CogniSight

IMAGE LEARNING AND RECOGNITION WITH NEUROMEM TRAINABLE CLASSIFIER

Inspect, Sort

Identify, Track

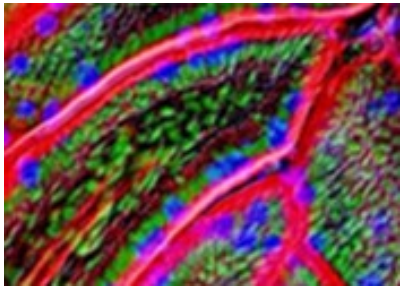

Detect, Count

Search, Tag

Match, Compare

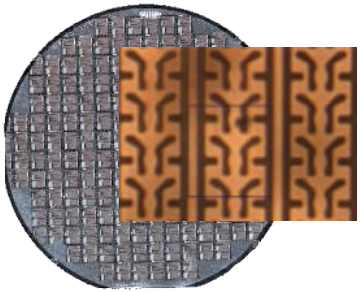

Find, Retrieve

Applications

What can a CogniSight engine do?

A CogniSight engine can be trained
- Off-line on collected images and video
- In real-time on live video

A CogniSight engine can learn under supervision
- From annotations made by domain experts on collection of reference images
- From live interrupts triggered by an operator

A CogniSight engine can learn no supervision
- To re-inforce a knowledge getting less confident
- To build a dictionnary of viewlets describing reference images
- To build a reactive image memory frame

A CogniSight engine can rely on single or multiple visual experts to made a decision
- Experts can be trained to be redundant to minimize false positive
- Experts can be trained to be complementary to waive uncertainties wisely

CogniSight hardware eco-system

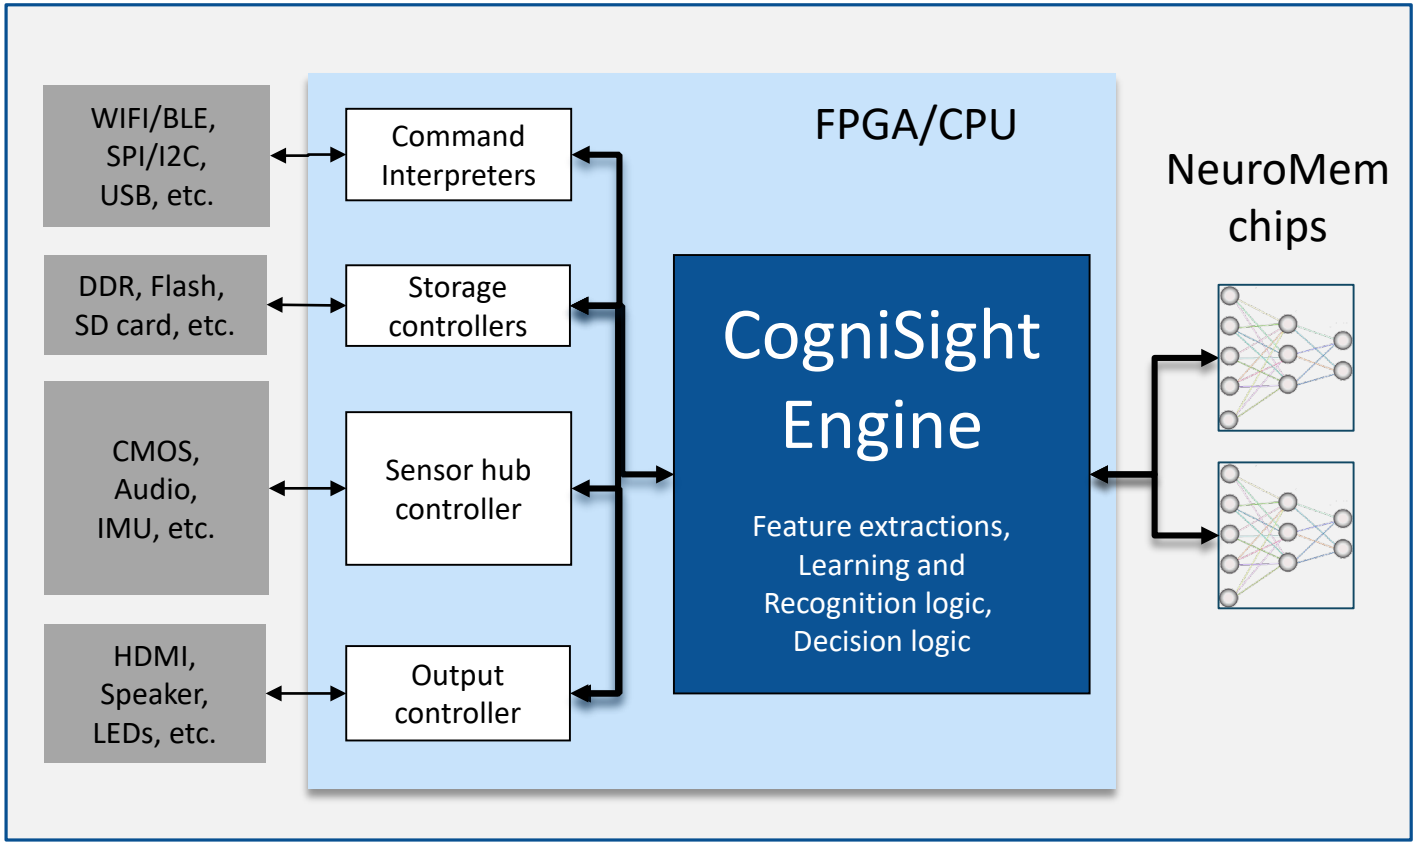

The big picture

Teaching is more practical than explaining “why” with rules and equations.

A good non-linear classifier enables such approach.

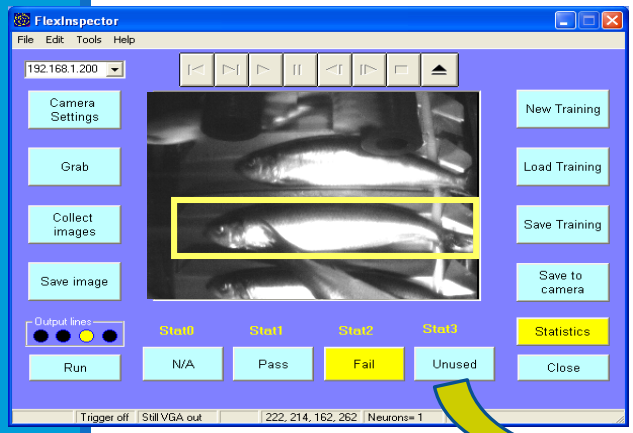

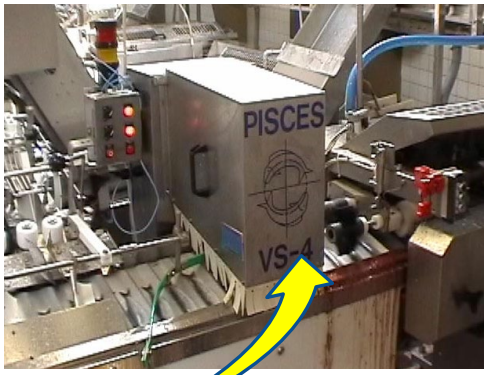

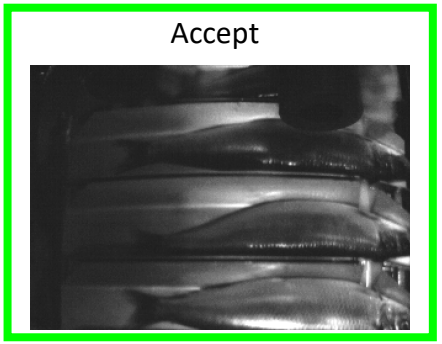

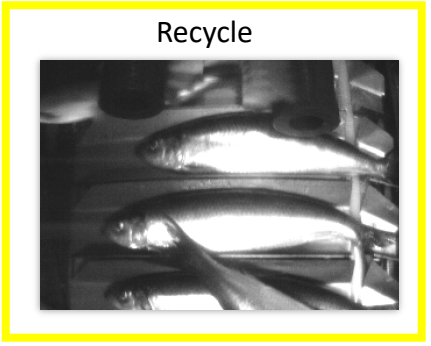

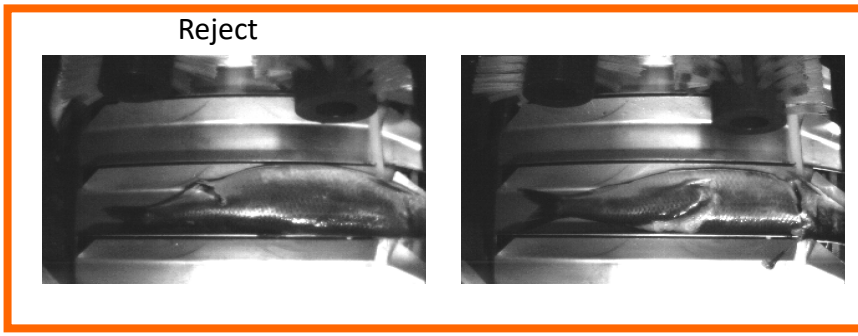

The teaching approach…

Let’s take the example of BGA counting….

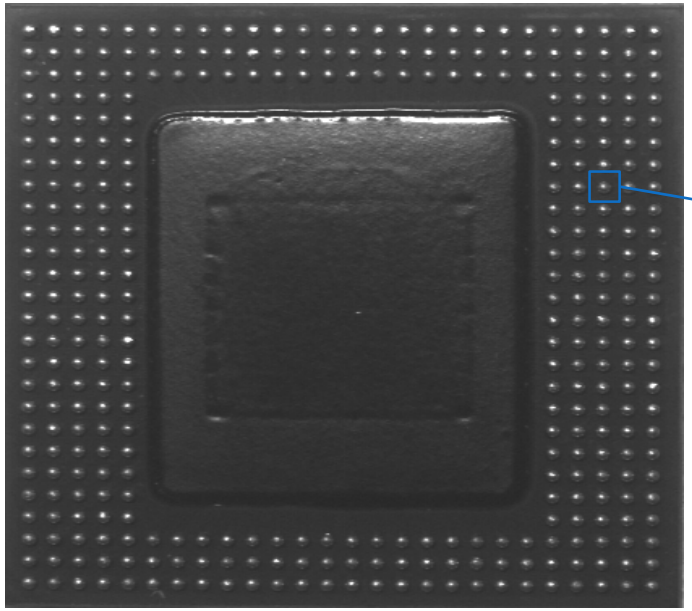

Pixel intensities

72	72	74	74	73	72	71	71	71	71	70	71
72	72	74	73	72	72	73	72	70	71	70	69
72	70	71	72	73	71	70	69	69	72	72	71
74	73	79	87	89	86	81	81	83	77	71	70
76	84	97	109	126	136	123	102	90	78	66	66
79	89	104	126	184	236	214	142	95	84	73	64
82	94	94	99	141	186	181	135	103	93	78	64
79	91	91	81	82	90	94	87	80	75	67	62
73	77	82	81	76	72	72	73	73	69	61	61
72	68	69	72	74	76	76	71	67	63	61	64

1) Conventional approach

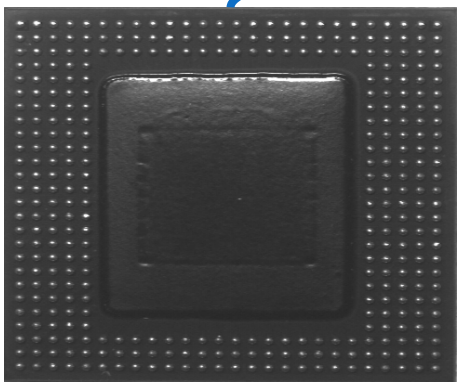

Image preprocessing

Thresholding

Morphological operations

Blob analysis

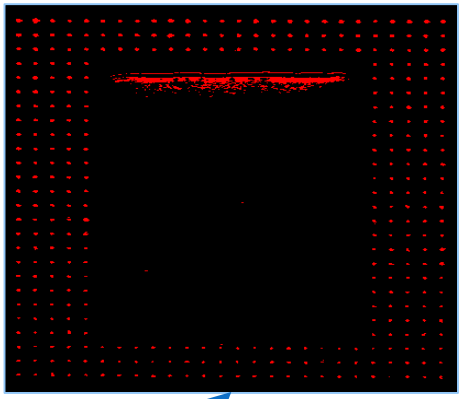

2) Non-linear NN approach

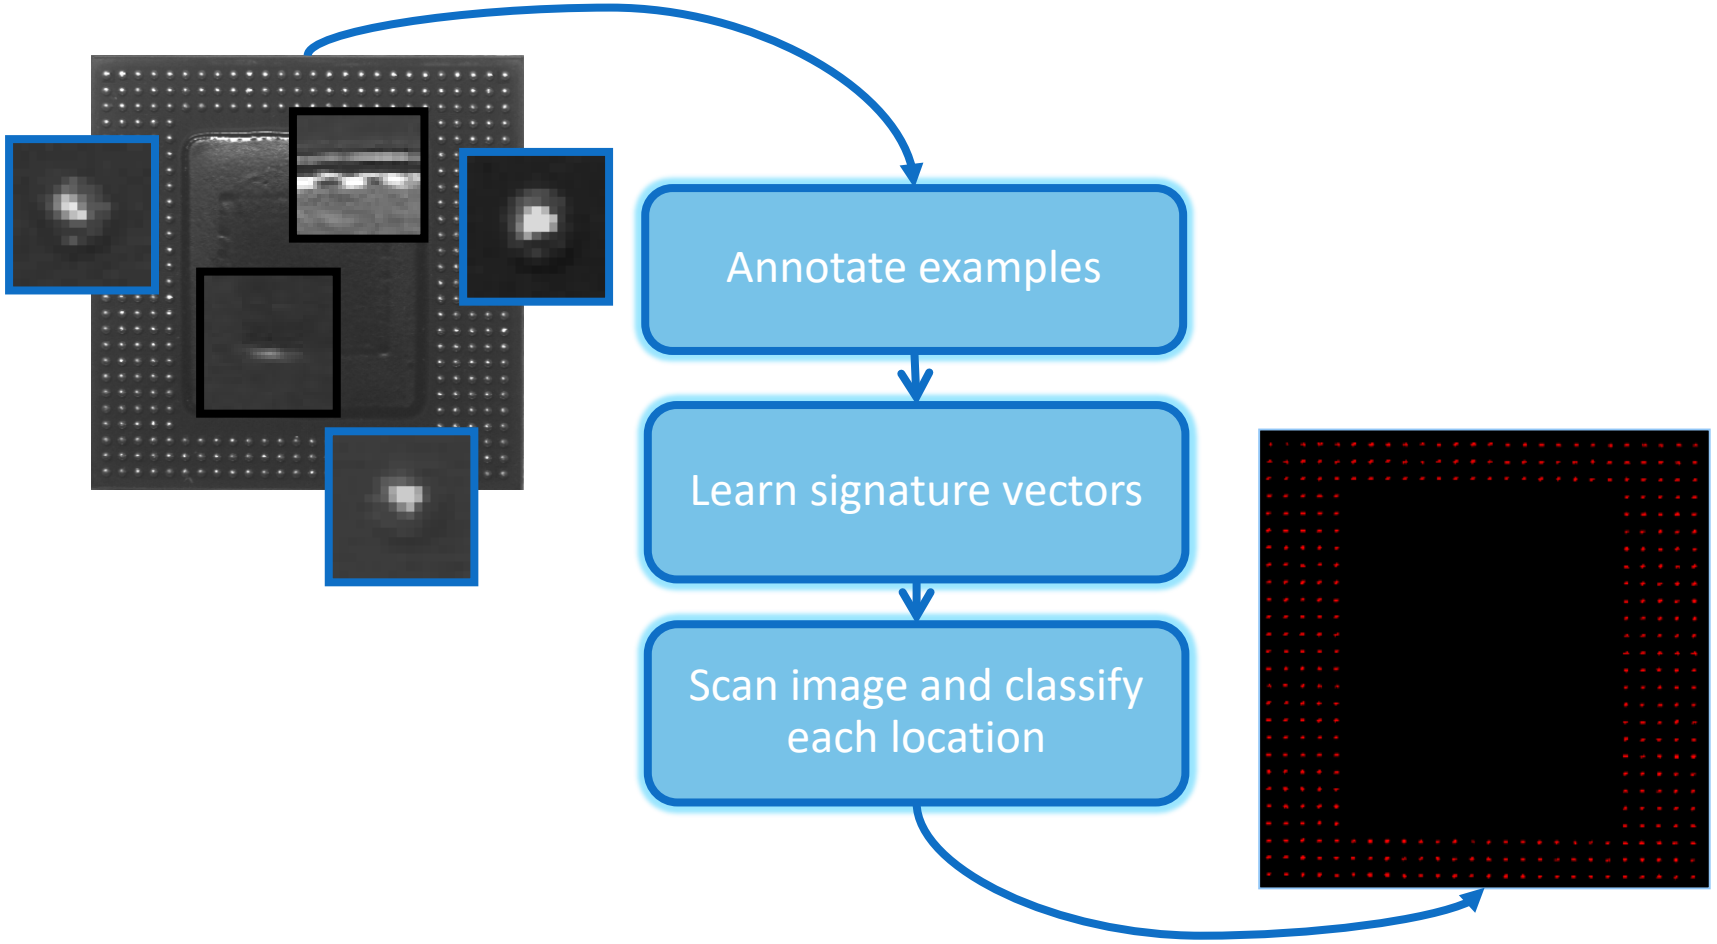

The data workflow

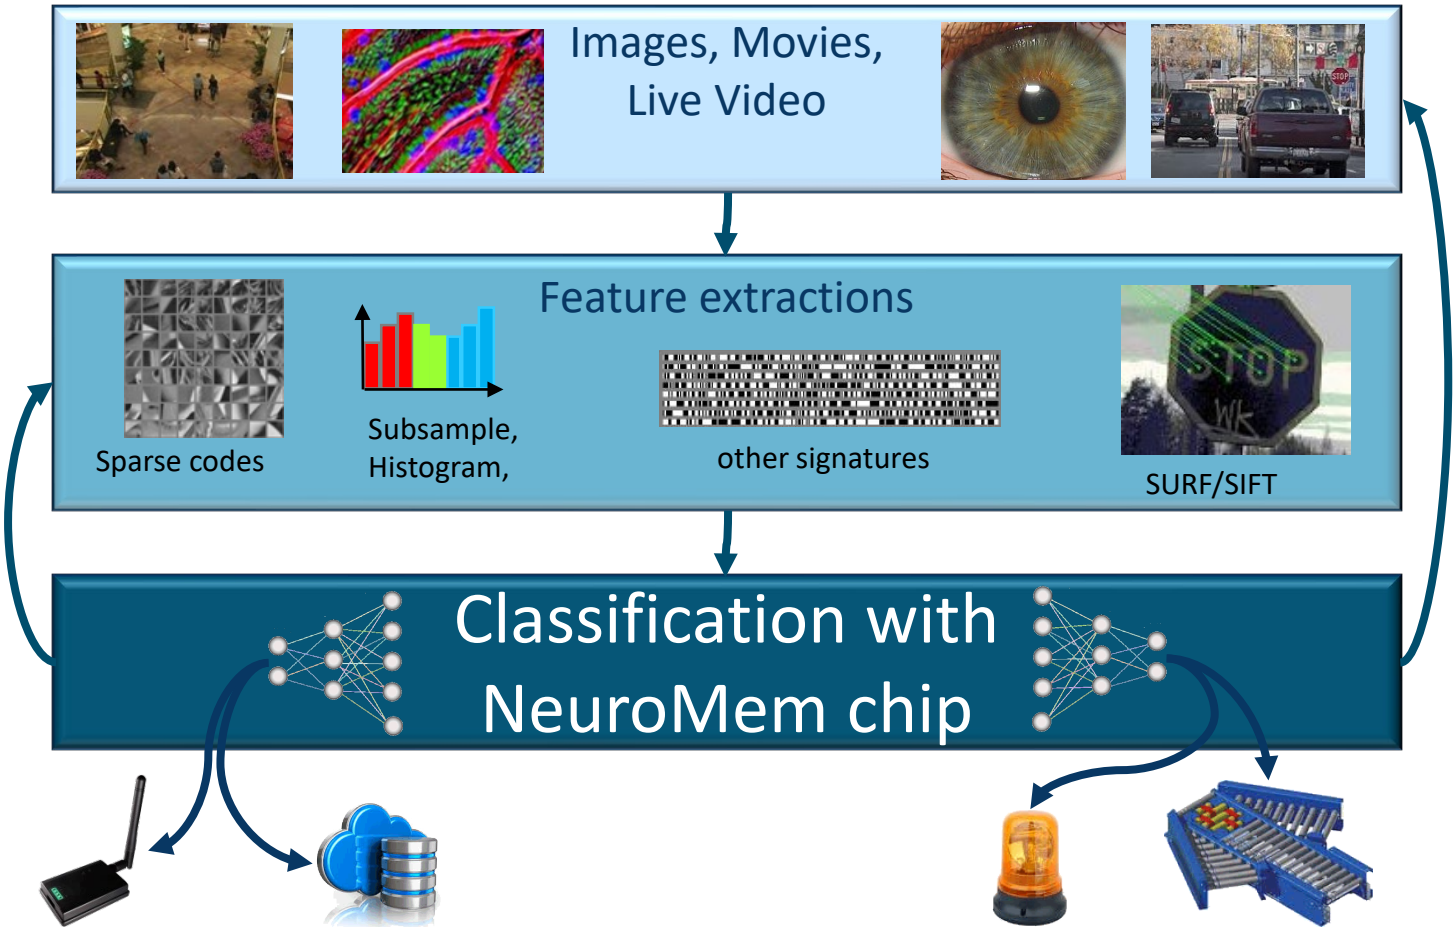

Application profiles

FUNCTIONAL GROUP

Discrete object classification

Object finding

Target tracking

Anomaly or Novelty detection

Texture recognition

Template matching

Object clustering

Compression

APPLICATION GROUP

Controlled environment

- Industrial automation
- Cooperative person identification

Environmental variations

- Outdoor lighting
- Changes of background

High cost of the mistake

- Requested throughput
- Requested accuracy

Image Learning

Teaching

- Supervised
- Semi-supervised
- Unsupervised

Features

- Shape
- Color
- Texture,
- Sparse codes

Learning Parameters

- Conservative, Moderate, Liberal
- Iterative, recursive

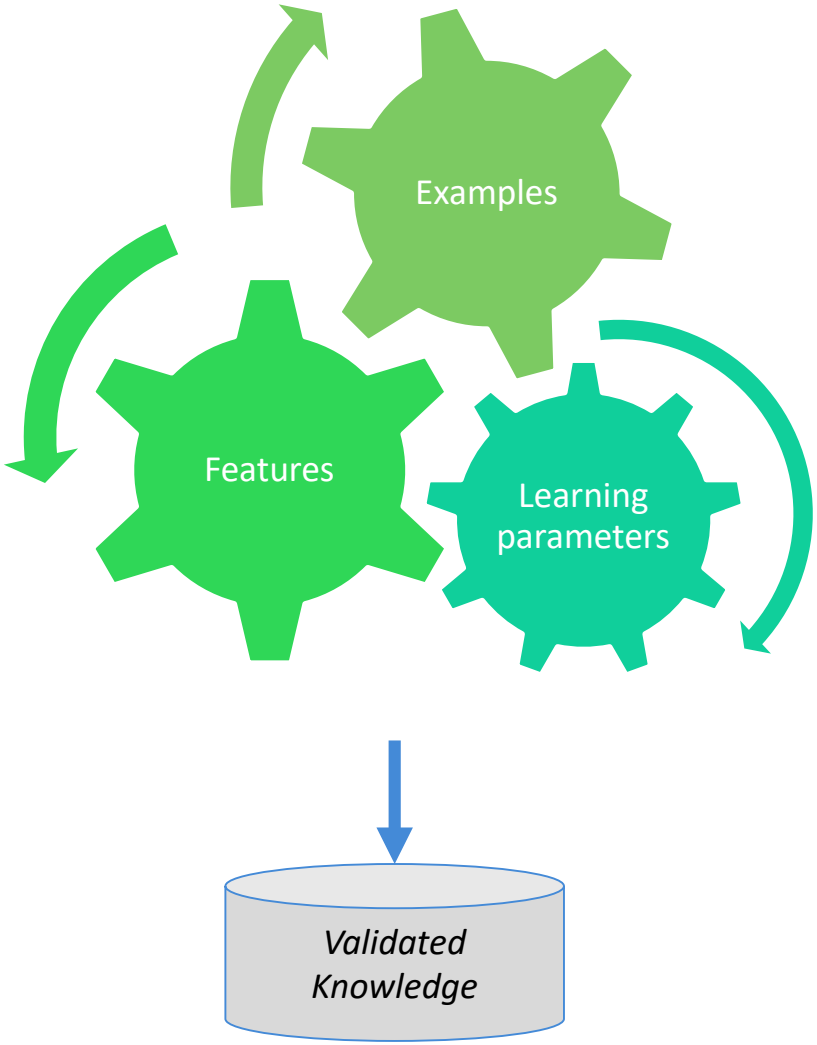

Learning Methods

SUPERVISED

Annotate finite and discrete objects under semi-controlled environment

Define of the smallest discrimination matrices including the objects

Define the best features discriminating the objects, possibly invariant to scale and orientation

UNSUPERVISED

Train a 1st layer of neurons to detect simple discriminant blocks of pixels in images such as edges, corner, patches of texture.

Train a 2nd layer of neurons to recognize the spatial distribution of these simple blocks describing larger shapes or objects.

How many neurons do I need?

Depends on the application, the variability of the data, etc.

Examples in image recognition…

Application	**Description**	**Estimated Neurons /Object**	**Total neurons**
Fish sorting	Classification of herrings (Accept, Reject, Recycle) passing on an in-line conveyor belt		≈200
Glass Inspection	Detection of anomalies of texture in patterned solar glass passing on a conveyor		≈800
Inkjet OCR	Reading of date code or serial numbers printed on a packaging	1-3 /digit	
Cooperative face recognition	Identifying a person facing front, positioned at a known distance of the camera, willing to remove her glasses if needed to be recognized	5/person	
Semantic analysis	Counting the occurrences of words from a dictionary in live tweets, posts and other text streams.	1/word	

Recognition

Where to look	• Fixed regions of search • Variable regions of search • Displacement and Steps
What to look for	• An object known by a given expert • A novelty
What to report	• Positions • Categories and confidence • Transform images
Decision rules	• Intra-expert consolidation • Inter-experts consolidation

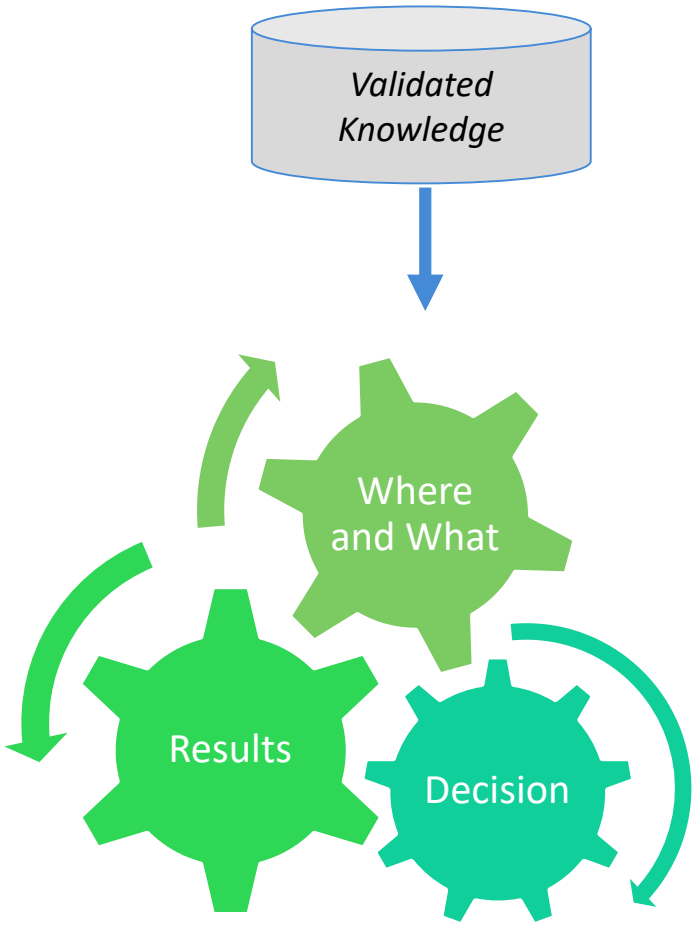

Recognition Methods

OBJECT LISTS

Produce the location of the recognized objects, along with their category and distance to the learned models.

The locations can be aggregated per category into a single point such as a centroid or into a bounding box.

Produce the location of all the anomalies or novelties found in the image.

TRANSFORM IMAGES

Map of the locations recognized by the neurons. The intensity can be representative of the category or distance of the neurons.

These maps can be used for visual interpretation, and also become input of a 2nd layer of network.

If the scanning step > 1, the maps are reduction of the source image.

Map of closest models: Replace each patch of pixels in the source image with the closest models stored in the neurons. A path for noise removal, image cleaning, smoothing, compression

Training multiple experts

An expert is trained with a specific feature
- Sensor related, Scale related, Feature related
- 1 expert = 1 context

Multiple experts can cooperate to deliver a global response
- Combinatorial rules, sequential rules

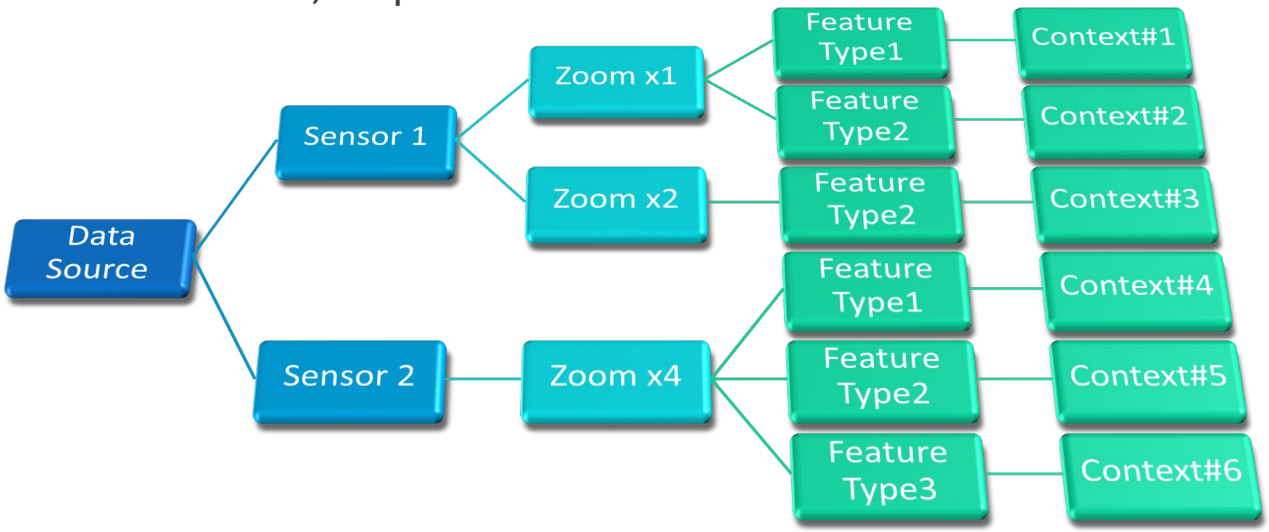

Consulting multiple experts

Sequential use

Locate license plate

Read license plate number

Match License plate number with owner

Parallel use

Detect new incoming person along edges

Track all detected persons

Multiple layers of experts

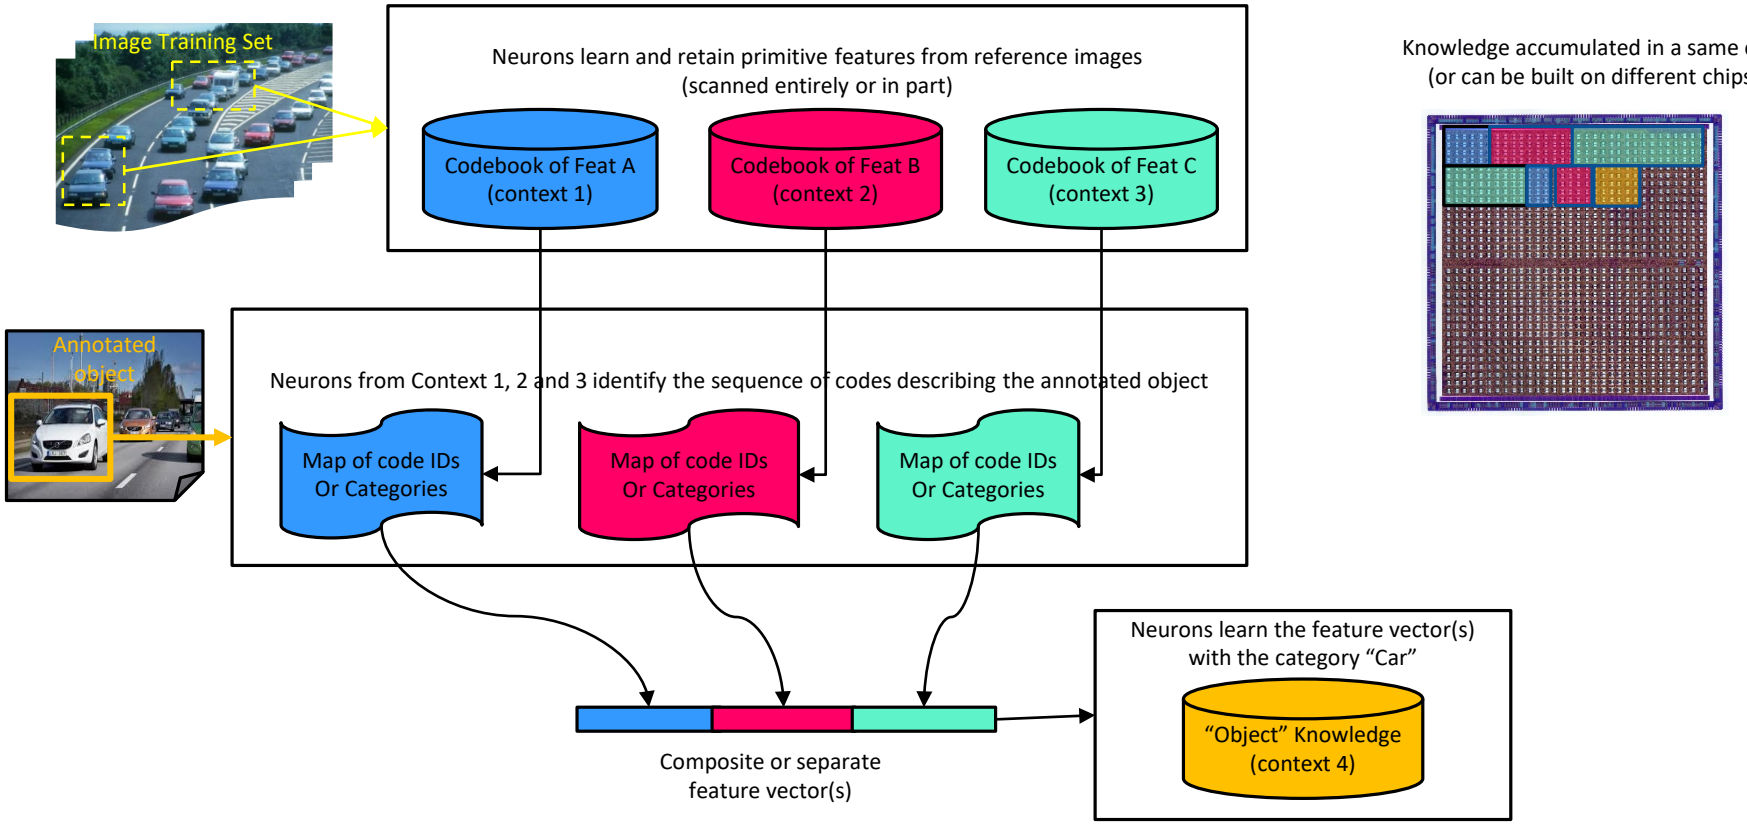

Classes of Knowledge Bases

Always-on or initial knowledge

- Used at the power-up and running continuously
- Intended to spare the battery and awake the device only when necessary.
- Triggers the use of the next contextual knowledge KN1
- Examples: Detection of a pair of eyes looking at the screen, a specific motion, etc.

Contextual knowledge

- Knowledge specific for an application
- Knowledge built and validated ahead of time
- Loaded from a knowledge file
- Examples: Recognizing a monument in a city, a product in a storefront

User knowledge

- Knowledge built on the device by the user
  - Taught by showing new examples
  - Expand the contextual knowledge, if any
- Examples: A philatelist can teach a new stamp to add to his collection