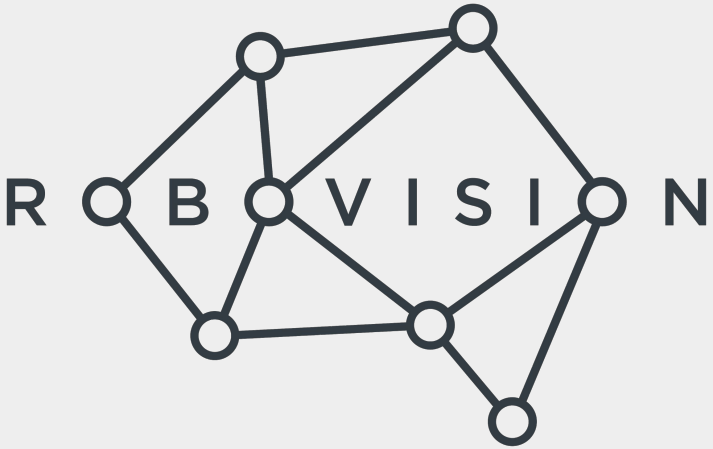

AI WINTER

IS *NOT* COMING

AGENDA

1. **getting real with AI**
2. AI winter
3. AI summer

Is AI a robotman standing in front of an extremely large (touch?)screen with random zeros and ones, with some circuitboards thrown in?

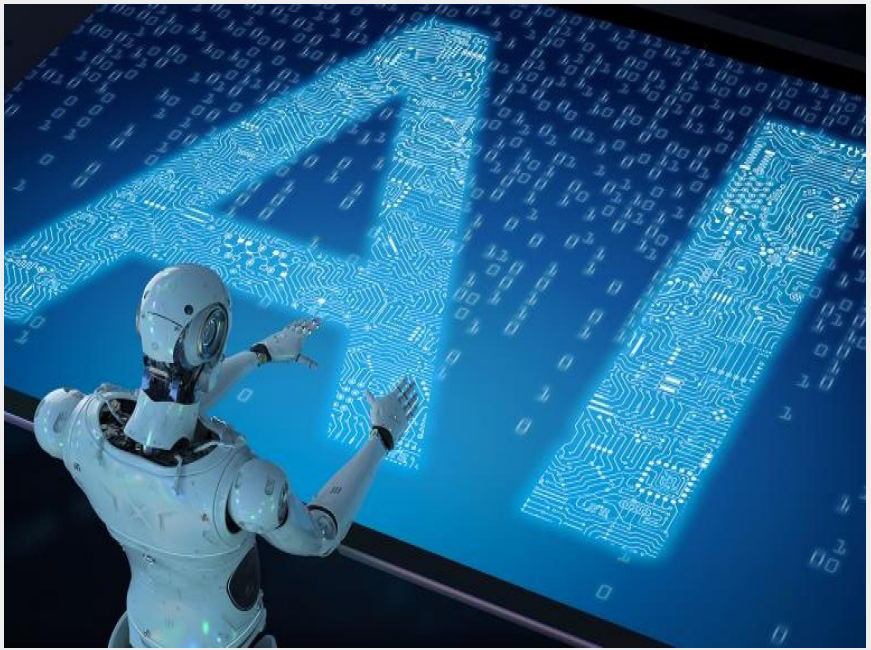

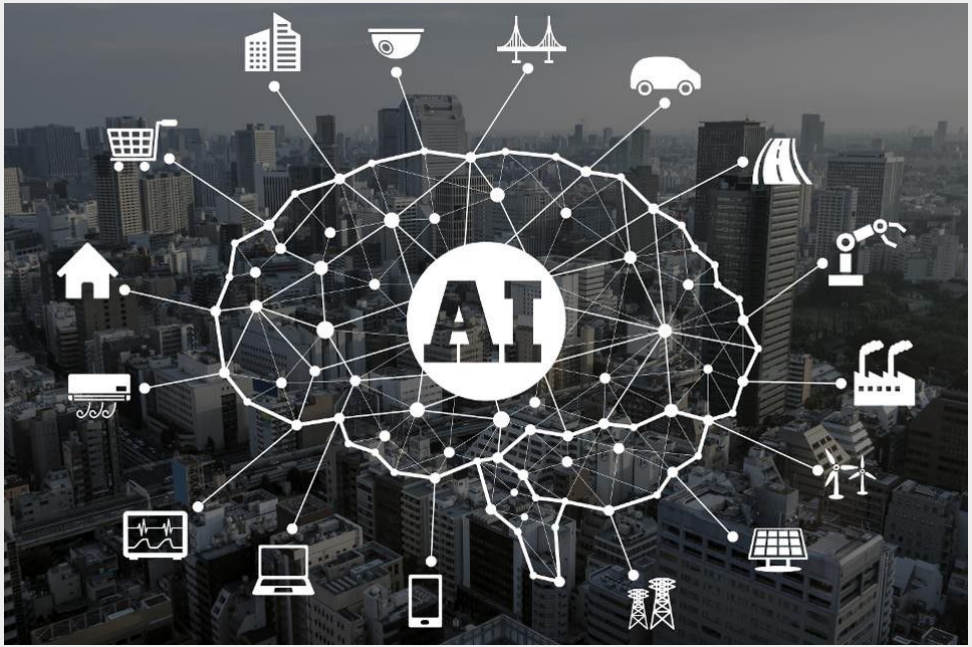

Is AI a brainlike structure that is connected with random images that don't tell us anything about the practical case?

WHAT IS ARTIFICIAL INTELLIGENCE?

PROBLEM

Intuitive knowledge is hardly transferable. Building an intelligent agent through explicit hardcoding of insights is unmanageable.

SOLUTION

A.I. is machine intelligence gathered by experience, without explicitly programming where to look.

DEEP LEARNING | creating artificial intelligence

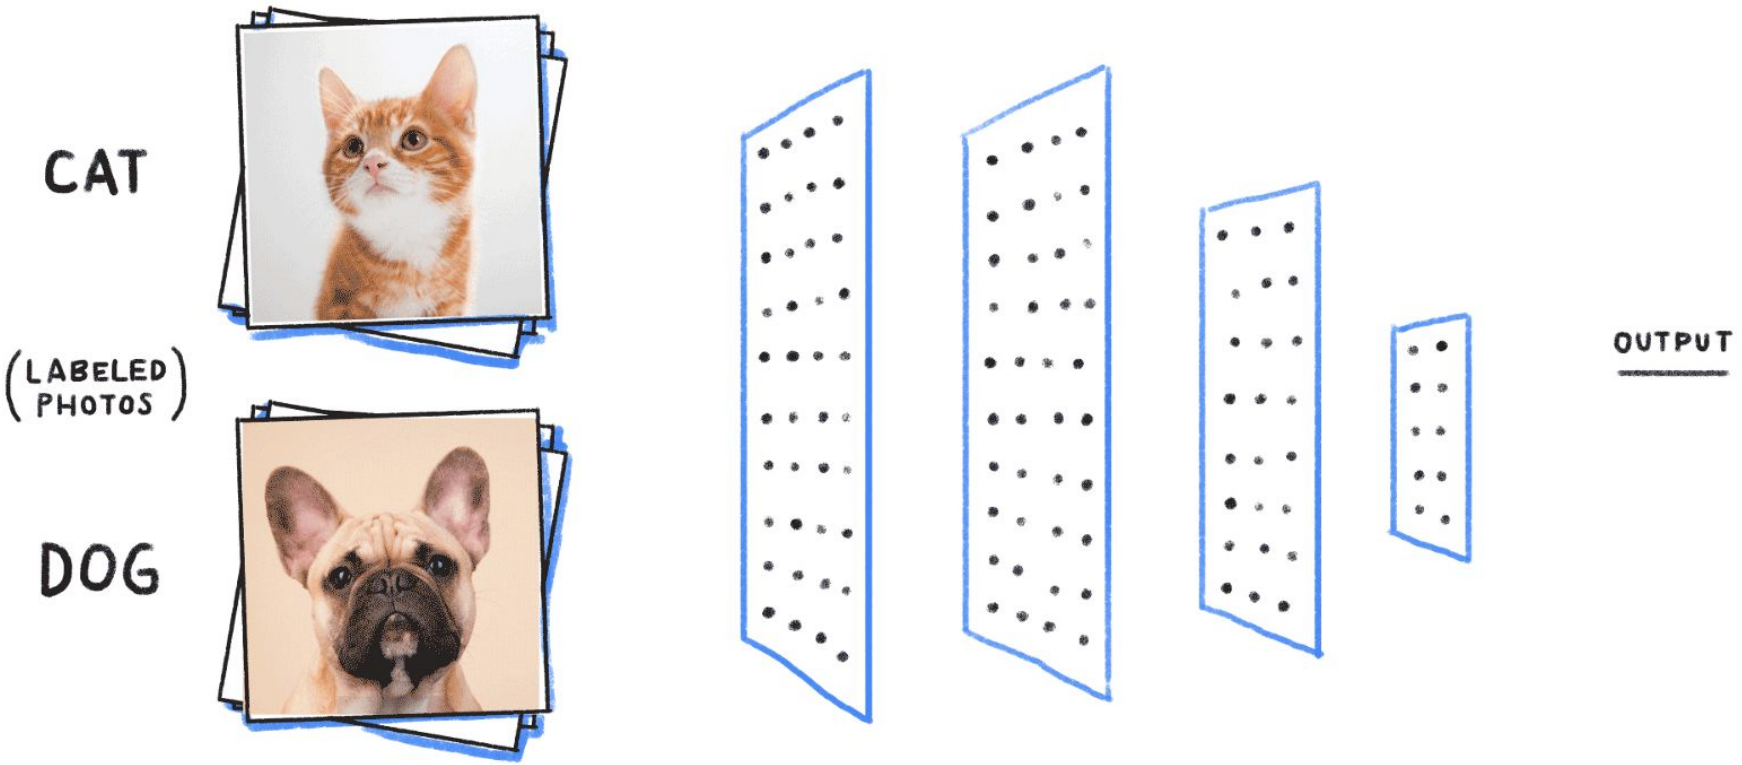

AGENDA

1. getting real with AI
2. **AI winter**
3. AI summer

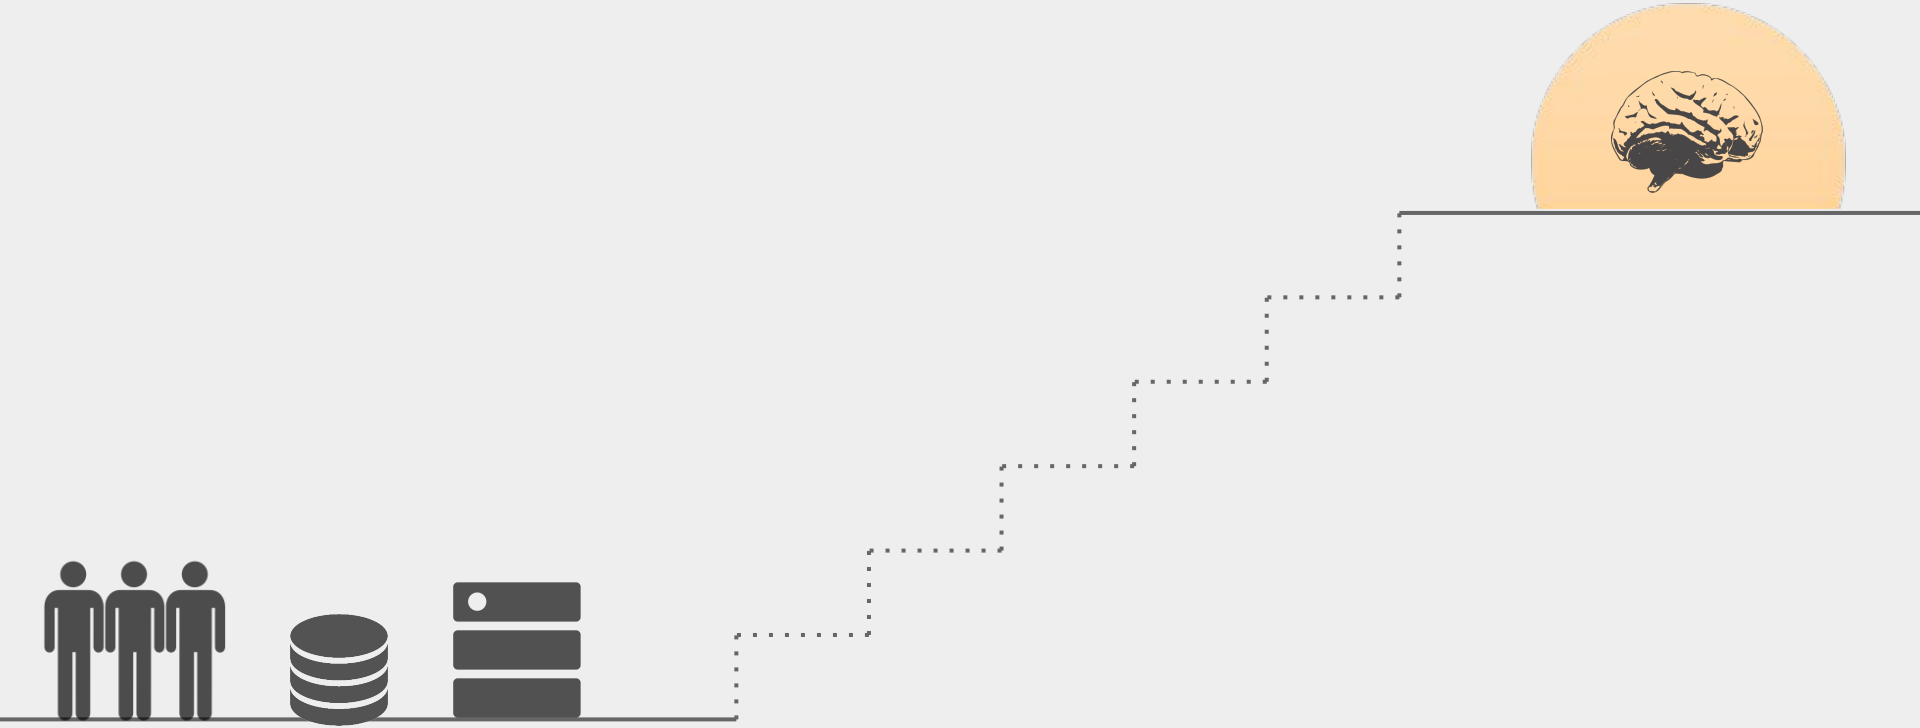

Every AI team feels the same pains.

Many projects get **lost in translation**.

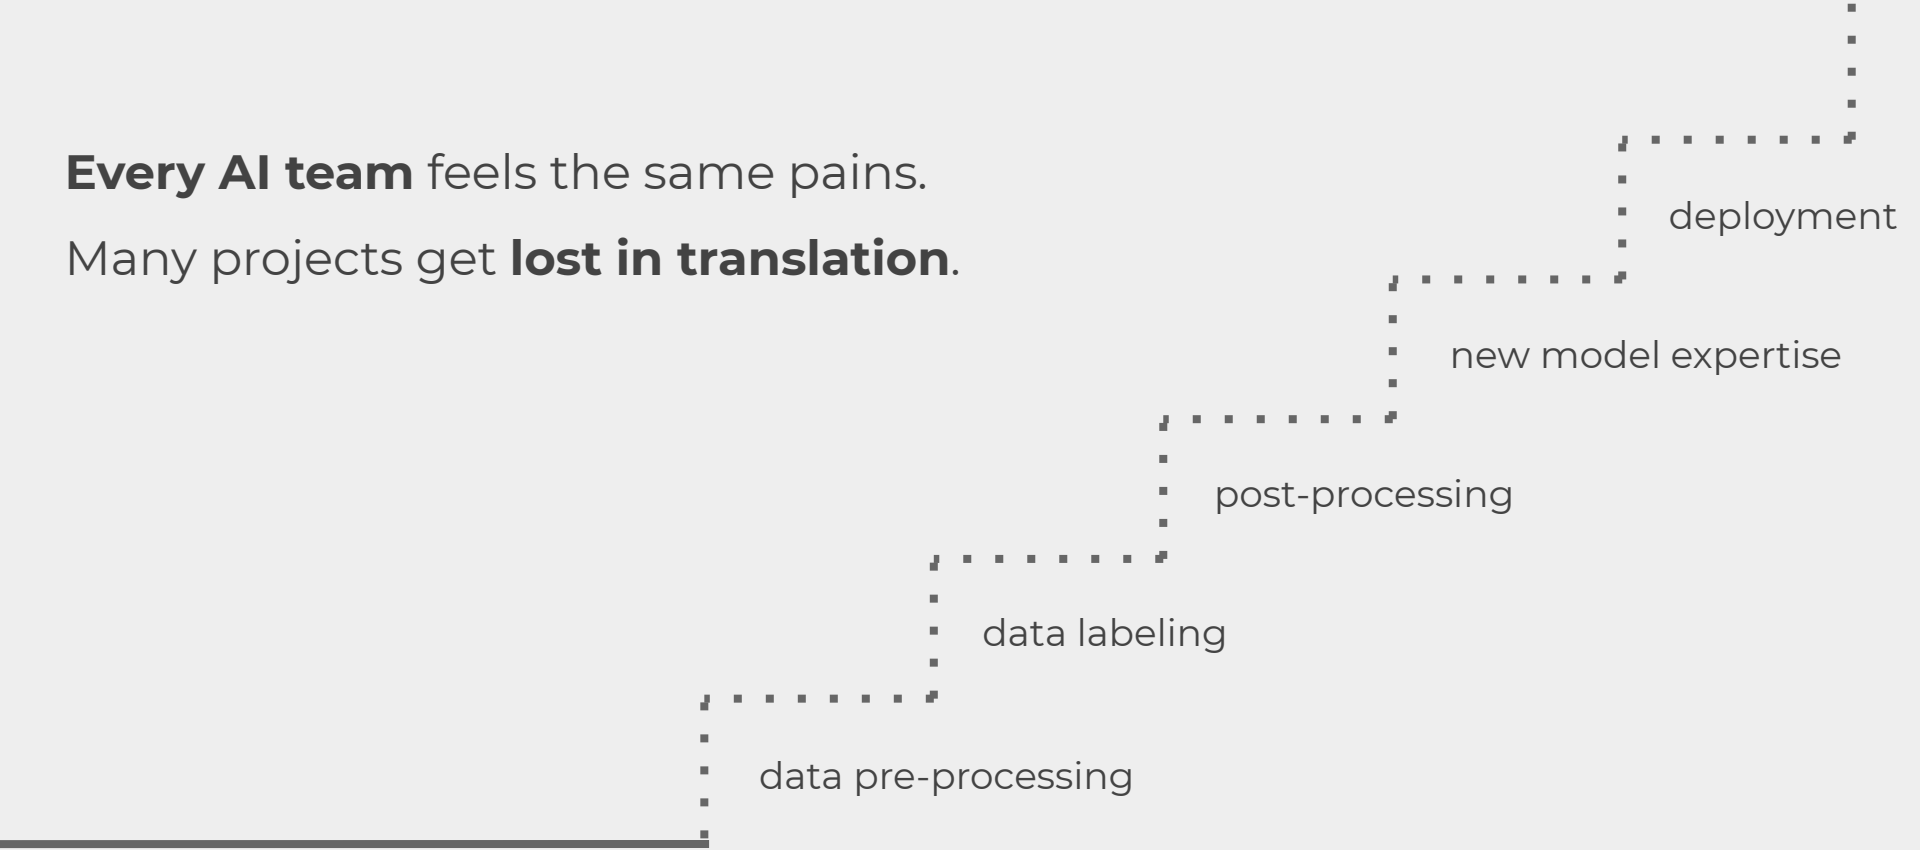

AGENDA

1. getting real with AI
2. AI winter
3. **AI summer**

DEEP LEARNING | decision type

Image classification

Is there a deer in the image?

Object detection

Where in the image is the deer?

Image segmentation

Where exactly is the deer? What pixels?

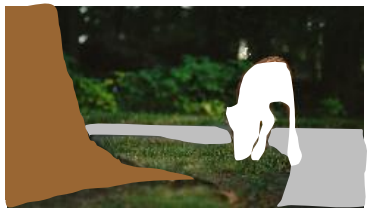

Image Similarity

Which images are similar to the query image?

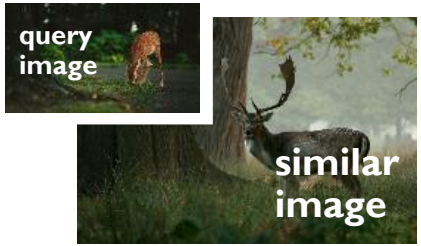

THE GAP | getting real with AI

FROM ...

“throw some AI on that problem”

“let’s do something with AI”

TO ...

“we did it! that wasn’t as hard as I thought!”

“I completely understand how it works now”

THE GAP | creating artificial intelligence

1 clear problem definition

2 break it down, and again

3 collect the necessary info

4 let AI take the decisions

FROM ...

“throw some AI on that problem”

“let’s do something with AI”

TO ...

“we did it! that wasn’t as hard as I thought!”

“I completely understand how it works now”

EXAMPLE | Amazon Go

EXAMPLE | Amazon Go

1

clear problem definition

"We need a system that knows which products are taken home by which customers."

2

break it down, and again

"We need a system that knows:
1. ***which products***
2. *are **taken** home*
3. *by **which customers**."*

3

collect the necessary info

1. *product recognition*
2. *action recognition*
3. *people detection + tracking*

4

let AI take the decisions

1. *deep learning model A*
2. *deep learning model B*
3. *deep learning model C + D*

EXAMPLE | Amazon Go

1	2	3	4
clear problem definition	break it down, and again	collect the necessary info	let AI take the decisions
			
Management Consulting	*Management Consulting*	*Data science*	*Robovision AI*

ROBOVISION AI PLATFORM | AI made easy

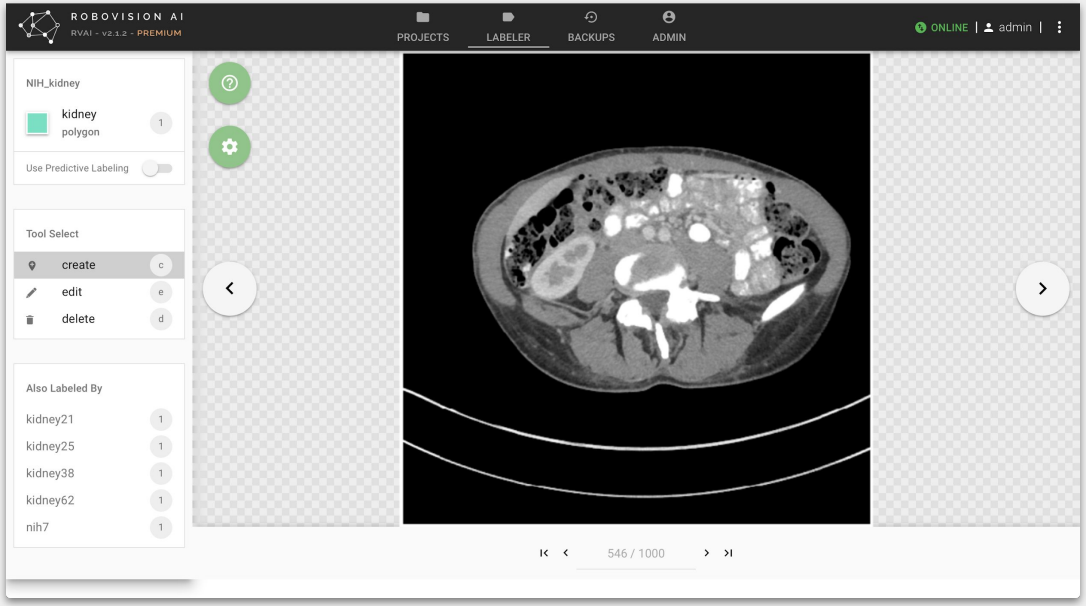

4

let AI take the decisions

Robovision AI

we put silicon intelligence to work in three main domains

agriculture

automation

manufacturing

quality control
automation
monitoring

people intelligence

safety
surveillance
security

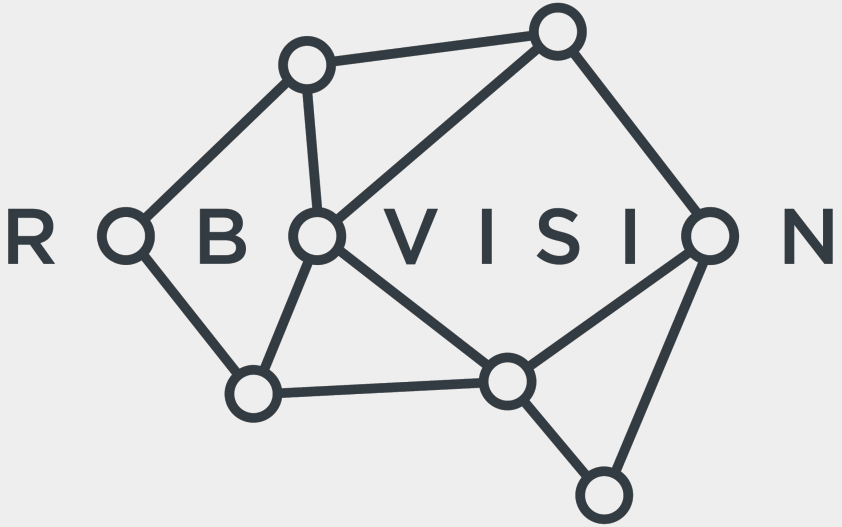

use case

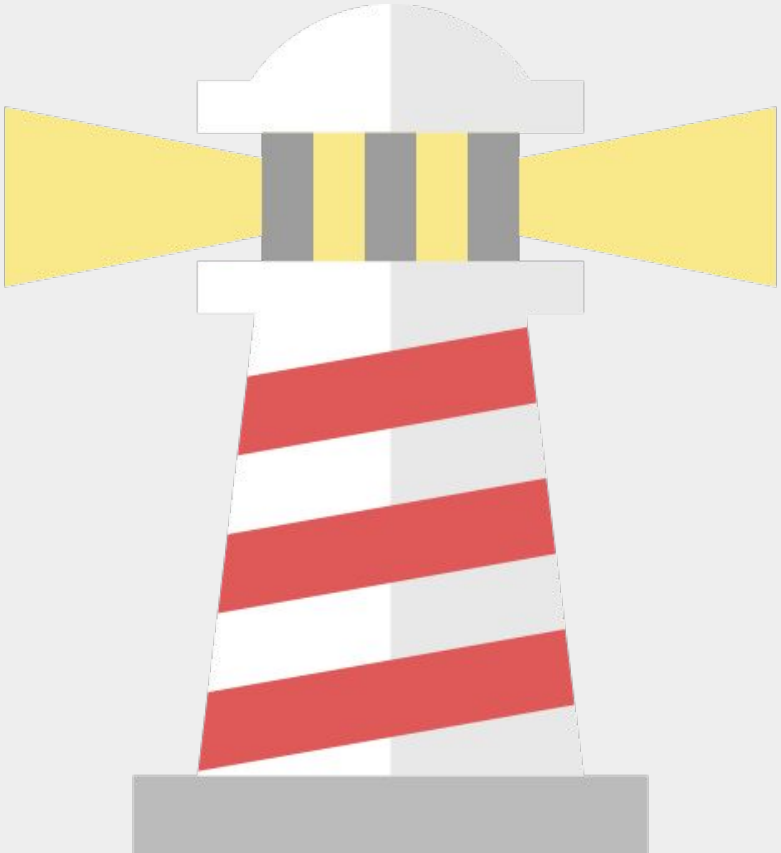

Staying relevant through lighthouses

We work with the NIH to develop 3D deep learning models of organs using crowd labelling.

National Institutes of Health

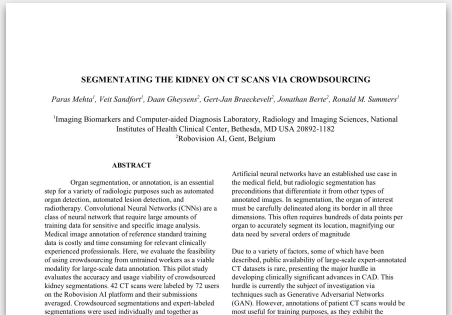

SEGMENTATING THE KIDNEY ON CT SCANS VIA CROWDSOURCING

Paras Mehta[1], Veit Sandfort[1], Daan Gheysens[2], Gert-Jan Braeckevelt[2], Jonathan Berte[2], Ronald M. Summers[1]

[1]Imaging Biomarkers and Computer-aided Diagnosis Laboratory, Radiology and Imaging Sciences, National Institutes of Health Clinical Center, Bethesda, MD USA 20892-1182
[2]Robovision AI, Gent, Belgium

ABSTRACT

Organ segmentation, or annotation, is an essential step for a variety of radiologic purposes such as automated organ detection, automated lesion detection, and radiotherapy. Convolutional Neural Networks (CNNs) are a class of neural network that require large amounts of training data for sensitive and specific image analysis. Medical image annotation of reference standard training data is costly and time consuming for relevant clinically experienced professionals. Here, we evaluate the feasibility of using crowdsourcing from untrained workers as a viable modality for large-scale data annotation. This pilot study evaluates the accuracy and usage viability of crowdsourced kidney segmentations. 42 CT scans were labeled by 72 users on the Robovision AI platform and their submissions averaged. Crowdsourced segmentations and expert-labeled segmentations were used individually and together as

Artificial neural networks have an established use case in the medical field, but radiologic segmentation has preconditions that differentiate it from other types of annotated images. In segmentation, the organ of interest must be carefully delineated along its border in all three dimensions. This often requires hundreds of data points per organ to accurately segment its location, magnifying our data need by several orders of magnitude

Due to a variety of factors, some of which have been described, public availability of large-scale expert-annotated CT datasets is rare, presenting the major hurdle in developing clinically significant advances in CAD. This hurdle is currently the subject of investigation via techniques such as Generative Adversarial Networks (GAN). However, annotations of patient CT scans would be most useful for training purposes, as they exhibit the

ORGAN SEGMENTATION | context

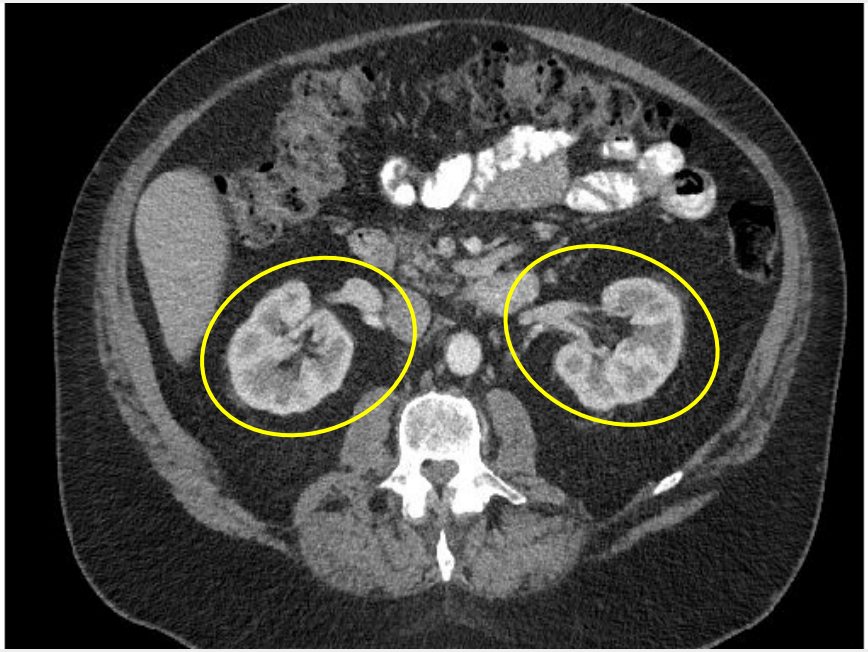

STEP 1 | data & annotations

STEP 2 | model selection & training + validation

	ID	Algorithm	Name	Progress
☐	3	instancesegmenta	NIH_kidney	DATA 1000 img. · LABELS 100% labeled · MODEL 81.46% performance
☐	28	instancesegmenta	NIH_val	DATA 919 img. · LABELS 0% labeled · MODEL No model trained
☐	33	instancesegmenta	NIH_training	DATA 0 img. · LABELS 0% labeled · MODEL No model trained

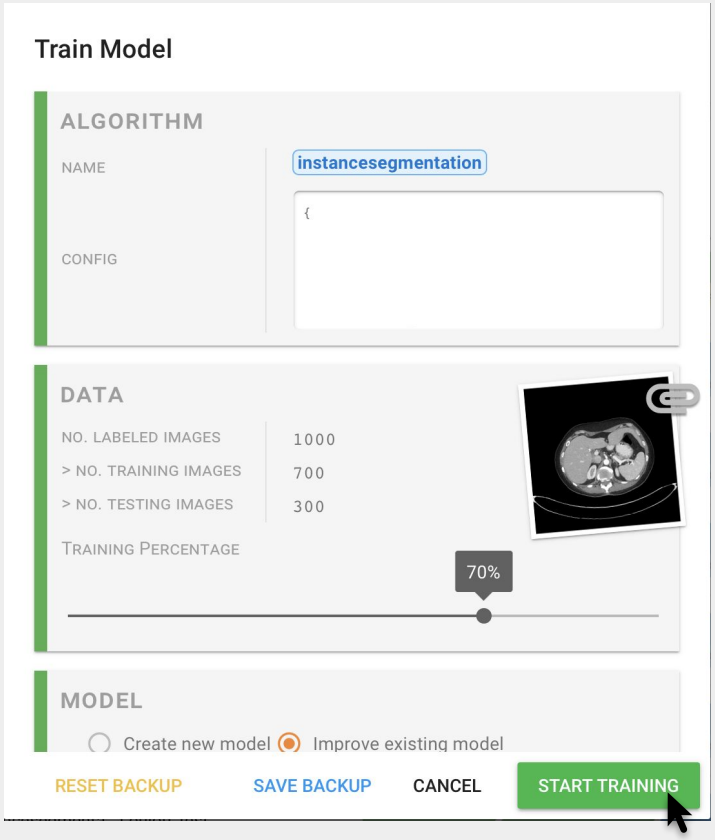

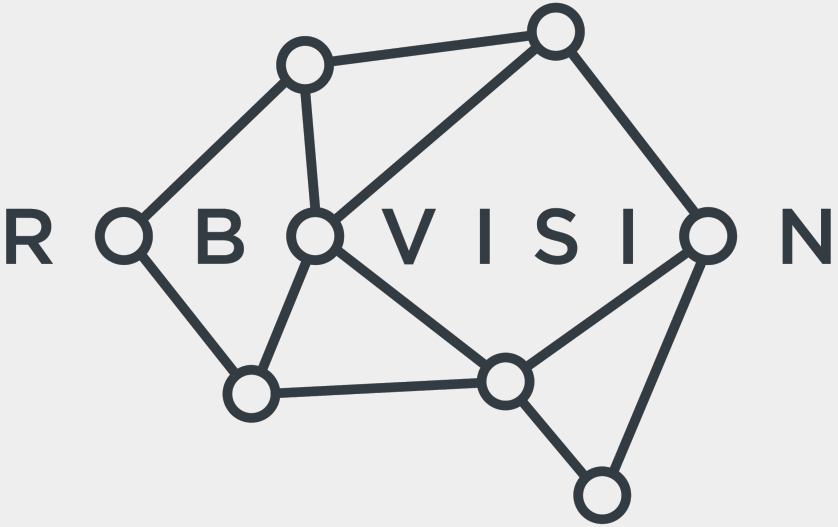

closing thoughts

THE NEXT WAVE | owning intelligence

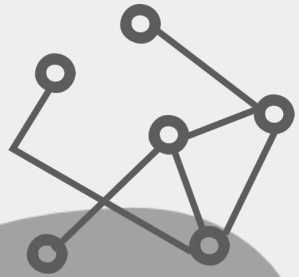

ROBOVISION

AISTORE

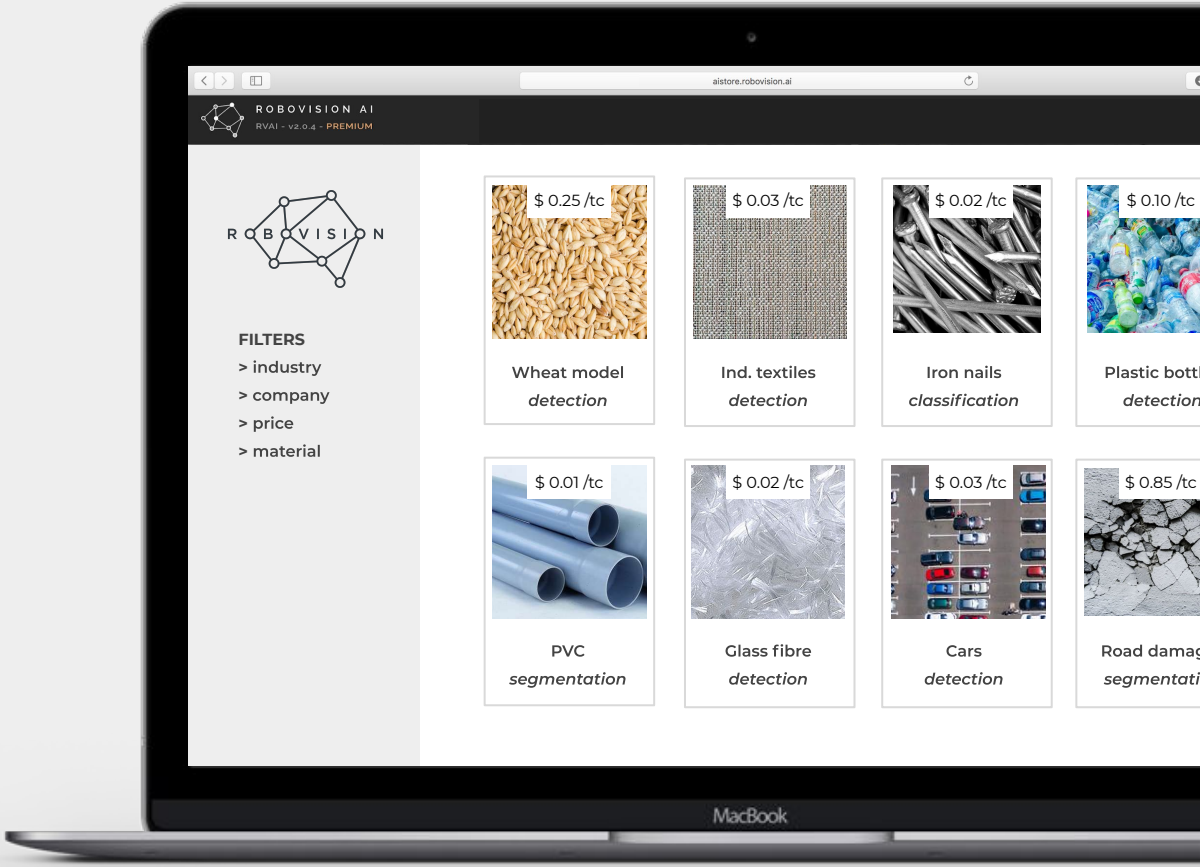

And we are already fueling AI model trading worldwide.

TIPS & TRICKS | qualifying use cases

1. If a person can easily spot it, then deep learning can as well
2. The quality of your data determines everything (see point 1)
3. Make sure you have a dataset of about 1,000-2,000 images to start with
4. If your model is not strong enough yet, then just collect more data
5. If the problem is too complex, break it down into smaller decisions
6. You can train your first model on the first day
7. Getting to 99% accuracy is 10x harder than getting to 90% accuracy
8. Use AI as a tool: you might not solve the total problem, but a hybrid might

Daan Gheysens
+32 475 385 528
daan.gheysens@robovision.eu

NVIDIA ELITE PARTNER

thank you.