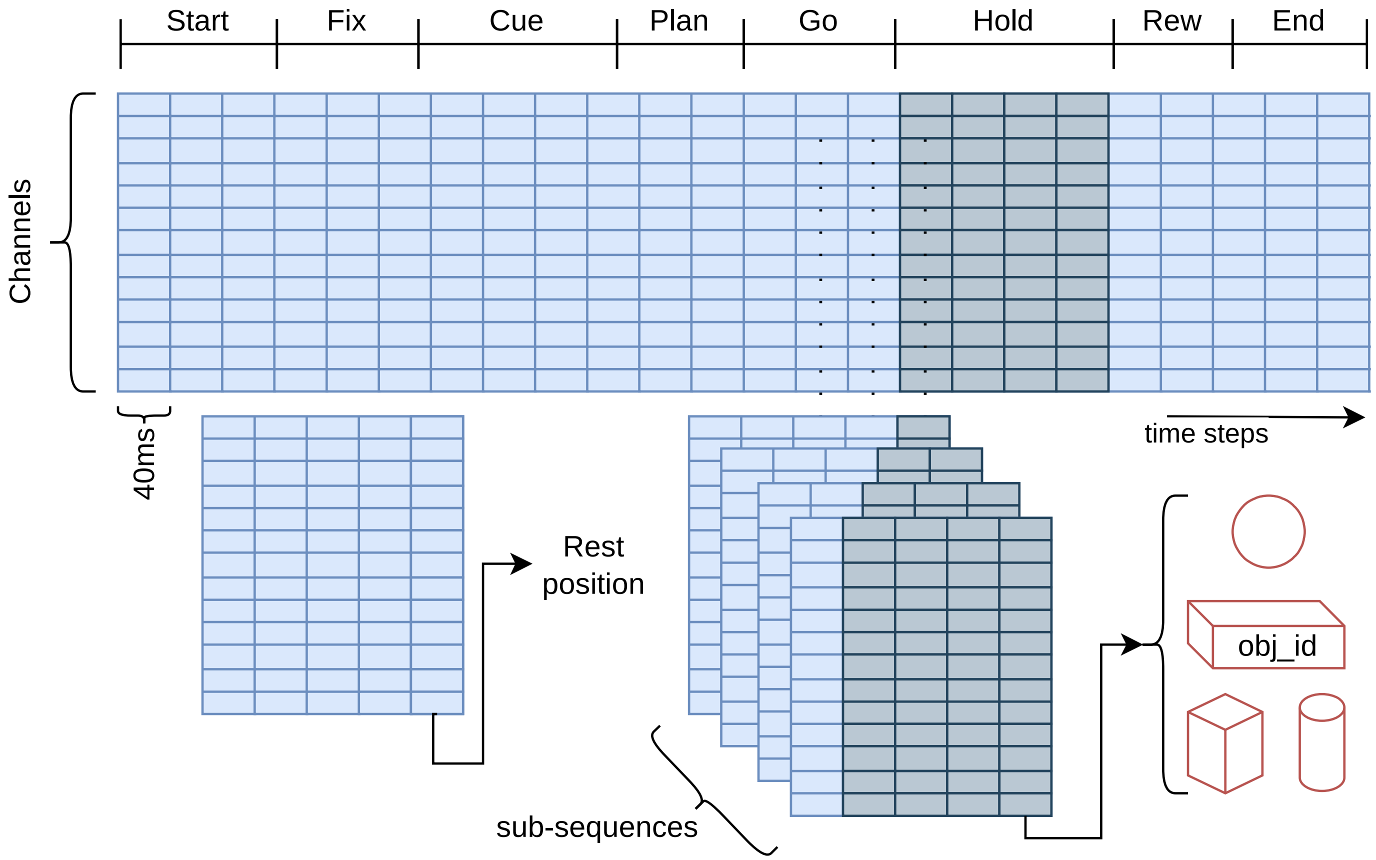

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

Decoding approach

Model selection

- LSTMs demonstrated good performance in literature
- Started from there (quite naïve in fact)
- **Bidirectionality** proved to be very effective
- Lots of **dropout** to prevent heavy overfitting
- Hyperparameter optimization

Hyperparameter	Values (L1=0.01, L2=0.01)	Selected M	Selected Z
LSTM layers	{ 1, 2, 3, 4 }	2	1
Hidden units	{ 16, 32, 40, 64 }	40	40
Dropout	{ 0, 0.2, 0.4, 0.6, 0.7, 0.8 }	0.8	0.7
Kernel regularisation	{ None, L1, L2, L1 + L2 }	L2	L2
Recurrent regularization	{ None, L1, L2, L1 + L2 }	L2	L1+L2
Initial learning rate	$\{ 10^{-3}, 2 \cdot 10^{-4}, 10^{-4} \}$	10^{-3}	10^{-3}

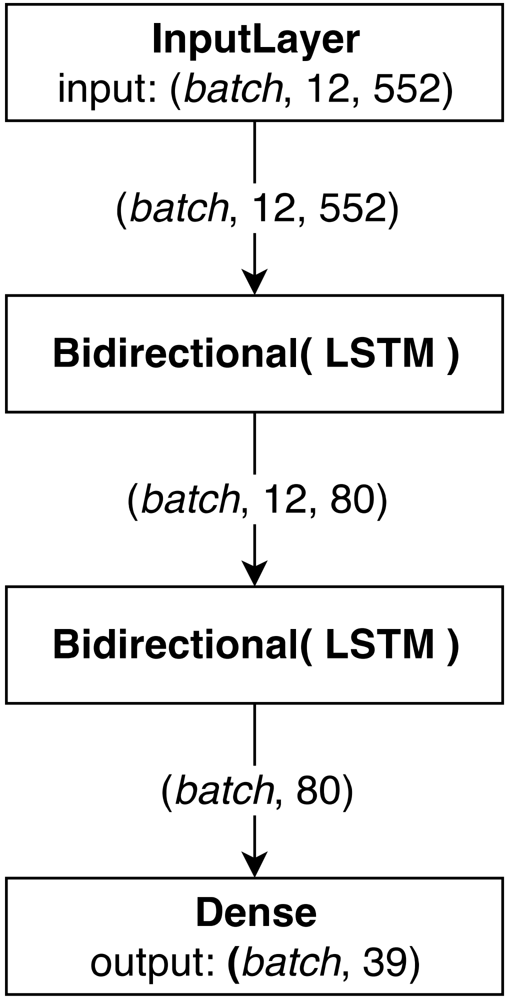

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

Results

Grasping phase detection

- Significant class imbalance here, with no grasp class ~10 times more represented than grasp class
- LSTM model reaches an accuracy of at least 98% for all datasets, the F1 score is always greater than 0.95
- Result relevant for a finite-state prosthesis control scenario
  - Each 40ms time step: predict prosthesis status (grasp vs. no-grasp)
  - 1% of unwanted movement for each time step
  - 0.1% of unresponsive prosthesis

(a) Confusion matrix for grasping phase detection (MRec40).

True	Predicted rest	Predicted grasp
rest	12824	173
grasp	25	1347

Unwanted Movement ~1%

Unresponsive prosthesis ~0.1%

(b) Grasping phase detection accuracy metrics.

Dataset id	Accuracy	F1 score
MRec40	99%	0.96
MRec41	99%	0.96
MRec42	99%	0.97
ZRec32	99%	0.96
ZRec35	98%	0.96
ZRec50	98%	0.95

Funded by the European Union

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

CRATOS

FONDAZIONE LINKS PASSION FOR INNOVATION linksfoundation.com COPYRIGHT ©2021 LINKS

Results

Object classification

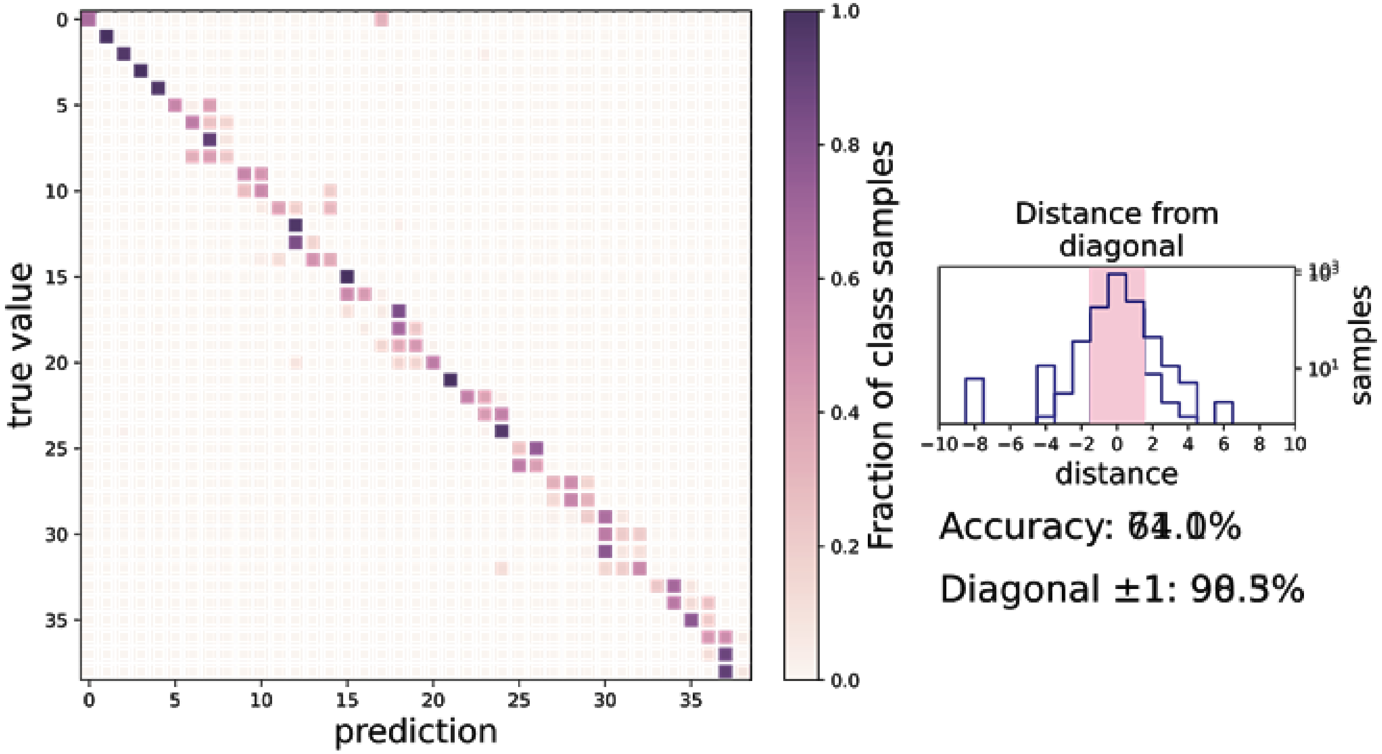

- Top 1% accuracy ranging from 61% to 74% across different sessions
- Neighbouring classes referred to similar objects
  - Relaxed accuracy defined as misclassification by 1 class (i.e., distance 1 from diagonal)

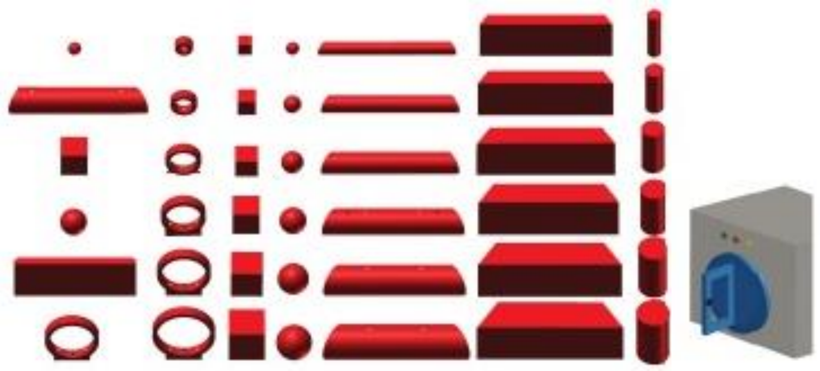

This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

linksfoundation.com
COPYRIGHT ©2021 LINKS

Results

Object classification – comparison with state of the art

Animal	Metric	Present work	Schaffelhofer, 2015	Fabiani, 2021
M	Accuracy	$69.7 \pm 4\%$	$62.9 \pm 3.6\%$	n/a
Z	Accuracy	$62.3 \pm 1.2\%$	$61.4 \pm 4.1\%$	n/a
Global	Accuracy Relaxed accuracy	$65.9 \pm 4.9\%$ $94.4 \pm 3.1\%$	62% 86.5%	22% 59%

- Schaffelhofer et al. report an average accuracy for the hold phase of 62% over a total of 10 recording sessions (against the six available for this work)
  - Offline naive bayesian classifier applied to the whole hold phase vs. sligding window
  - Validated with a leave-one-out (LOO) approach
  - Dataset fraction used for training was significantly higher
  - Significantly better (M) and slightly better (Z) accuracy results in a harder set-up. Significantly better relaxed accuracy
- Fabiani reported offline and online accuracy figures
  - This work outperforms both cases, despite the use of similar LSTM architechtures
  - Improvement likely due to bidirectional networks and stronger regularization

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044
CRATOS
FONDAZIONE LINKS PASSION FOR INNOVATION
linksfoundation.com
COPYRIGHT ©2021 LINKS

Results

Robustness to reduction of training data

- Real life scenario involves frequent re-training, also during recording sessions
- Neuroscientist's requirement is to limit training to a minimum, and let the model work for most of the time
    - Even more relevant for patients
- Smaller training sets were evaluated
- The model is still outperforming the previous SotA down to 40% of training data

Training + validation set	Accuracy	Relaxed accuracy
80%	74.1%	98.3%
70%	70.1%	97.1%
60%	69%	93.7%
50%	62.8%	94%
40%	63.8%	92.8%
30%*	59.2%	87.5%
20%**	51%	81%

*validation is 30% of training set to ensure at least one representative per class
**validation is 40% of training set

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

CRATOS
FONDAZIONE LINKS PASSION FOR INNOVATION
linksfoundation.com
COPYRIGHT ©2021 LINKS

Remarks and future work

- LSTM can match and outperform previous approaches on a known dataset
  - More accurate prosthesis control can be enabled by modern architectures, at the expense of some explainability
- Results for relaxed accuracy are very promising for the final application
  - Continuous control of few degrees of freedom (regression on few variables)
  - Finite-state control of 5-6 grasp types (classification on fewer classes)

Future work

- Neural data evolves with time: robustness/adaptability of models is key to reduce patients and physicians effort
- Fine-tuning with limited data already proven effective
- Semi-supervised training with usage data is an open problem, inputs from the community are welcome!

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

COPYRIGHT ©2021 LINKS

Thanks for your attention!

Acknowledgments

Funded by the European Union
This project has received funding from the European Horizon 2020 research and innovation programme under GA No 965044

COPYRIGHT ©2021 LINKS