The invention of the network camera

and the VLSI technology behind

Stefan Lundberg Axis

Oct 2018

Axis Mission

Axis

- Founded in 1984 at IDEON in Lund
- IT company focused on Network Video Solutions
- 3054 employees (Q2, 2018)
- Sales 2017, 8600 MSEK
- Listed on NASDAQ OMX under acquisition by Canon
- Head office in Lund, close to LTH
- Own presence in more than 50 countries
- 90000 partners in 179 countries

Original products

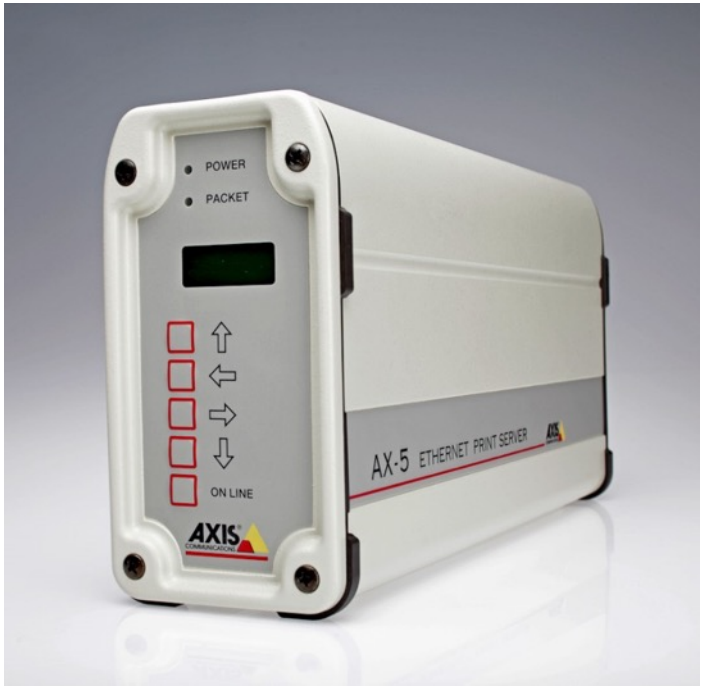

The market's broadest portfolio

Fixed cameras

Thermal cameras

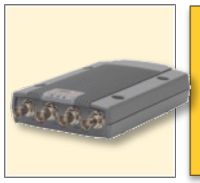

Encoders/ decoders

Fixed dome cameras

Onboard cameras

Software & recording

PTZ cameras

Accessories

Physical access control

AXIS Q3709-PVE Network Camera

Multi-sensor, multi-megapixel - 180° overview. One camera.

> 180° panoramic **overview**

> Smooth video of movements at up to **30 fps in 3 x 4K**

> Efficient one-camera **installation**

Small business – AXIS Camera Companion

Video surveillance made simple

- Optimized solution for 1-4 camera systems, with support for up to 16 cameras
- Recording on SD cards, removes the need for central recording
- Mobile app for freedom of use
  - Axis Secure Remote Access for a secure and easy-to-install way to access the system remotely
  - Axis Mobile Streaming for efficient bandwidth usage
- VMS and mobile viewing app **free** for download

AXIS A8004-VE Network Door Controller

Reliable audio visual identification and entry control

> A perfect complement to any surveillance installation to effectively control entry:
  - Identify and talk to visitors
  - Record what happens at the entrances

> Suitable for small- and mid-size installations as well as advanced enterprise systems


AXIS
COMMUNICATIONS
ARTPEC-6
1618076
1638LP900

AXIS Q1659

Professional photography meets video surveillance

AXIS Q1659 Network Camera offers **ultra-high image resolution** for surveillance in superb detail at 8 frames per second, achieving **unprecedented levels of detail** for observing open spaces and across long distances.

It features digital single-lens reflex (DSLR) imaging technology and offers a choice of **Canon EF/EF-S lenses**, depending on individual user needs. Equipped with an EF lens mount, the camera enables easy lens changes.

AXIS Q1659 is compatible with a large number of **video management systems**, and it supports **Zipstream** that significantly reduces bandwidth and storage requirements.

> 20 MP resolution
> Canon EF mount and EF/EF-S lenses
> Axis Zipstream technology
> PoE and SFP slot for optional fiber connection
> Optional accessory housing; AXIS T93C10 Outdoor Housing

AXIS D2050-VE Network Radar Detector

Why a radar? What's the problem

- Many customers are searching for a solution that can minimize repeatedly false detections.
- Video Motion Detection (VMD) can be challenging to use outdoor.
  - During night time and
  - Bad weather conditions
  - Spiders and bugs are present.
- Radar can serve as a complement to video surveillance
  - Suitable for area protection and makes it possible to positively detect movement with high accuracy

AXIS D2050-VE

- 24GHz FMCW Radar
  - Frequency modulated continous wave radar
- Range 50 meters / 164 feet
  - Coverage 3000 m2 / 3600 sq yd
- 120 degree horizontal detection coverage (tbd)
- 3,5 m mounting height
- PoE class 4 (IEEE 802.11at)
- Minimizes false detections
  - Rain, snow, insects, shadows, night

PTZ live autopilot

> Radar data can control Pan-Tilt-Zoom cameras
  - Enables PTZ camera to be pre-positioned when operator take control
  - Follows object during day, night or bad weather conditions

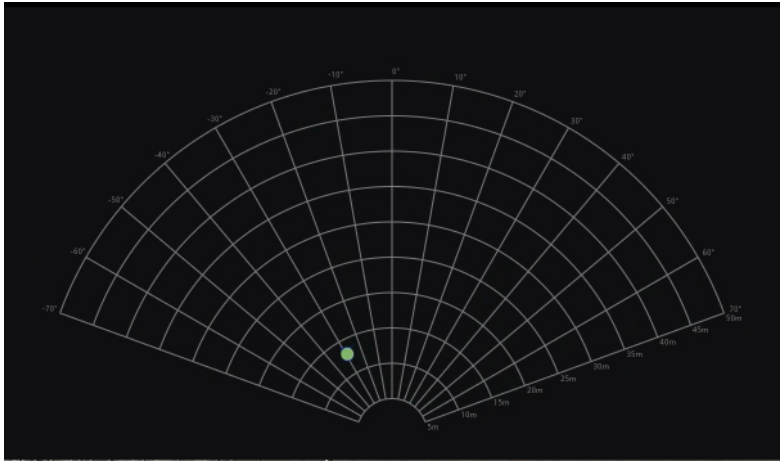

30+ years of intelligent networks

Network audio systems

Physical access control

Network video

Storage

Network print servers

IBM protocol converters

2012

1984 1992 1996 1998 2001 2011 2013 2015

Europe

Asia

USA

Global

Our business model

Why is the sales model so important ?

- Everyone can trust Axis
- Integrators never compete with their main partner
  - But they do compete with each other
  - Axis never “steals from the channel”
- Distributors always make money
  - Axis often has no set-up to bypass
- Rebate system to ensure correct pricing
  - Deal pricing, partner pricing, etc.
- Make sure this model fits
  - Analytics, use partners
  - ACS – use for low end, not large systems

Regional development Q2, 2018

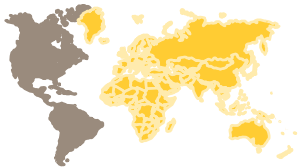

Americas

57%

of sales

> Total sales of SEK 1,480 M (1,180)

> Growth of 25%

> Local growth of 31%

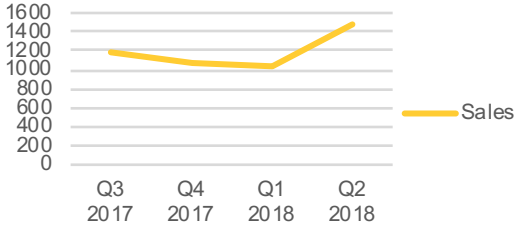

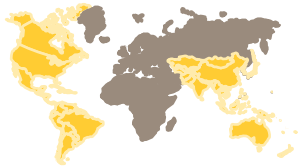

EMEA

32%

of sales

> Total sales of SEK 831 M (696)

> Growth of 20%

> Local growth of 14%

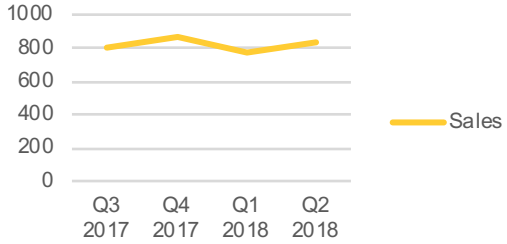

Asia

11%

of sales

> Total sales of SEK 290 M (244)

> Growth of 19%

> Local growth of 22%

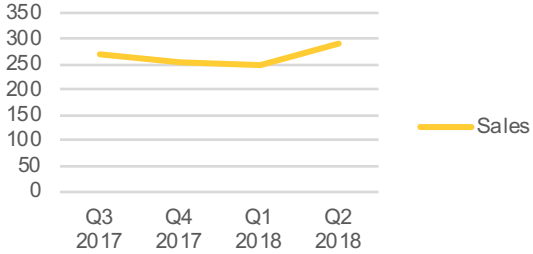

Axis annual sales 1996-2017

Worldwide presence – close to the customers

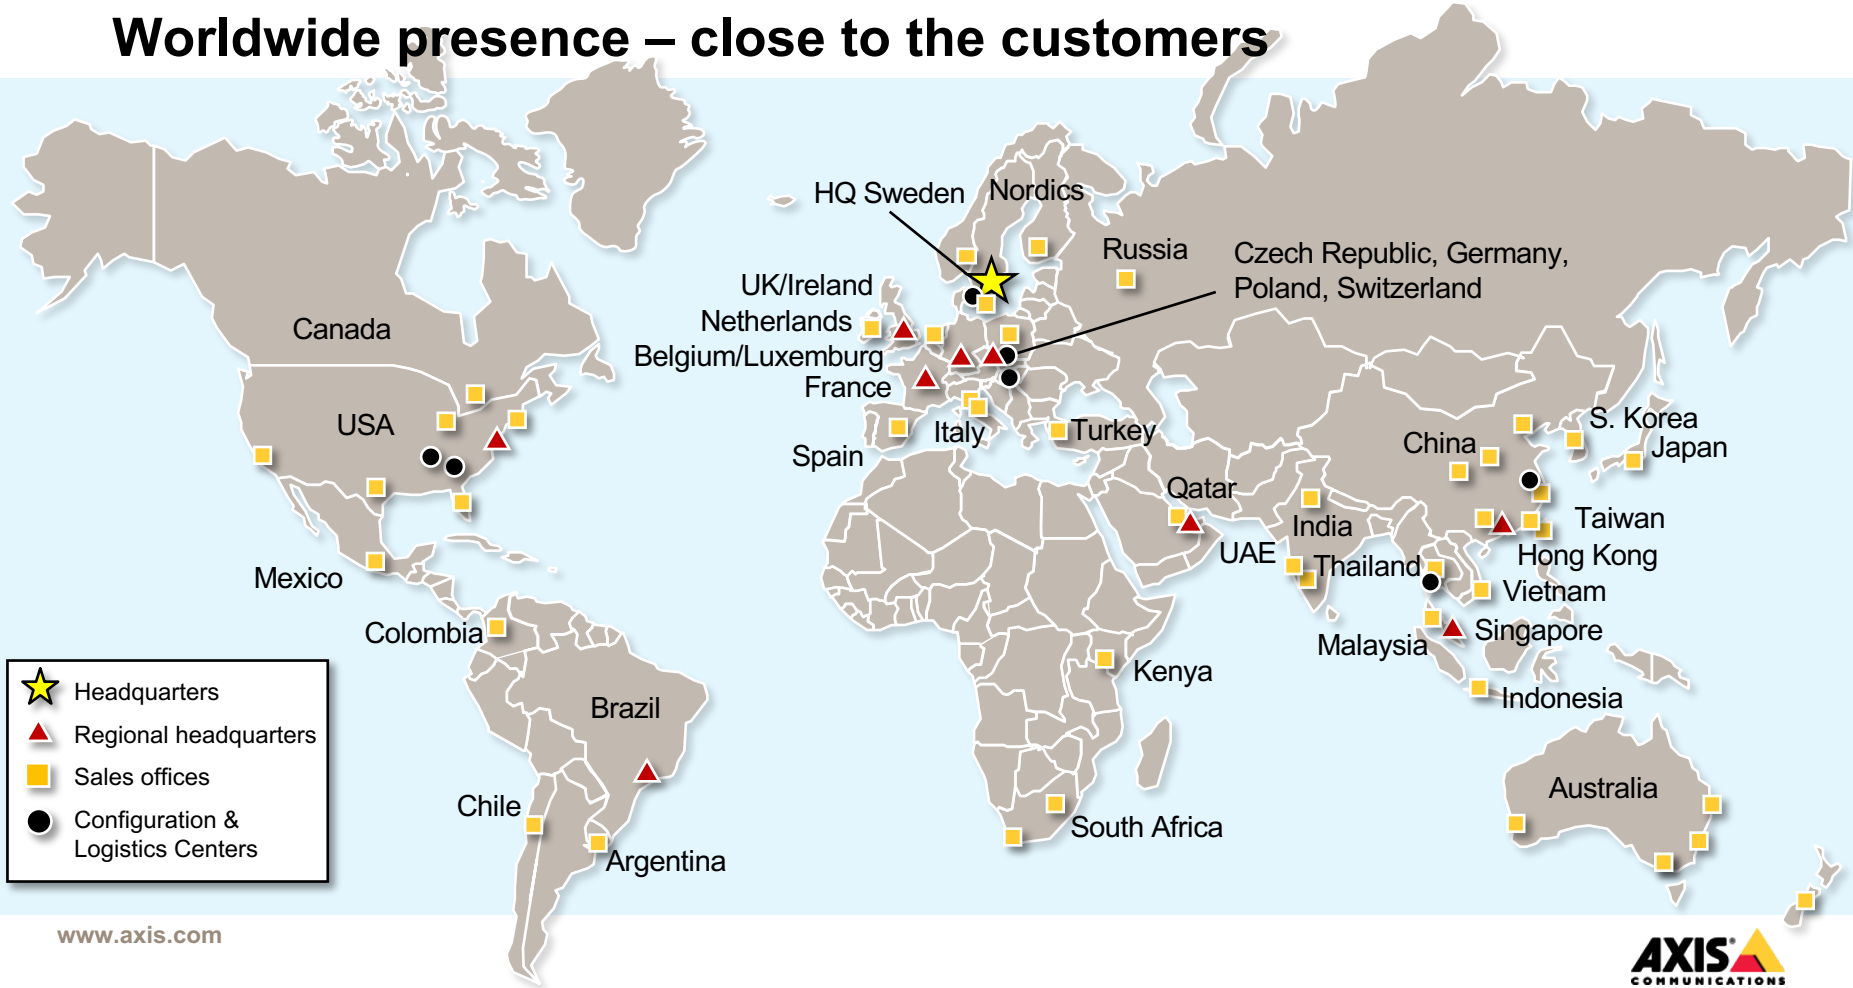

Taking long-term responsibility by thinking big

Sustainability highlights

> Continued to enhance and reduce product packaging

> Optimized logistics

> Energy-efficient, more environmentally friendly products

> Anti-Corruption policy

Axis – continuously driving innovation

1996

World's first network camera

1998

World's first video encoder

1999

World's first network video chip

2004

First MPEG-4 and Motion JPEG compression camera

2008

First H.264 compression standard for network camera

2009

First network cameras with HDTV, and with remote focus & zoom functions

HDTV NETWORK VIDEO

SMPTE

2010

First thermal network camera

2011

Lightfinder technology

2012

Unique high-performance WDR camera

2012

First network camera with active cooling

2013

Physical Access Control

2015

Zipstream technology & Sharpdome technology

2015

Open standard network loudspeaker & Open IP-based door station

2016

Pan, Tilt, Roll, Zoom (PTRZ) technology & laser focus technology

2017

Network radar technology

AXIS COMMUNICATIONS

20 years of network cameras 1996-2016

Axis invented the world's first network camera

Technology leader with many industry firsts

Continuing to drive innovation for a smarter, safer world

The fourth phase - the invention of the network camera

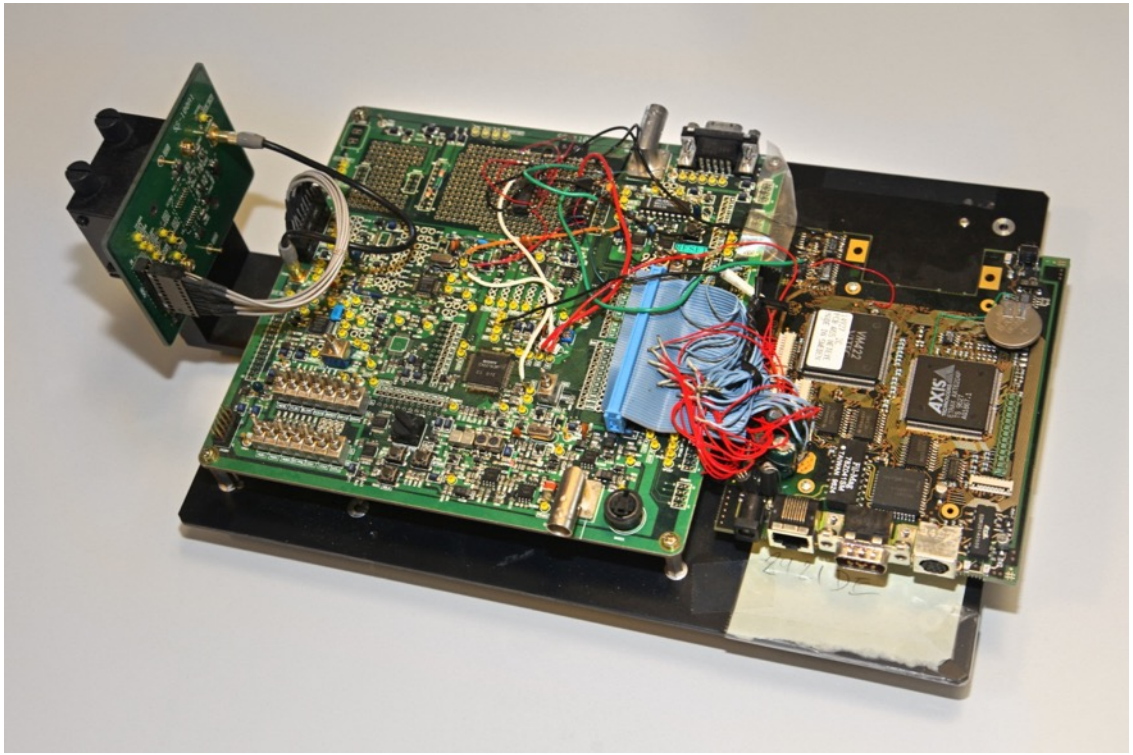

1996 - AXIS NetEye 200

- The world's first network camera
- Launched at Interop Atlanta, September 18th, 1996
- Performance
  - 1 image/second in 352*288 pixel
  - 3 frames/minute in 0.4 Mpixel

More than 10,000 sold!

One of our first camera customers…

> Steve Wozniak, co-founder of Apple

> Was in a car accident during a tech support call

> No injuries and problem was solved

World's first video encoder in 1998

> AXIS 240 – the world's first video encoder

> Andy Rubin, CTO & founder of Android at Google, in 1997 when he was testing AXIS 240
prototypes

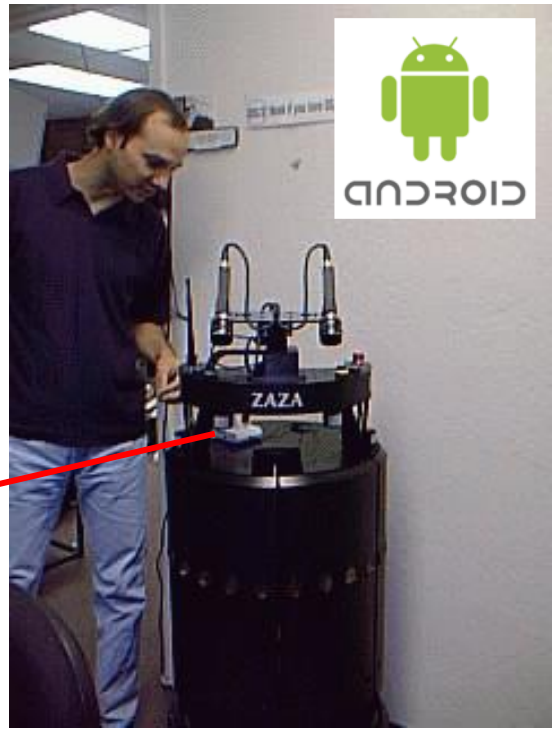

Lightfinder technology, Wall Street Journal Technology Innovation Award

AXIS 221

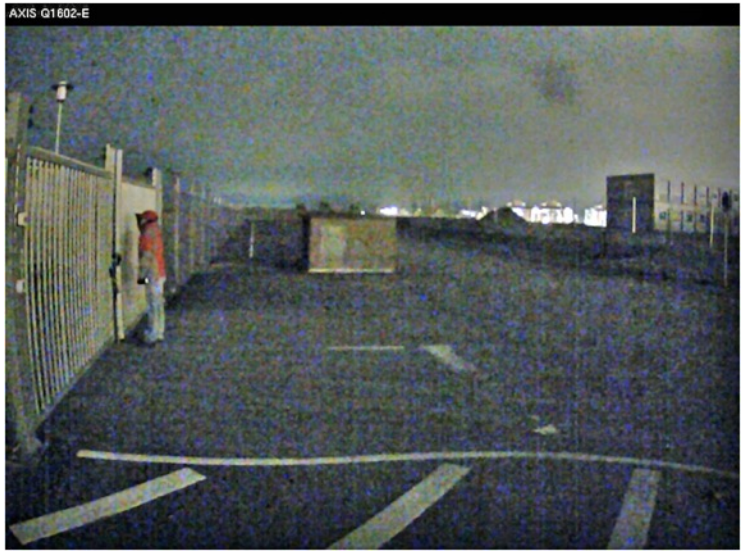

AXIS Q1602

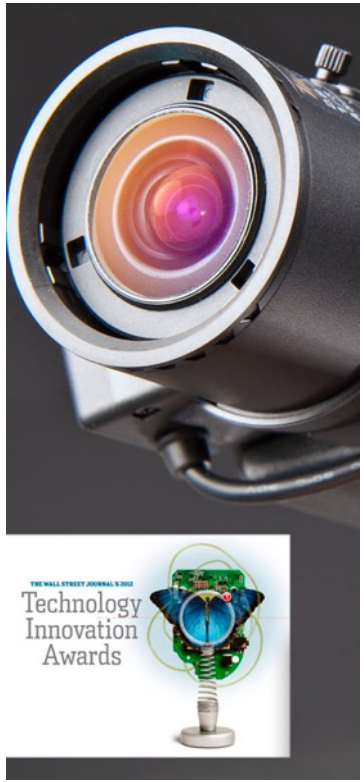

Perimeter surveillance with no street illumination during night time and approximately 0.1 lux.

In-house R&D – where the coming successes are made

- Substantial R&D investment
  - 15% of revenue
- Strong network camera patent portfolio
- Canon is very strong on patents
- 3 corner stones
  - Innovation
  - Openness
  - Quality

More than 800 engineers work in Lund

Mechanics

Software

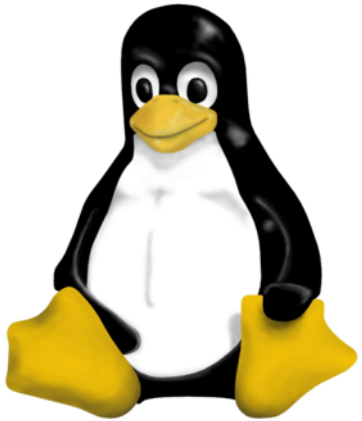

Electronics

ASIC´s

”ASIC” stands for ”Application Specific Integrated Circuit”.

ASIC development at AXIS

Axis – continuously driving innovation in network video

1999
ARTPEC-1
Worlds first Network camera ASIC

1996
Off the shelf Components (VITEC)

2003
ARTPEC-A

2004

ARTPEC-2
Dualstream MPEG-4/MJPEG

2008

ARTPEC-3
HDTV/H.264

2009
ARTPEC-B

2011

ARTPEC-4
Lightfinder technology

2012
Ambarella
(ARTPEC-C)

2014

ARTPEC-5
Forensic capture
Zipstream

2017

ARTPEC-6
Forensic WDR
Radar

AXIS ASIC history

> AXIS has developed SoC ASICs since late 80ths

> 80s ASICs
  - where used for network protocol converting

> 90s ASICs
  - In-house developed RISC CPU (ARM-style) where added and the proprietary Canon printer protocol Page 21 where added forming the ETRAX chip family.
  - A separate image processing ASIC ARTPEC was developed for the camera business.

> During 2000 to 2010 the ASICs became large unified SoC where more and more external IPs where.

> Today AXIS is focusing on development of differentiating functions such as image processing, analytics, scaling, overlay and the overall specification of the SoC.

AXIS ASIC history

- From start all functionality was developed in-house by a few designers
- There where no market for general building blocks (IP blocks) such as CPUs, Memory Ctrl, Interfaces etc but this has changed over the years,
- Today we need about 60 man years and still half of the functionality is developed by external parties.
- Number of transistors has multiplied by 2500 since then.
  - Moores Law: Number of transistors doubles every 18th months
  - AXIS has historically been one to two step behind latest technology
  - But are now closing in to be able to fulfill requirements

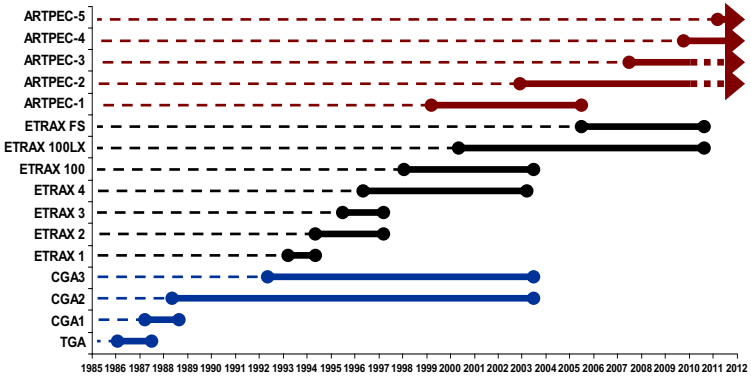

Typical Axis SoC content

- > Embedded CPU running Linux
- > Image processing pipeline
- > Image scaler with dewarping
- > Image compression subsystem
- > Crypto accelerator
- > Ethernet controller
- > I/O controller
- > Interfaces etc

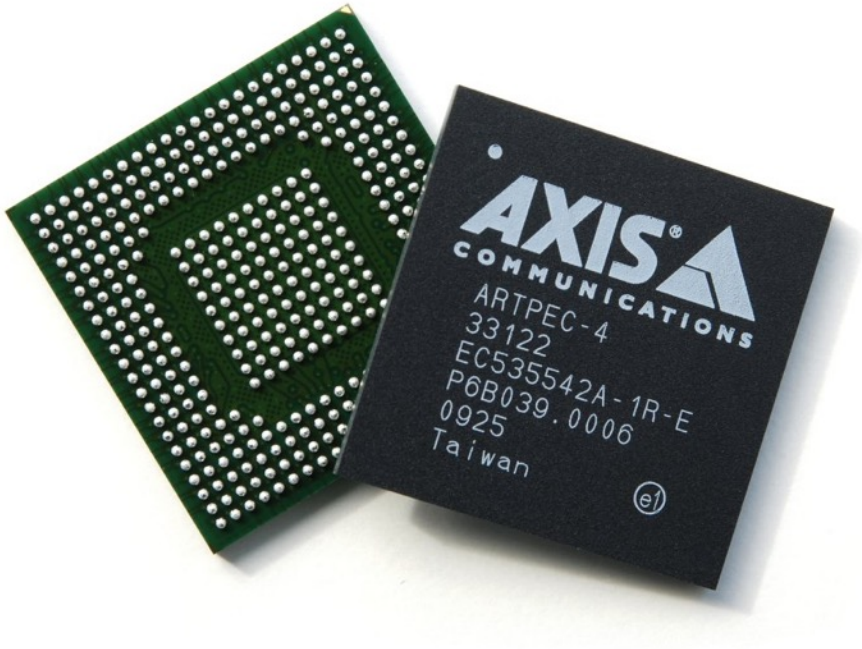

An ASIC often combines the different functions needed for a complete system on one single chip. This is called ”System on a Chip” or ”SoC”.

ARTPEC 1 - 6

> Six generations of dedicated network camera ASIC’s

ARTPEC-1

- Worlds first network camera ASIC
- uCLinux on ETRAX CPU
- CCD/CMOS IPP, MJPEG compression
- Latch based design, 2-phase clocking
  - Clock gating to save power
- 50MHz
- 32kbyte SRAM
- 160 pin PQFP

ARTPEC-1 product

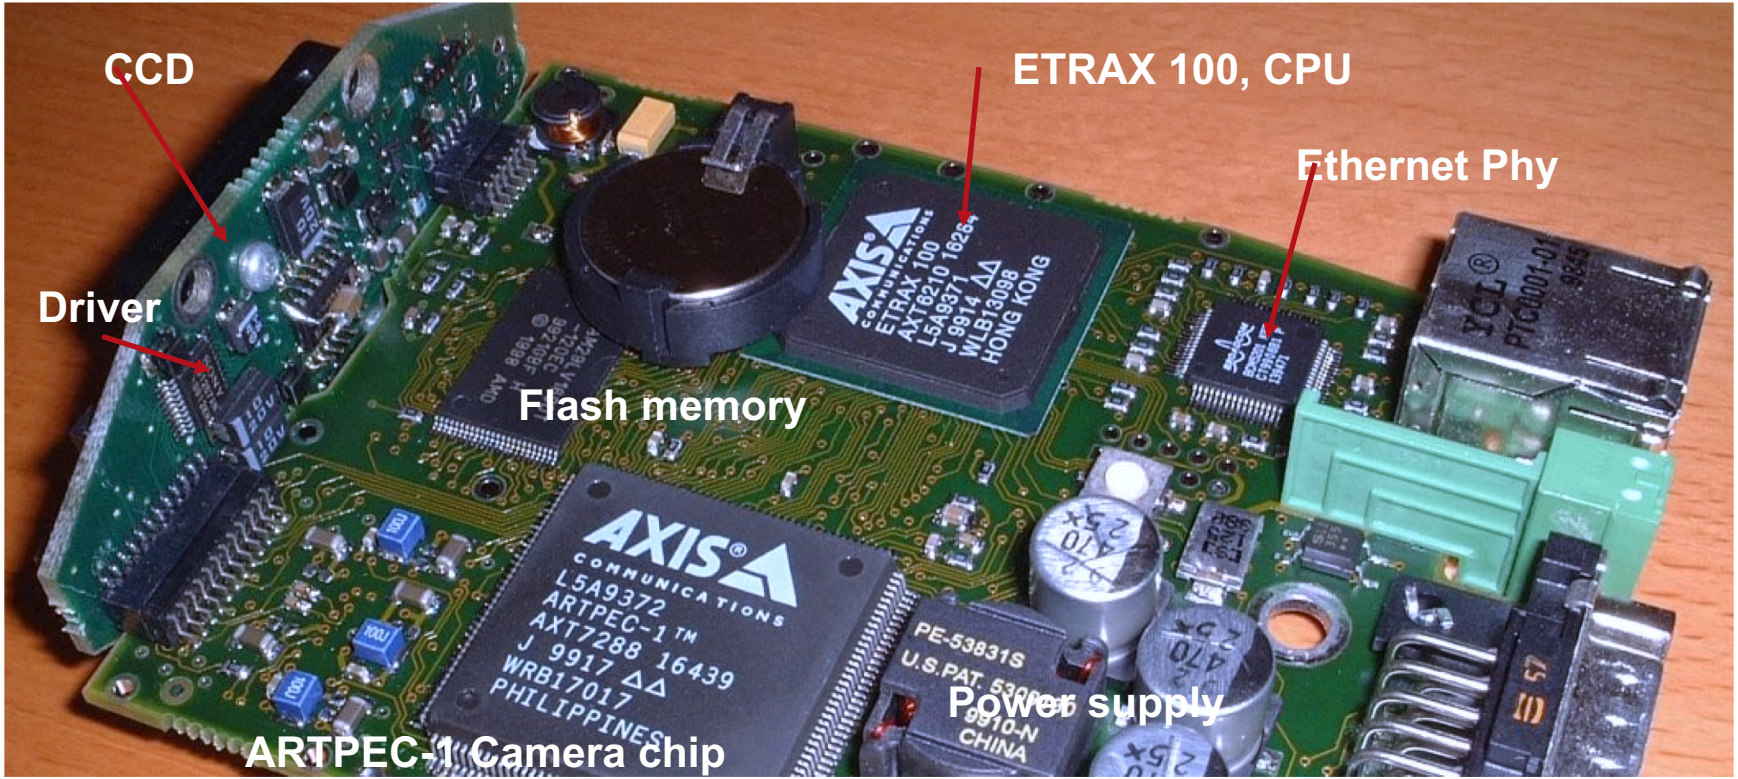

ARTPEC-2

- Worlds first dualstream MPEG-4/MJPEG network camera ASIC
- CCD/CMOS IPP 45MPixel/s
- Latch based design, 2-phase clocking
  - Clock gating to save power
- 100MHz
- 40kbyte SRAM
- 208 pin PQFP

ARTPEC-3

- > Worlds first HDTV/H.264 network camera SoC
- > HDTV 1080p30
- > 1-phase clock
- > 200MHz
- > 16kB + 16kB Cache
- > LFBGA400 17x17mm

ARTPEC-4

- > Lowlight/WDR network camera SoC
- > HDTV 1080p30
- > MIPS
- > 400 ball TFBGA

ARTPEC-5

- Faster general purpose CPU
  - Dual core
    - Independent execution units
    - Dual hw threads per core
    - Parallel computing architecture
- Higher memory throughput
- Faster Video Analytics
- New camera features
  - Improved image quality
  - WDR
- Zipstream

ARTPEC-6

- High quality 4k30 products
- Faster general purpose CPU
  - ARM
- Higher memory throughput
- Faster Video Analytics
  - GPU
- New camera features
  - Improved image quality
  - Forensic WDR
- HDMI

Electronic design obstacles, history and future

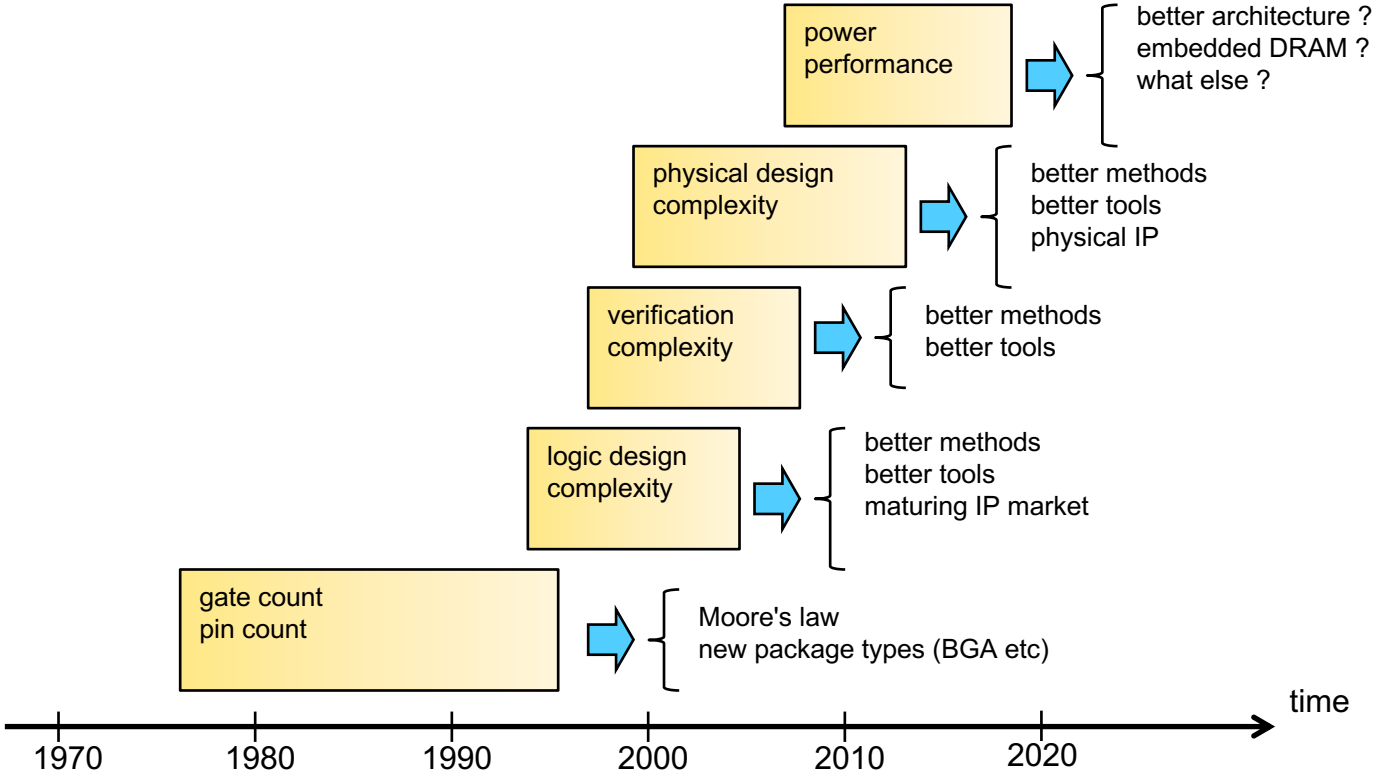

Why ASIC development?

> Benefits:
  - Product performance
  - Product size and power (Size 70 to 100 times less than FPGA in same technology)
  - Unique features (hard to copy by competitors)
  - Secured access to technology (still there are not many external ”Platforms” for video market)
  - Unit cost reduction

> Challenges:
  - Time to market
  - Development costs
  - Complexity
  - Project risk (Size/Time/Spec)
  - Technology risk

Description of AXIS ASIC development

- We use partners/ASIC vendors for the back-end design and manufacturing of our ARTPEC chips. We have used both partners that has their own fab and fabless companies.
- Design and verification (VCS) using SystemVerilog, main reason to SystemVerilog is verification but we are also using advanced structures that makes design work more effective.
- We have done VMM based random verification since 2005 and are now moving to UVM. Modelling is TLM2.0 and C based
- We do Verilog netlist handoff with timing constraints to our partner. We do both synthesis (DC) and floorplan (DC Graphical) to ensure quality and decrease number of iterations

Chip development

- The written code will be translated by running a synthesis tool to a specific cell of the selected library of a specific process (e.g 28nm TSMC Low power library)
- When all functionality is described in System Verilog and synthesized we call it a Netlist.
- We can simulate functionality, performance (speed), power (dynamic and static) and do an approximate placement and routing (connection).

```
module ff (q, d, clk)
output q; input d, clk;
reg q;
always @(posedge clk) q = d;
endmodule
```

EDA tools

- Tool supplier
  - We have been using Synopsys tools for a long time.
  - The license model is Time Based License with pool of tools
  - Mix changed based on project flow
- Design & Implementation - SystemVerilog
- Synthesis – Design Compiler
- Verification - VCS
- Floor planning – DC Graphical
- STA – Primetime
- DFT – we do not do DFT but we prepare our design for DFT
- Using other point tools for specific design purposes
  - Power - Power Compiler
  - High-level modulation - QUEMU, TLM2.0
- Continually evaluation tools and vendors

How to handle risk and development cost?

- Luckily specification of mobile devices is more similar to Surveillance Camera than old Mobile phone i.e. always ”On”, which gives us benefit from using same IP
- Increased design complexity must be managed
  - Improved design tools and design methods
  - Modularization and reuse
  - Make use of External competence for Top Level integration and General System implementation
  - Make use of external chips for parts of our product portfolio that has specific needs
- NRE (external design and manufacturing cost) cost increase dramatically with newer technologies
  - Carefully specify new designs to reach enough volumes
- Make use of external chips in some products to ensure second source

Two teams

- **Chip Platforms**
  - Manager Lars Branzen
  - Product Manager
  - Project Manager
- Competence Groups
  - Backend group
  - System Architecture group
  - FPGA

- **Chip Platforms IP**
  - Manager
  - Architect
  - Project Manager
- Competence Groups
  - Front End Design Group
  - Verification Group

Responsibilities

- ASIC and FPGA methodology for System-On-Chip
- Models for early SW development (TLM/System-C)
- Suggest chip platform technology solutions
- Chip platform roadmap
- Chip product ownership
  - User documentation
  - HW support during lifetime
- ASIC and IP purchasing
- FPGA development for early SW development, algorithm validation and product functionality extensions

Chip Project

Chip development

- > Axis develop critical functionality
- > ASIC Vendor integrate CPU subsystem
- > The netlist is then assembled by the ASIC Vendor who will add functionality for test and manufacturing followed by exact placement and routing.
- > After iterations between Axis and Vendor (changing floorplan, RTL, specification etc.) the final version of the circuit will be written out in GDSII format which will be used for producing production Masks.
- > Prototype production
- > Engineering Samples sent to to Axis

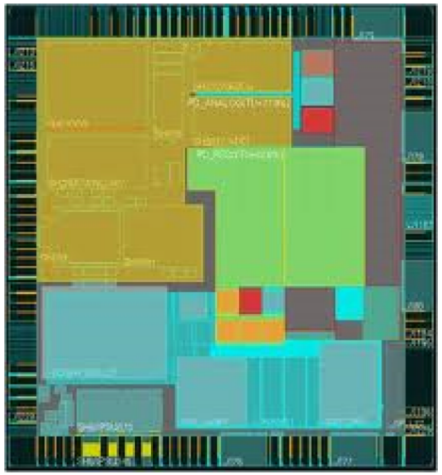

Chip development as a Program

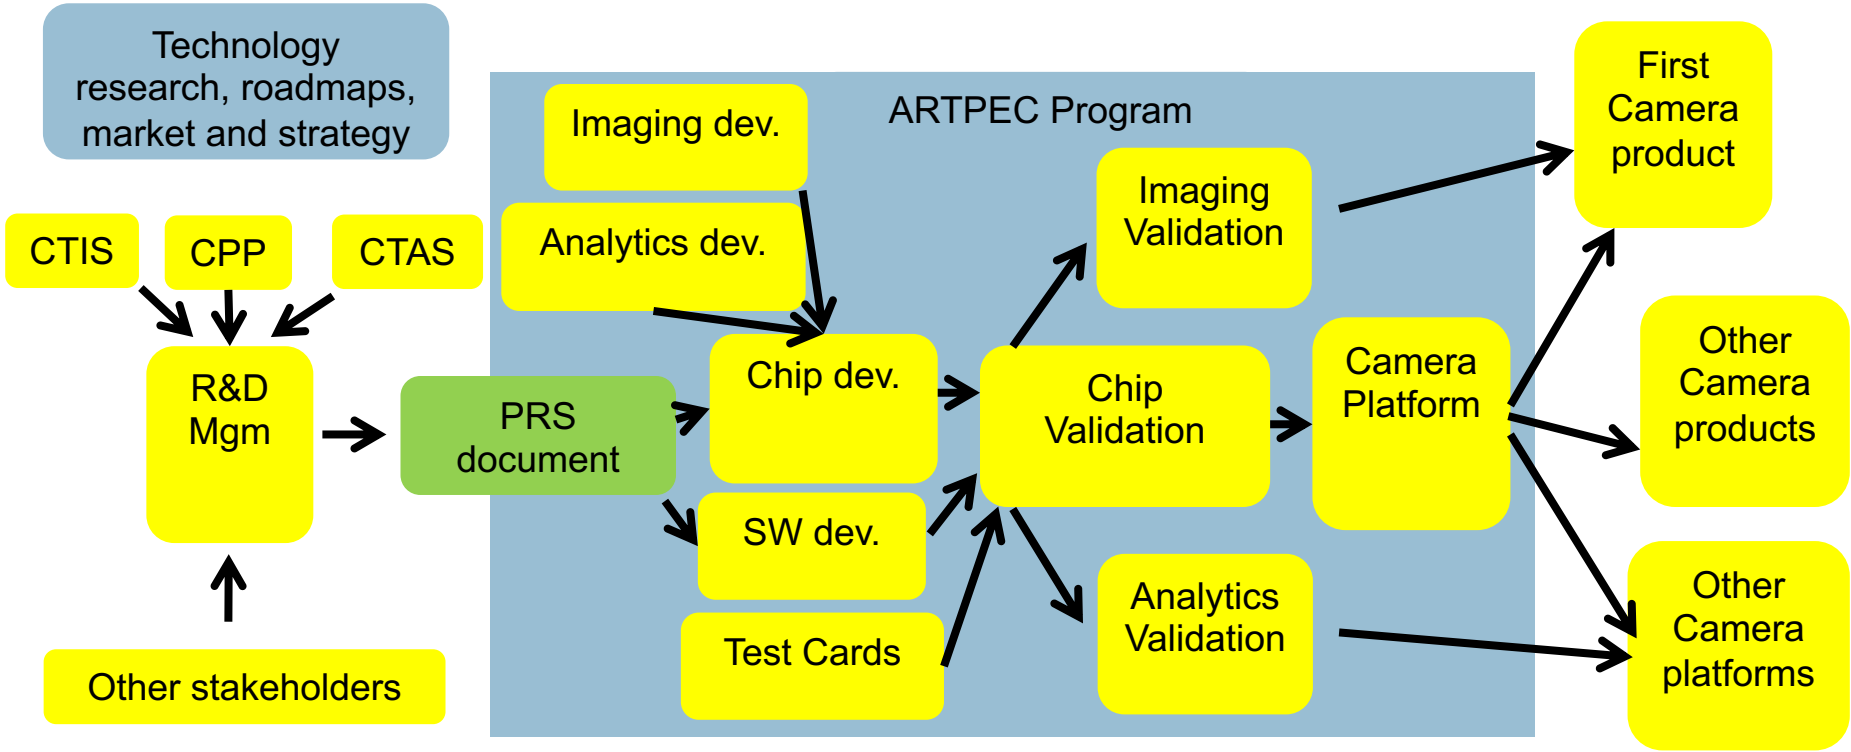

Chip IP project

A parallel project develop new specific functions (IP)

> Driven by CP-IP team
  - ~15 ASIC design engineers + 5-10 consultants

> Assisted by
  - Core Technologies Imaging Systems (CTIS)
    - Algorithms and technologies
  - Core Technologies Analytics&System (CTAS)
    - Algorithms and technologies
  - Core Product Platforms team (CPP)
    - HW/SW interface reviews
    - Linux driver development
    - Prototype/sample validation
  - Core Technologies – Media and Graphics (CTMG)
    - Image Coding, Audio functions and algorithms, GPU, Scaling and Overlay requirements

Chip project

Vendor selection and integration

> Driven by our Chip Platforms team
  - ~7 ASIC design engineers

> Assisted by:
  - Tech-Ref and PCB-CAD
    - Electronics
      - Package (ball-out)
      - PCB layout (incl. X-talk and SI-analysis)
    - Mechanics
      - Thermal design and analysis

Chip project

- Modular design flow
  - Design implemented in the RTL (Register Transfer Level) language System Verilog
    - Parallel processes
    - Synchronous design style
  - Fully verified sub-designs, comprehensive random testing (UVM/VMM)
  - Synthesis, STA and DFT clean sub-designs
  - Design guidelines and checklists
  - Documentation and reviews
  - Predictable integration in ASIC project

Chip Project - Typical ASIC project organization

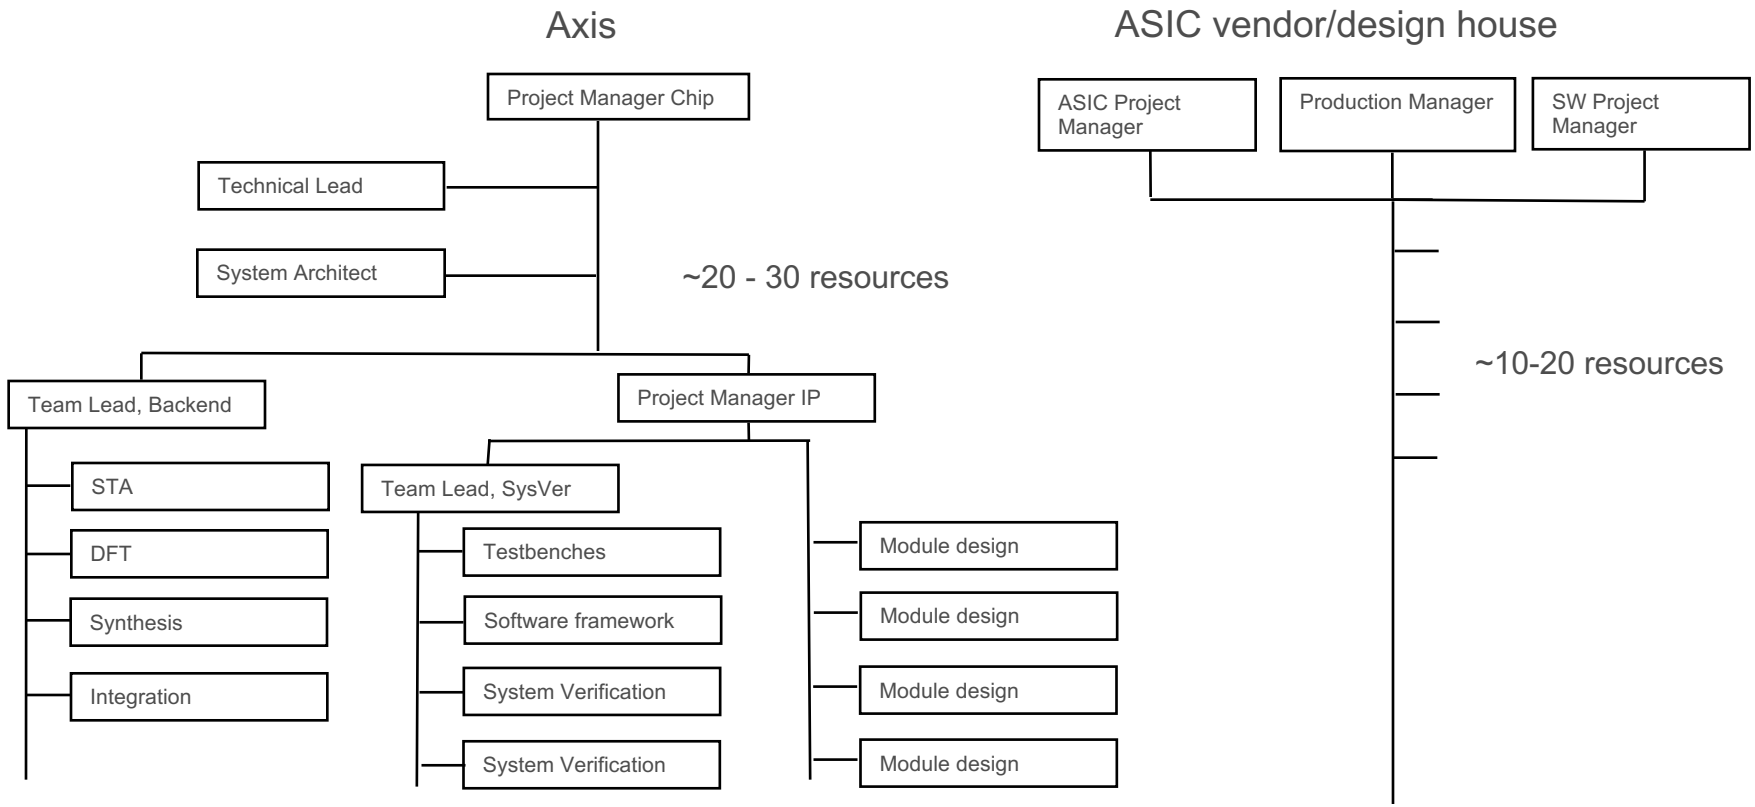

Chip project - Cooperation with ASIC vendor/design house

- Generally 3 main phases in project
  - 0.5 – initial netlist
  - 0.9 – trial netlist
  - 1.0 – final netlist
- Traditional ASIC design flow for ARTPEC
  - Verilog netlist and SDC (Static timing Design Constraints) handover
- Joint work on
  - Package / pinout
  - Power simulations
  - IP integration
  - DFT
  - Floorplanning
  - Timing closure

18-month ASIC design cycle

Project A ending phase

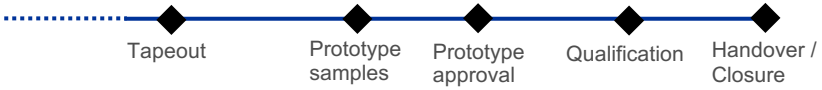

Project B main phase

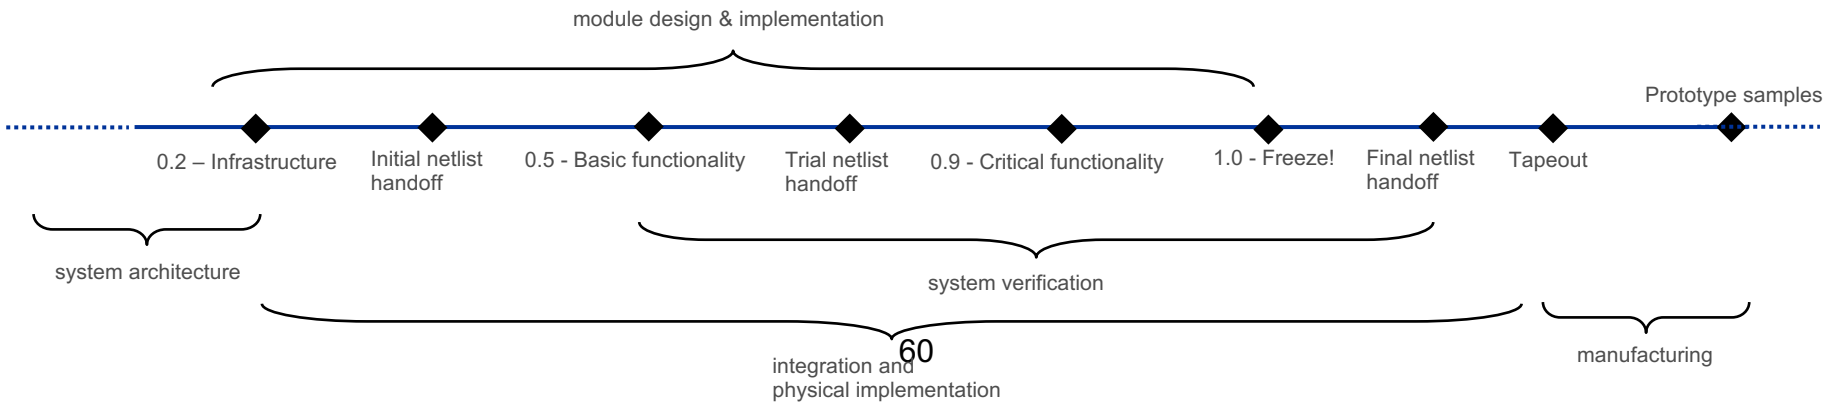

Project C starting phase

Deep learning

Definitions

Arthur Samuel defined machine learning as a "Field of study that gives computers the ability to learn without being explicitly programmed" 1959.

Deep Learning is a class of machine learning algoritms that use a cascade of neural-inspired non-linear mathematic for feature extraction and transformations. The term was introduced 1986 by Rina Dechter, but modern implementation much later.

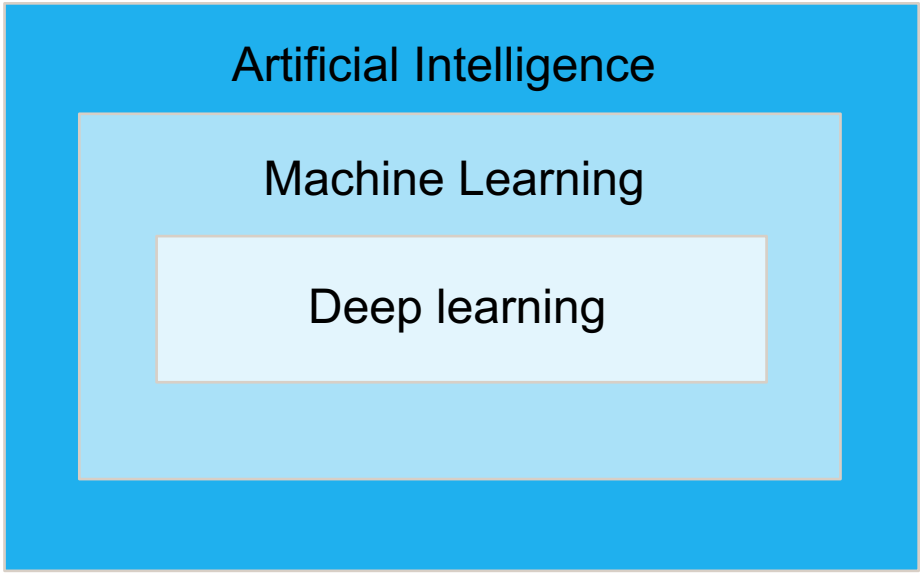

Machine learning

Picture -> algorithm -> “It’s a lamp!”

Variation is the problem when doing classification

- Viewport, Scale, Deformation, Occlusion, Illumination, Background clutter, Intra-class differences

Not easy to program with classic rule based methods

Deep learning

Inspired by how the human brain learns to see.

Hand-crafted features

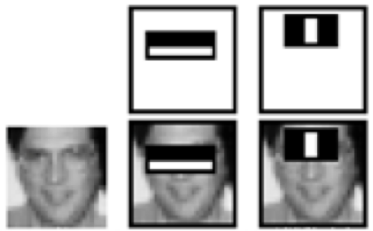

Data-driven features

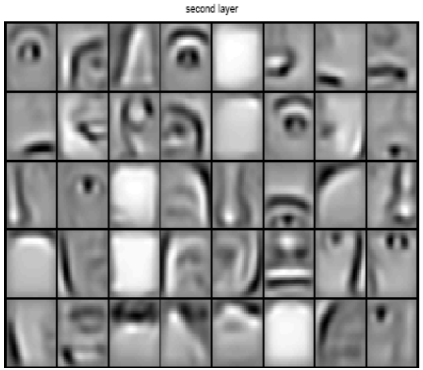

Image classification, data driven approach

> Problem: Based on **N** measures, sort into **K** number of classes

CARS

CLOCKS

HOUSES

Classifier types

- Parametric classifier
  - Calculations generate some sort of equation to be used later. Initial data can be discarded.
  - Fast to run on new data when the classifier is used.
- Non-parametric classifiers
  - Quickly the algorithm collects large amount of data that need to be stored for later use.
  - Large amount of data need to be processed for new data when the classifier is used.

Algorithm examples:

Logistic Regression
Linear Discriminant Analysis
Perceptron
Naive Bayes
Simple Neural Networks

k-Nearest Neighbors
Decision trees like CART
Support Vector Machines

Binary classification

- Each sample belongs to one of two classes
- This example is 2-Dimmensional, eg two features (x,y) are measured.

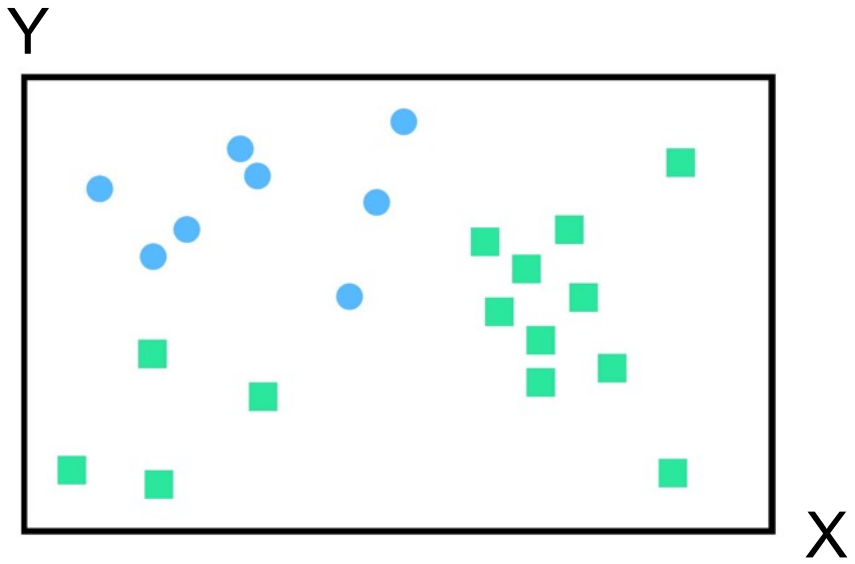

Multi-class classification

> Clustering is a way to group similar samples.
  - Unsupervised learning (No labels needed)
> How do we know how many categories there are?
> How to decide the value of hyperparameter k?
  - In many cases we need to select and test.
  - k-means clustering

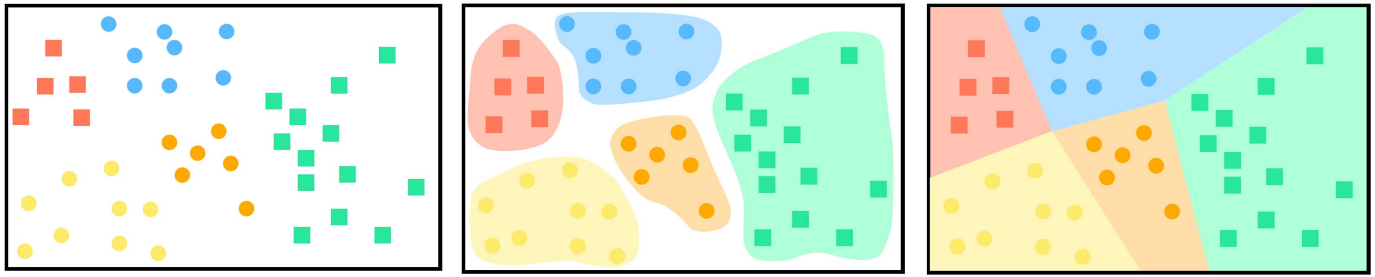

Nearest Neighbor Classifier (Not a CNN)

- Compare a new image with all images in the dataset and pick label from the closest.
- Supervised learning (Labels needed)
- Problem we need to keep all training data and rerun for every time we need to classify (It's an lazy algorithm)
- Part 1: Measure the difference pixel by pixel (Not trivial)
  - Pixel-wise distance (L1)
  - Euclidean distance (L2)

k-Nearest Neighbor Classifier (Not a CNN)

> Part 2: Pick the closest neighbor, or in case of kNN pick the winning vote among the k closest Labels.

> Not very useful for general image classification due to impact from all sort of variations

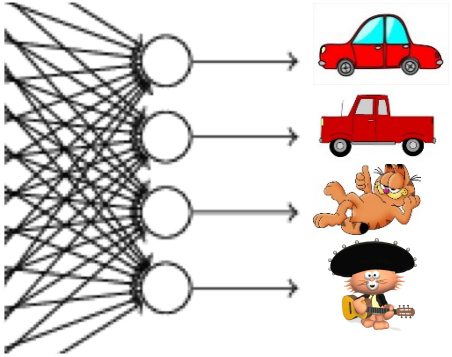

124	18	8		80	34	17		44	16	9	
34	204	5	-	120	180	13	=	86	24	8	Sum 374
168	67	198		132	198	218		36	131	20	

Linear classification

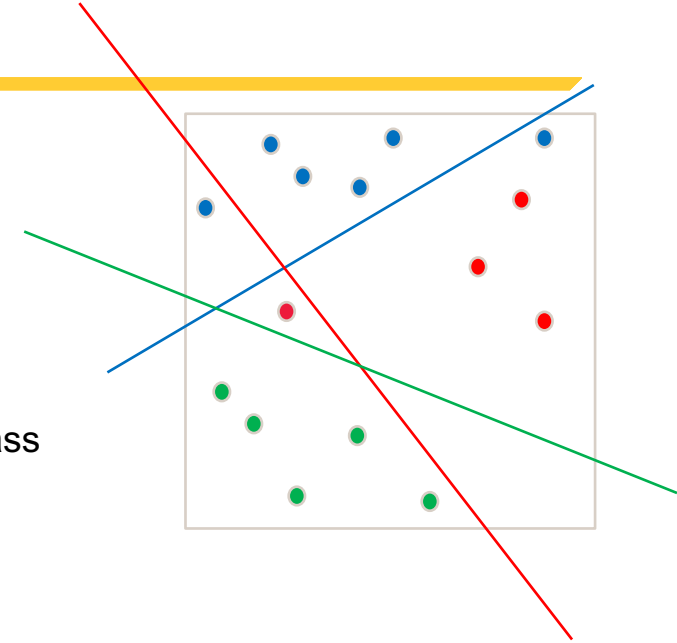

- Classification (Draw boundaries through the collected data)
- Linear classification
  - Score function is a weighted sum of all values
  - Write as a Matrix multiplication f(x1;W,b)
  - 2-dim example:
    - On a plane: straight lines will be the borders for each class
  - Loss function
    - To measure the error (eg how bad is this classification)
  - There are different loss functions
    - SVM classification (hinge loss)
    - Soft-max classification (cross-entropy loss)
      - Provides probabilities
- Major benefit: When the parameters (W,b) is known the dataset might be discarded

Optimization

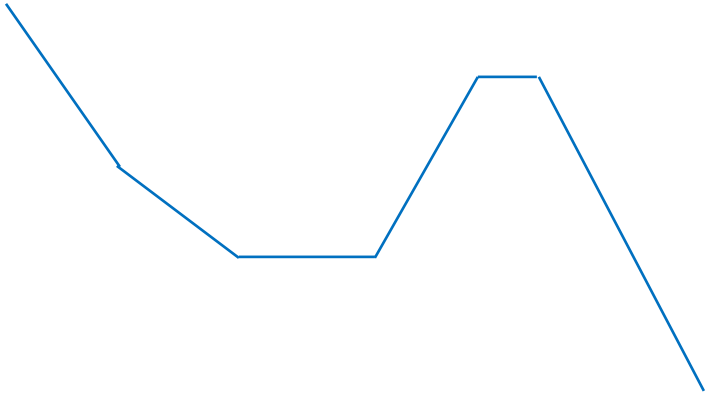

- How to find W (and b)
- Random search (bad idea)
- Follow the gradient (good idea)
  - Numerical gradient
  - Analytic gradient
- Algorithm:
  - Start with random set
  - Refine parameters using iteration and a step-size
  - Use the analytics gradient and a gradient descent algorithm
- Backpropagation (The key to training neural networks)
  - Use a network with simple nodes where you can solve the gradient

The artificial neuron

Brain neuron

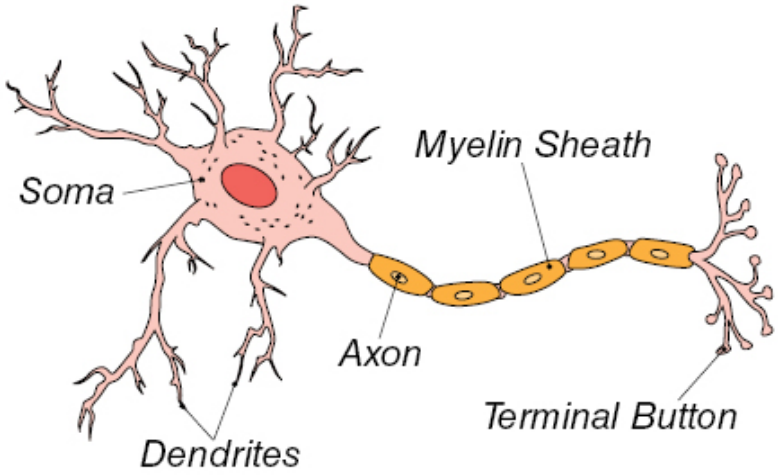

Artificial neuron

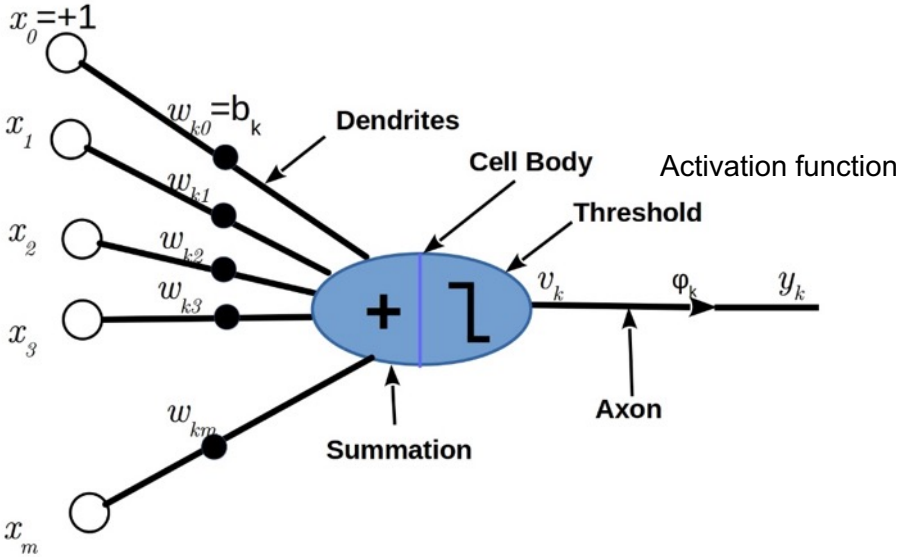

Coarse model of biological neuron

Neural network architectures

- Neural networks are neurons in a graph
  - Input layer
  - Hidden layer(s)
  - Output layer
- No loops
- Reason to use layers
  - Express them as vector matrix multiplications
  - Groups of neurons will approximate non-linear function
- Good properties
  - Cheap to use
  - Difficult to train
- Modern CNN have 10-20 layers and 100M parameters
  - Deep learning

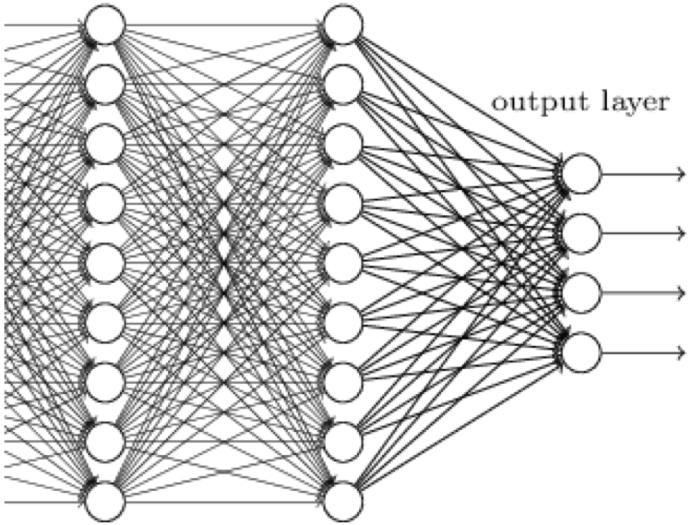

The architecture: Feature extraction

Simple filters get combined into more complex shapes in later layers

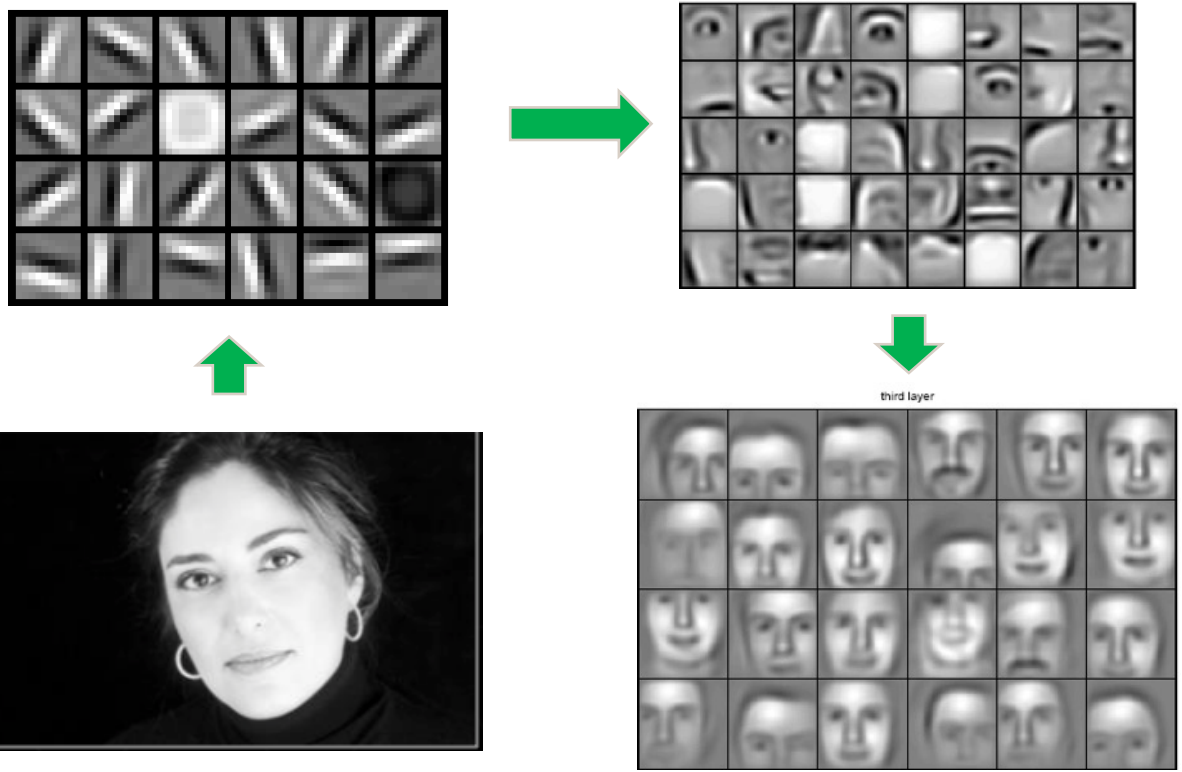

The architecture: Example, modified Alexnet (ZF-5)

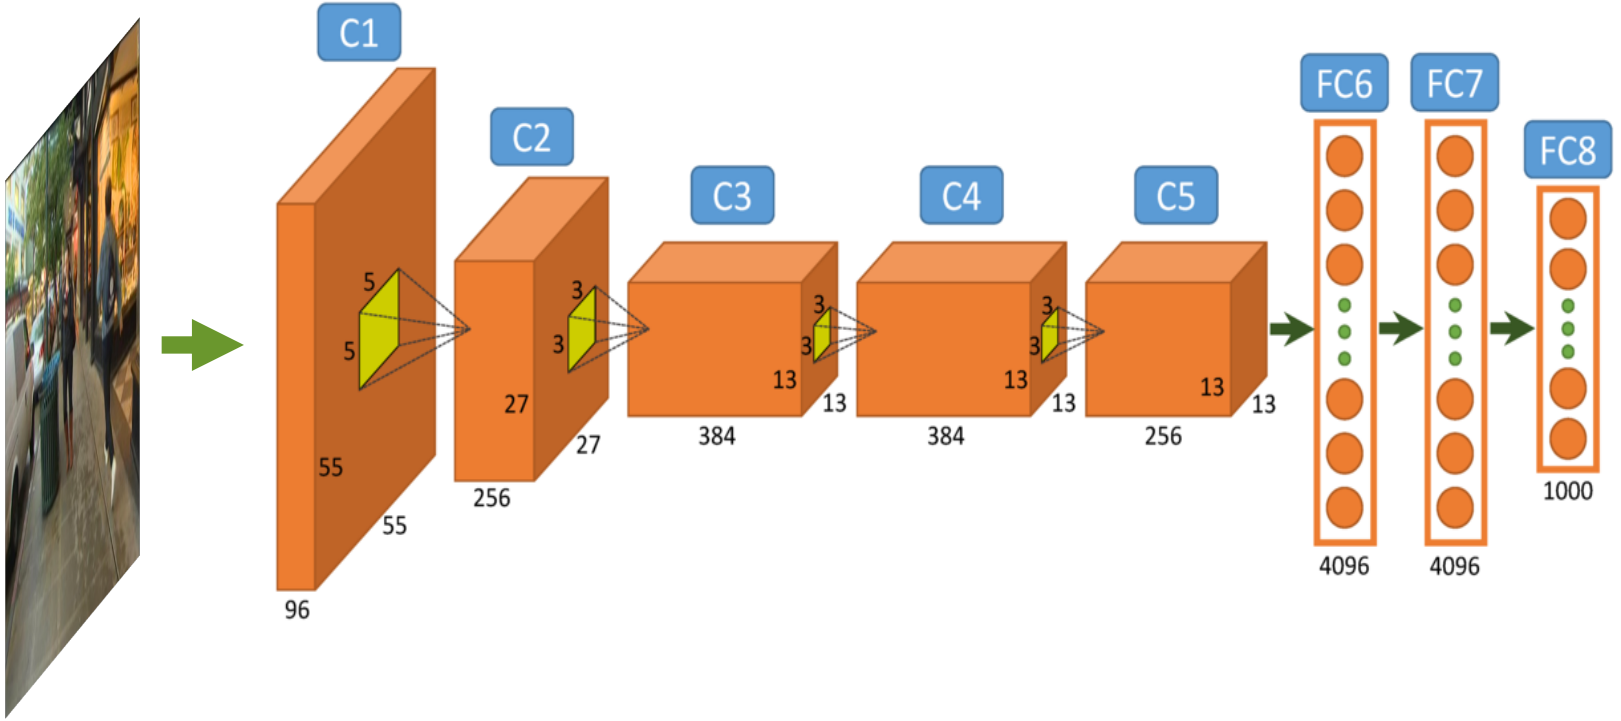

Training the network

- Initialize every weight in the system with **small random numbers**
  - Uniform distributed
  - Gaussian distributed
  - Special algorithms exists based on this distributions
- Starting with all weights the same does not work well

Training phase (in principle)

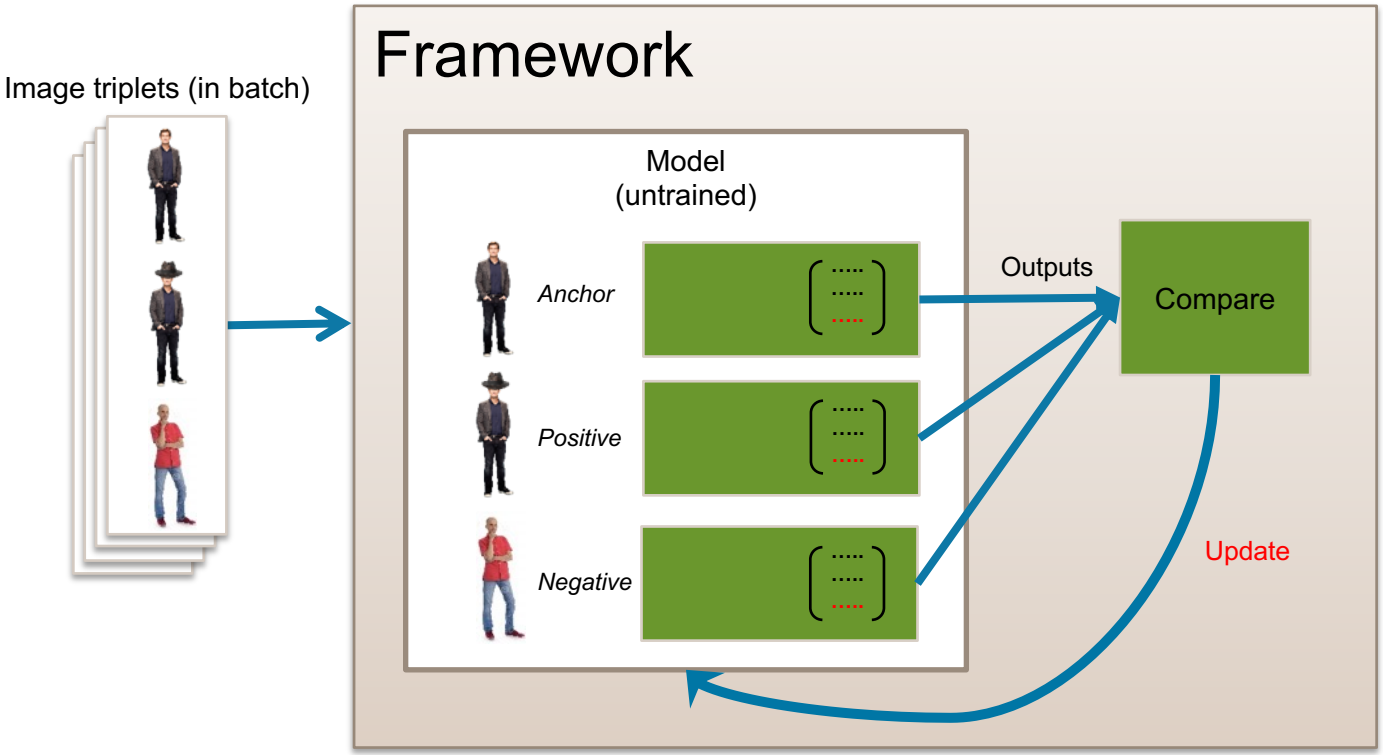

Framework used during training
- Executed on internal PC
- Done before deployment
- HW intense

Deployment (in principle)

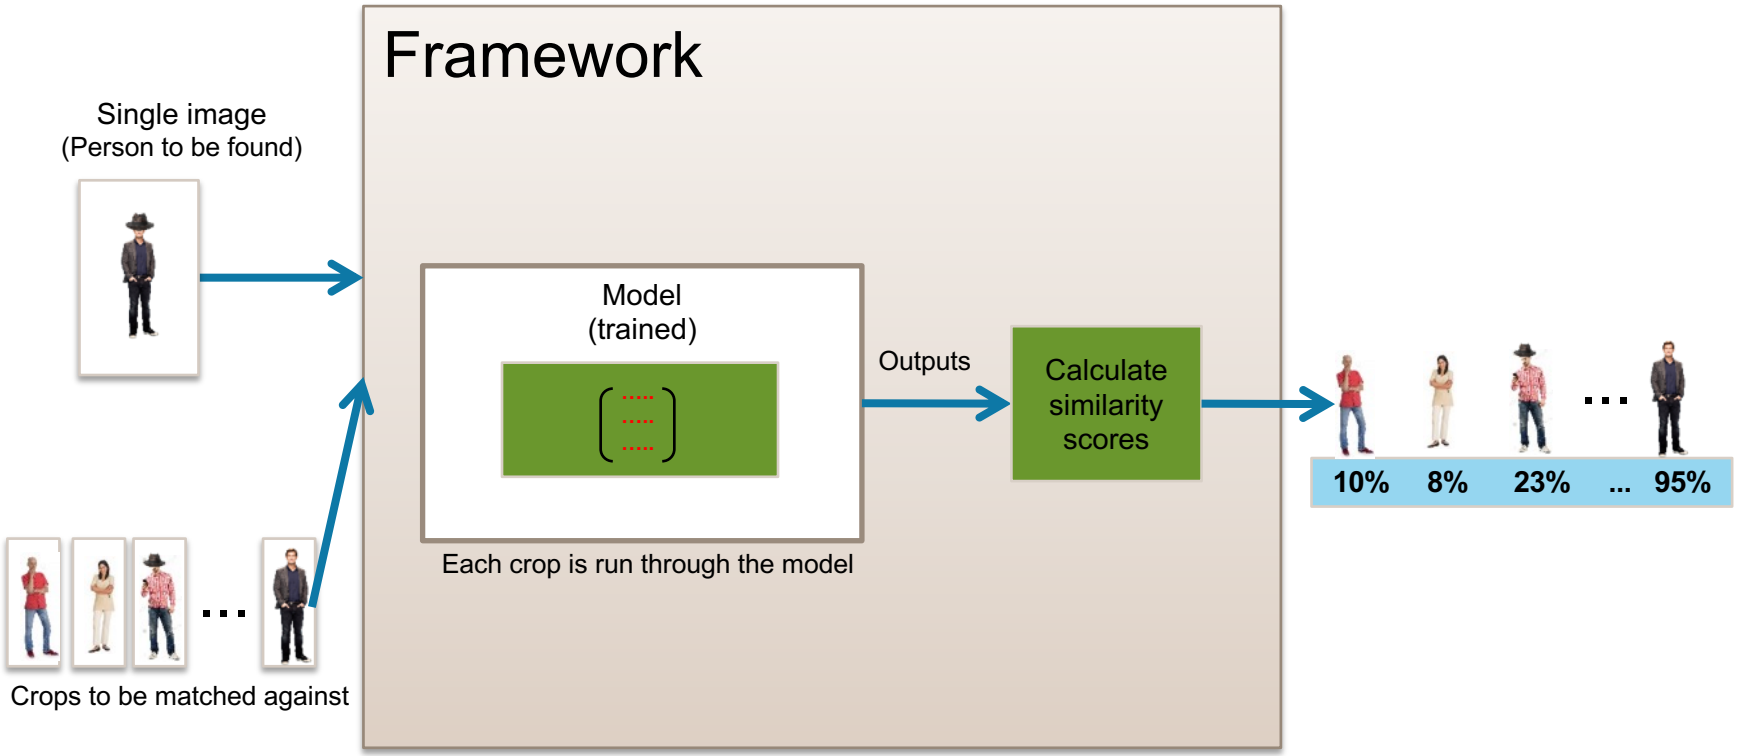

Training data

- Typical very large data sets are needed
- Divide the data training set into parts, with different usage
  - Training set, test set, validation set
  - Avoid data contamination (Using the validation set for training)
  - Validation set is only used **once** (for validating the final results)
- Overfitting/underfitting
- Stop training when validation error is flatting out

On chip CNN

- Easy to build the forward path
- Requires a lot of memory storage and bandwidth
- Currently a parameter(data) movement problem
- Hot research area
- Algorithm optimizations
  - Pruning
  - Compression
  - Retraining

Confusion matrix

Ground Truth \ Prediction	Positive	Negative
Positive	True Positive	False Negative
Negative	False Positive	True Negative

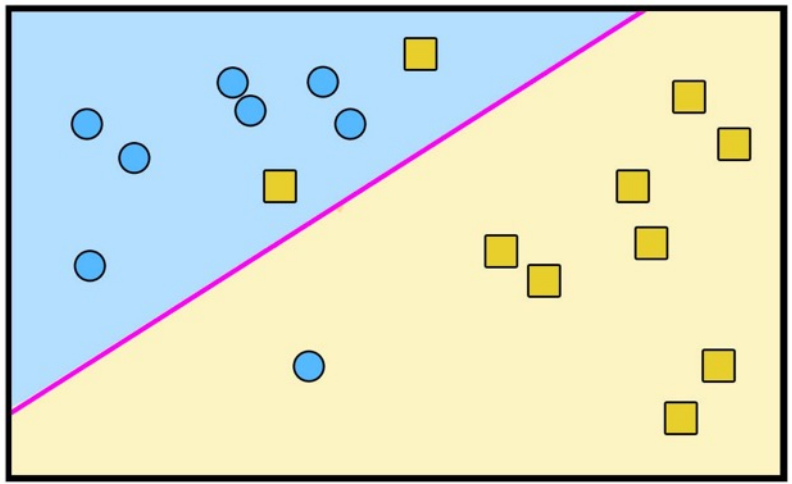

Characterize the classifier

$$\text{Accuracy} = \frac{TP+TN}{All}$$

$$\text{Recall} = \frac{TP}{TP+FN}$$

$$\text{Precision} = \frac{TP}{TP+FP}$$

$$\text{F1 Score} = \frac{2*TP}{2*TP+FP+FN}$$

Use precision and recall together

- Accuracy is a common measure but:
  - Comparing machine learning implementations requires both prescision and recall to be evaluated
  - Perfect score for Precision only is not enough
  - Perfect score for Recall only is not enough
- F1-score is a “harmonic mean” combining precision and recall.
- Important: This information is not enough if the probabilities are not equal.

Example: Finding wanted criminals at the stadium entrance

- Assume 1% of visitors to be wanted due to minor crimes
- Assume 10000 coming to the stadium
- The face detector performance is not that bad, show in the confusion matrix to the right.
- We do not want to miss criminals
- It is better to have some extra false alarms

		Prediction	
		Positive	Negative
Ground Truth	Positive	TP 0.99	FN 0.01
	Negative	FP 0.02	TN 0.98

But the confusion matrix is not correct!

- In this case this confusion matrix can't be used
- One important critical piece of information is missing
- We are looking for only 1% of all the visitors!
- The chance of being a criminal in this example is not 50/50

		Prediction	
		Positive	Negative
Ground Truth	Positive	TP 0.99	FN 0.01
	Negative	FP 0.02	TN 0.98

Deriving the proper confusion matrix

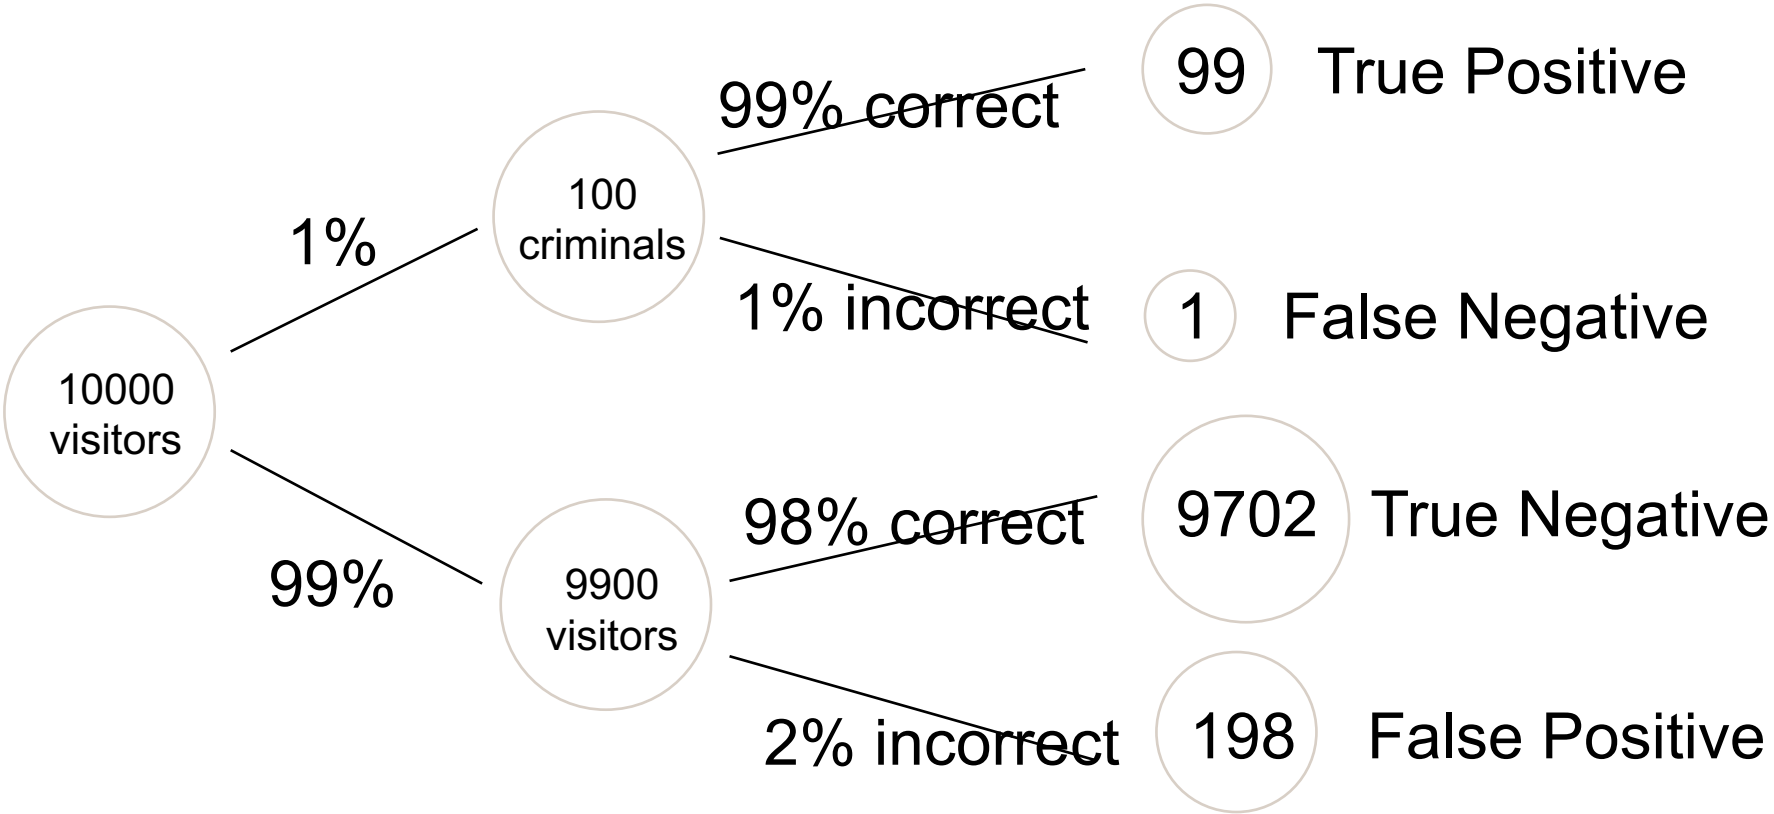

Proper confusion matrix assuming 1% of visitors wanted

- Precision TP/(TP+FP) = 33%
- Recall TP/(TP+FN) = 99%
- Our face detection algorithm has 99% probability to correctly find criminals
- Still, 2/3 of the time it detects someone the it is a false alarm.
- Most of our alarms are wrong!
- Things get even worse with 0.01% criminals…
- And if we independently captures 100 images of every person…
- This setup can only be used for capturing murders where someone's live would be in danger.

		Prediction	
		Positive	Negative
Ground Truth	Positive	TP 99	FN 1
	Negative	FP 198	TN 9702

Learn more about CNN/Deep learning

Stanford CS class CS231n:

Convolutional Neural Networks for Visual Recognition

http://cs231n.github.io

Stefan Lundberg

Beyond deep learning – High dimensional (HD) computing

- Compute-oriented solutions will be replaced by data-oriented solutions
- Classic methods are now plateauing
- Deep learning are powerful but very expensive to train and debug
- The need:
  - Learning based approach, statistical touch, not digital/discrete, redundant, flexible data representation, patterns, randomness as a feature.

HD compute

- System operates on high-dimensional vectors
  - Algorithms operating on vectors
- Traditional compute operates on bits and numbers
  - Logical circuit operating with bits and arithmetic
- Example of problems that can be solved with HD-compute:
  - Language identification
  - Music classification
  - Electromyography gesture classification
  - Scene Analysis and Spatial Reasoning

Learn more about High-dimensional compute.

Lund Circuit Design Workshop 2018
Invited speaker: Jan Rabaey, UC Berkley US

cdworkshop.eit.lth.se JanRabaey HD-Computing

EE380: Computer Systems Colloquium Seminar Computing with High-Dimensional Vectors Speaker: Pentti Kanerva, Stanford CSLI & UC Berkeley Redwood Center for Theoretical Neuroscience

https://www.youtube.com/watch?v=zUCoxhExe0o

Deep learning based applications


GENERATIONS /
SIGGRAPH 2018
VANCOUVER
12-16 AUGUST

VisemeNet: Audio-Driven Animator-Centric Speech Animation

YANG ZHOU, University of Massachusetts Amherst
ZHAN XU, University of Massachusetts Amherst
CHRIS LANDRETH, University of Toronto
EVANGELOS KALOGERAKIS, University of Massachusetts Amherst
SUBHRANSU MAJI, University of Massachusetts Amherst
KARAN SINGH, University of Toronto

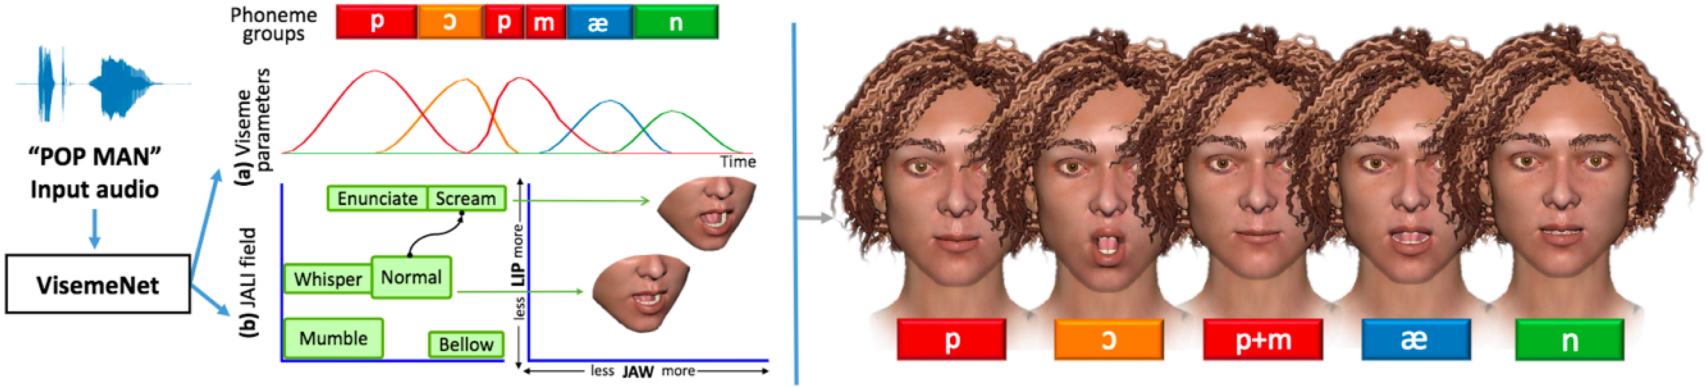

Fig. 1. VisemeNet is a deep-learning approach that uses a 3-stage LSTM network, to predict compact animator-centric viseme curves with proper co-articulation, and speech style parameters, directly from speech audio in near real-time (120ms lag).

Deep Video Portraits

HYEONGWOO KIM, Max Planck Institute for Informatics, Germany
PABLO GARRIDO, Technicolor, France
AYUSH TEWARI and WEIPENG XU, Max Planck Institute for Informatics, Germany
JUSTUS THIES and MATTHIAS NIESSNER, Technical University of Munich, Germany
PATRICK PÉREZ, Technicolor, France
CHRISTIAN RICHARDT, University of Bath, United Kingdom
MICHAEL ZOLLHÖFER, Stanford University, United States of America
CHRISTIAN THEOBALT, Max Planck Institute for Informatics, Germany

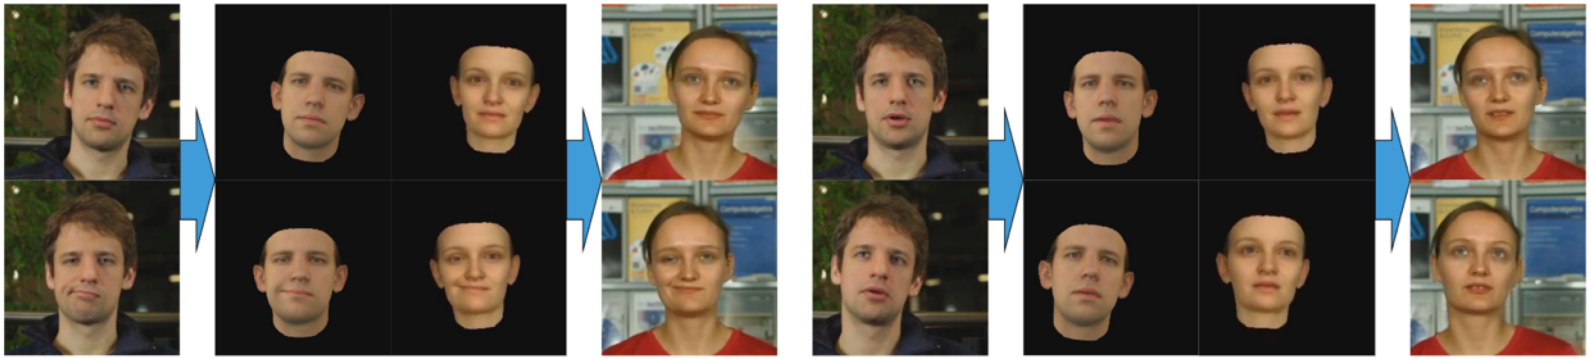

Fig. 1. Unlike current face reenactment approaches that only modify the expression of a target actor in a video, our novel deep video portrait approach enables full control over the target by transferring the rigid head pose, facial expression and eye motion with a high level of photorealism.

Deep video portraits results 1

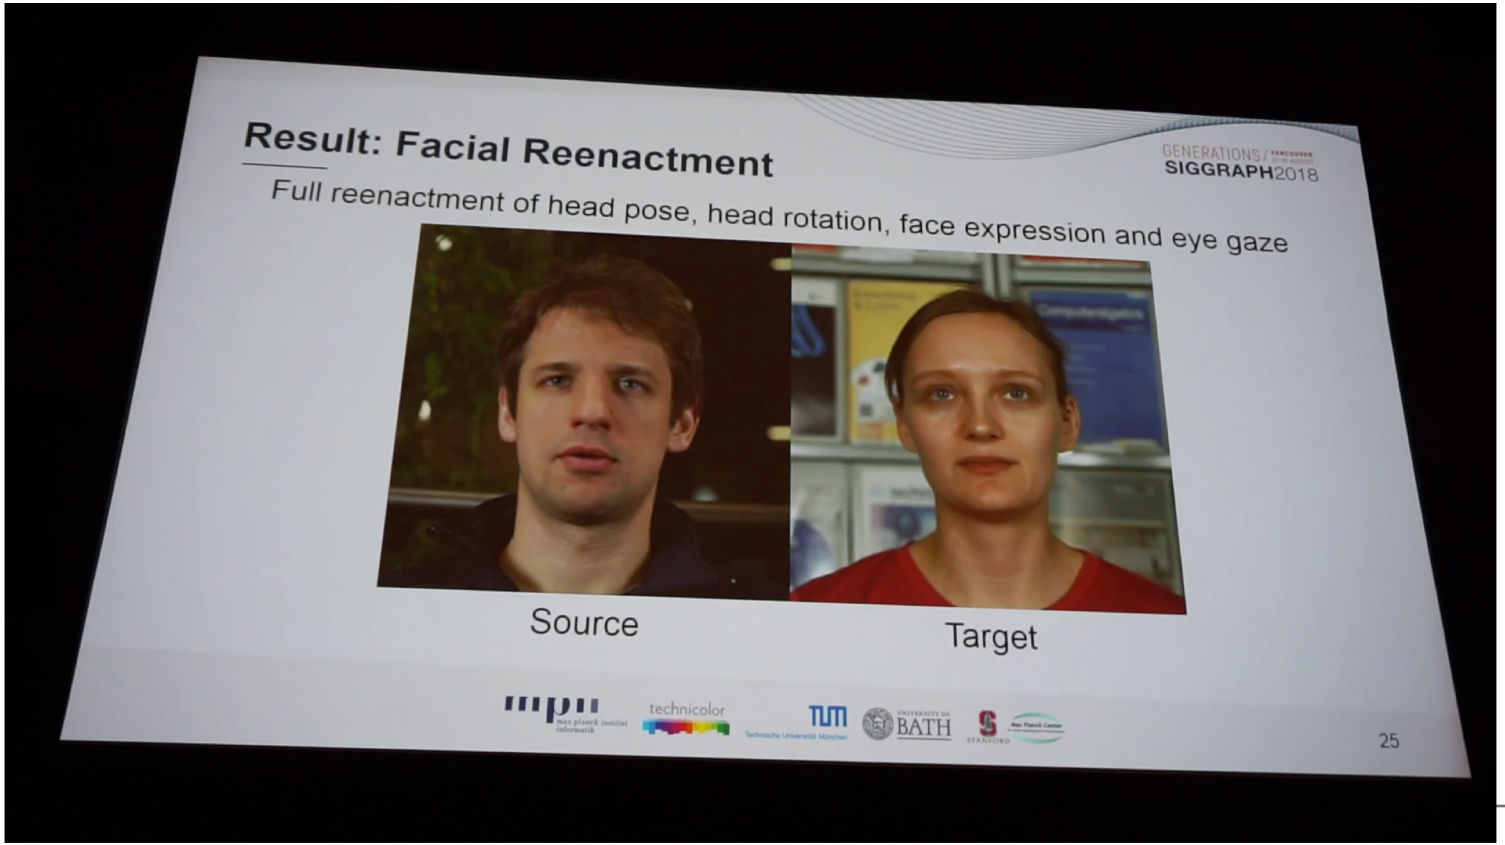

Deep video portraits results 2

High-Fidelity Facial Reflectance and Geometry Inference From an Unconstrained Image

SHUGO YAMAGUCHI*, Waseda University and USC Institute for Creative Technologies
SHUNSUKE SAITO*, Pinscreen, University of Southern California, and USC Institute for Creative Technologies
KOKI NAGANO, Pinscreen
YAJIE ZHAO, USC Institute for Creative Technologies
WEIKAI CHEN, USC Institute for Creative Technologies
KYLE OLSZEWSKI, Pinscreen, University of Southern California, and USC Institute for Creative Technologies
SHIGEO MORISHIMA, Waseda University
HAO LI, Pinscreen, University of Southern California, and USC Institute for Creative Technologies

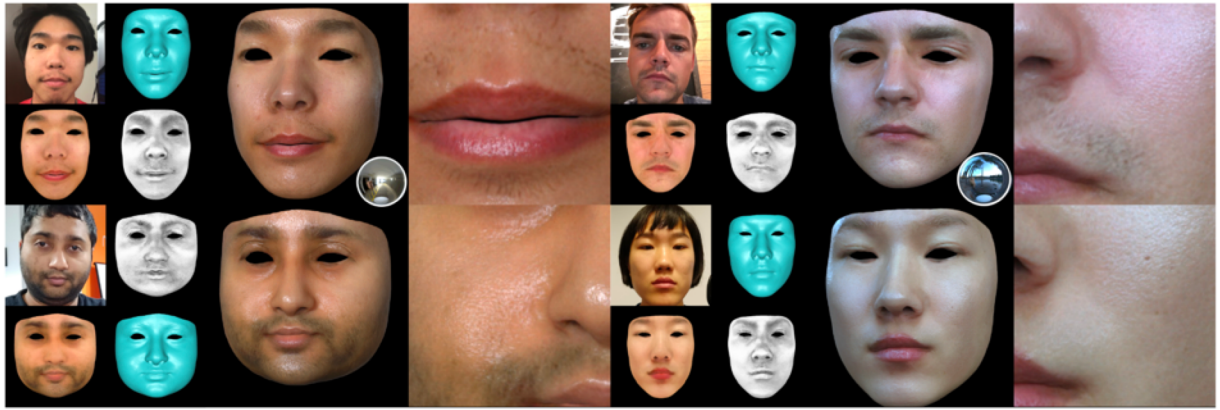

Fig. 1. Our system infers high-fidelity facial reflectance and geometry maps from a single image (diffuse albedo, specular albedo, as well as medium- and high-frequency displacements). These maps can be used for high-fidelity rendering under novel illumination conditions.

3D model from single 2D image

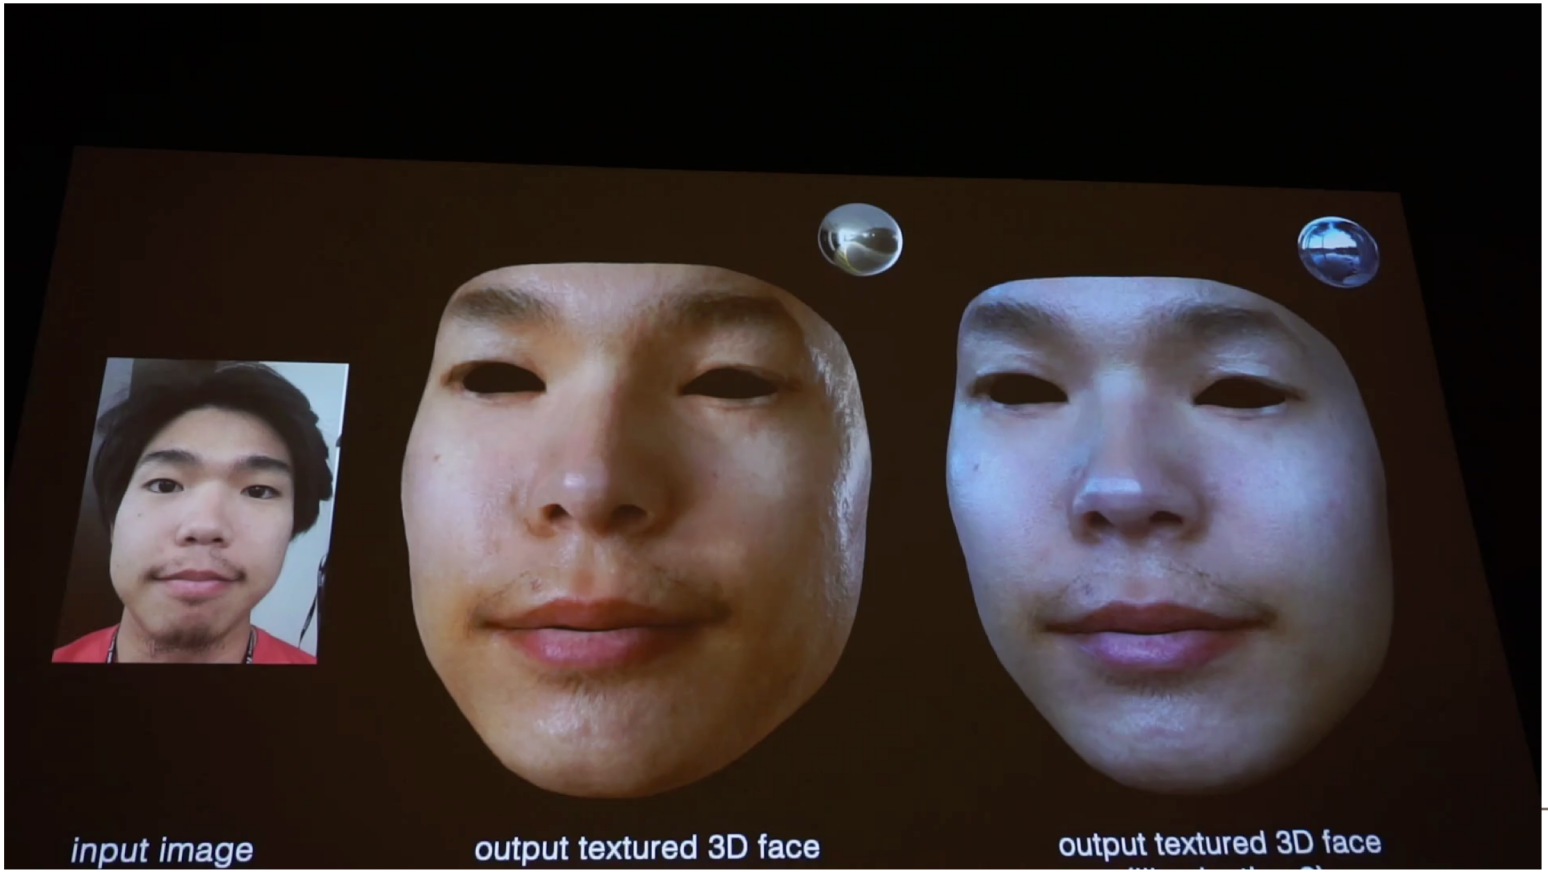

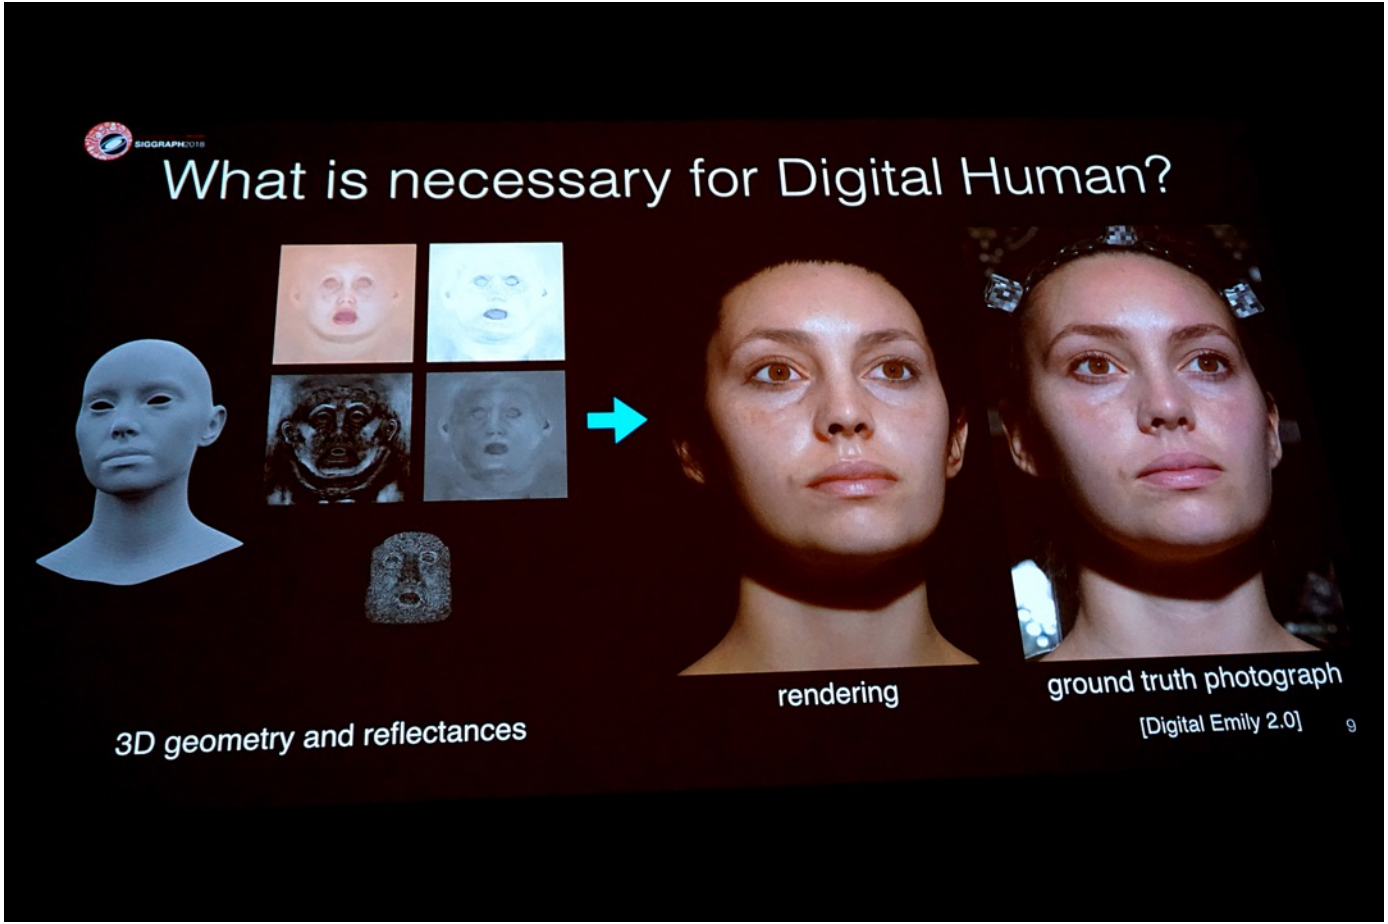
SIGGRAPH2018
What is necessary for Digital Human?
3D geometry and reflectances
rendering
ground truth photograph
[Digital Emily 2.0]
9

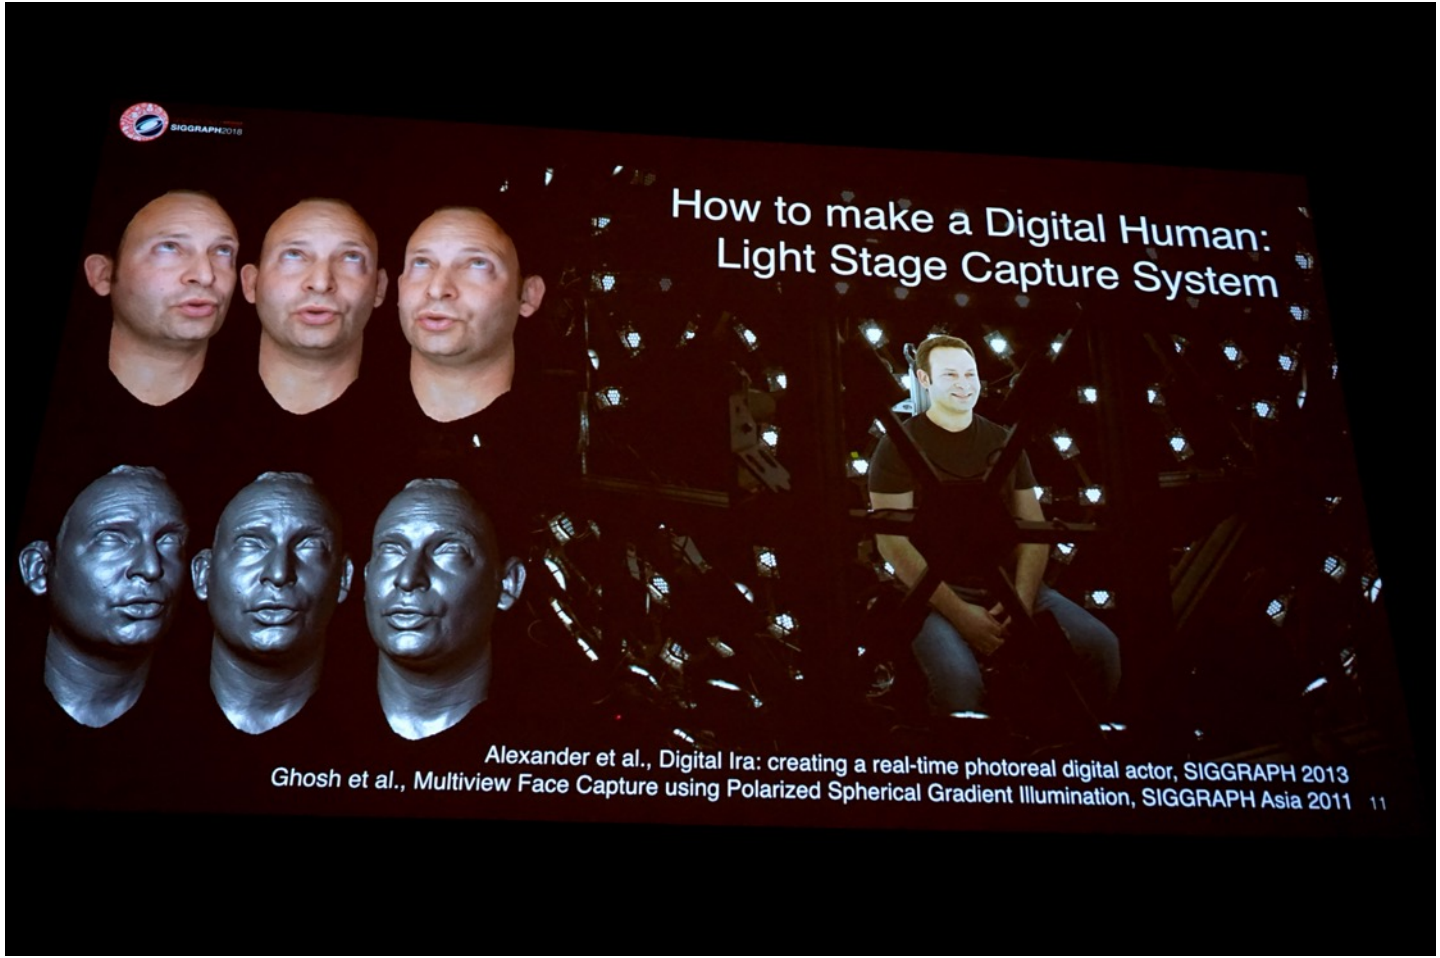
SIGGRAPH2018
How to make a Digital Human:
Light Stage Capture System
Alexander et al., Digital Ira: creating a real-time photoreal digital actor, SIGGRAPH 2013
Ghosh et al., Multiview Face Capture using Polarized Spherical Gradient Illumination, SIGGRAPH Asia 2011
11

Texture Completion (In-painting)

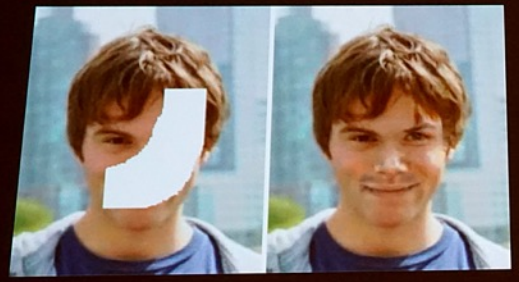

Image In-painting
(Iizuka et al. 2017)

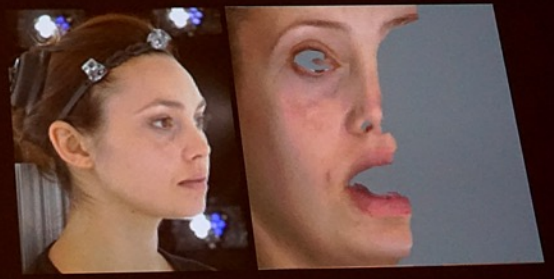

Our Completion Target

Iizuka et al., Globally and Locally Consistent Image Completion, SIGGRAPH 2017

Our Training Data

scanned images

base mesh

diffuse albedo

specular albedo

displacement

Light Stage Data (20 id)

Chicago Face Database
(600 id)


GENERATIONS
SIGGRAPH2018
Deep Exemplar-based Colorization
Mingming He*1 Dongdong Chen*2 Jing Liao3 Pedro V. Sander1 Lu Yuan3
(* Equally Contributed)
1Hong Kong UST
2University of Science and Technology of China
3Microsoft Research
SIGGRAPH2018

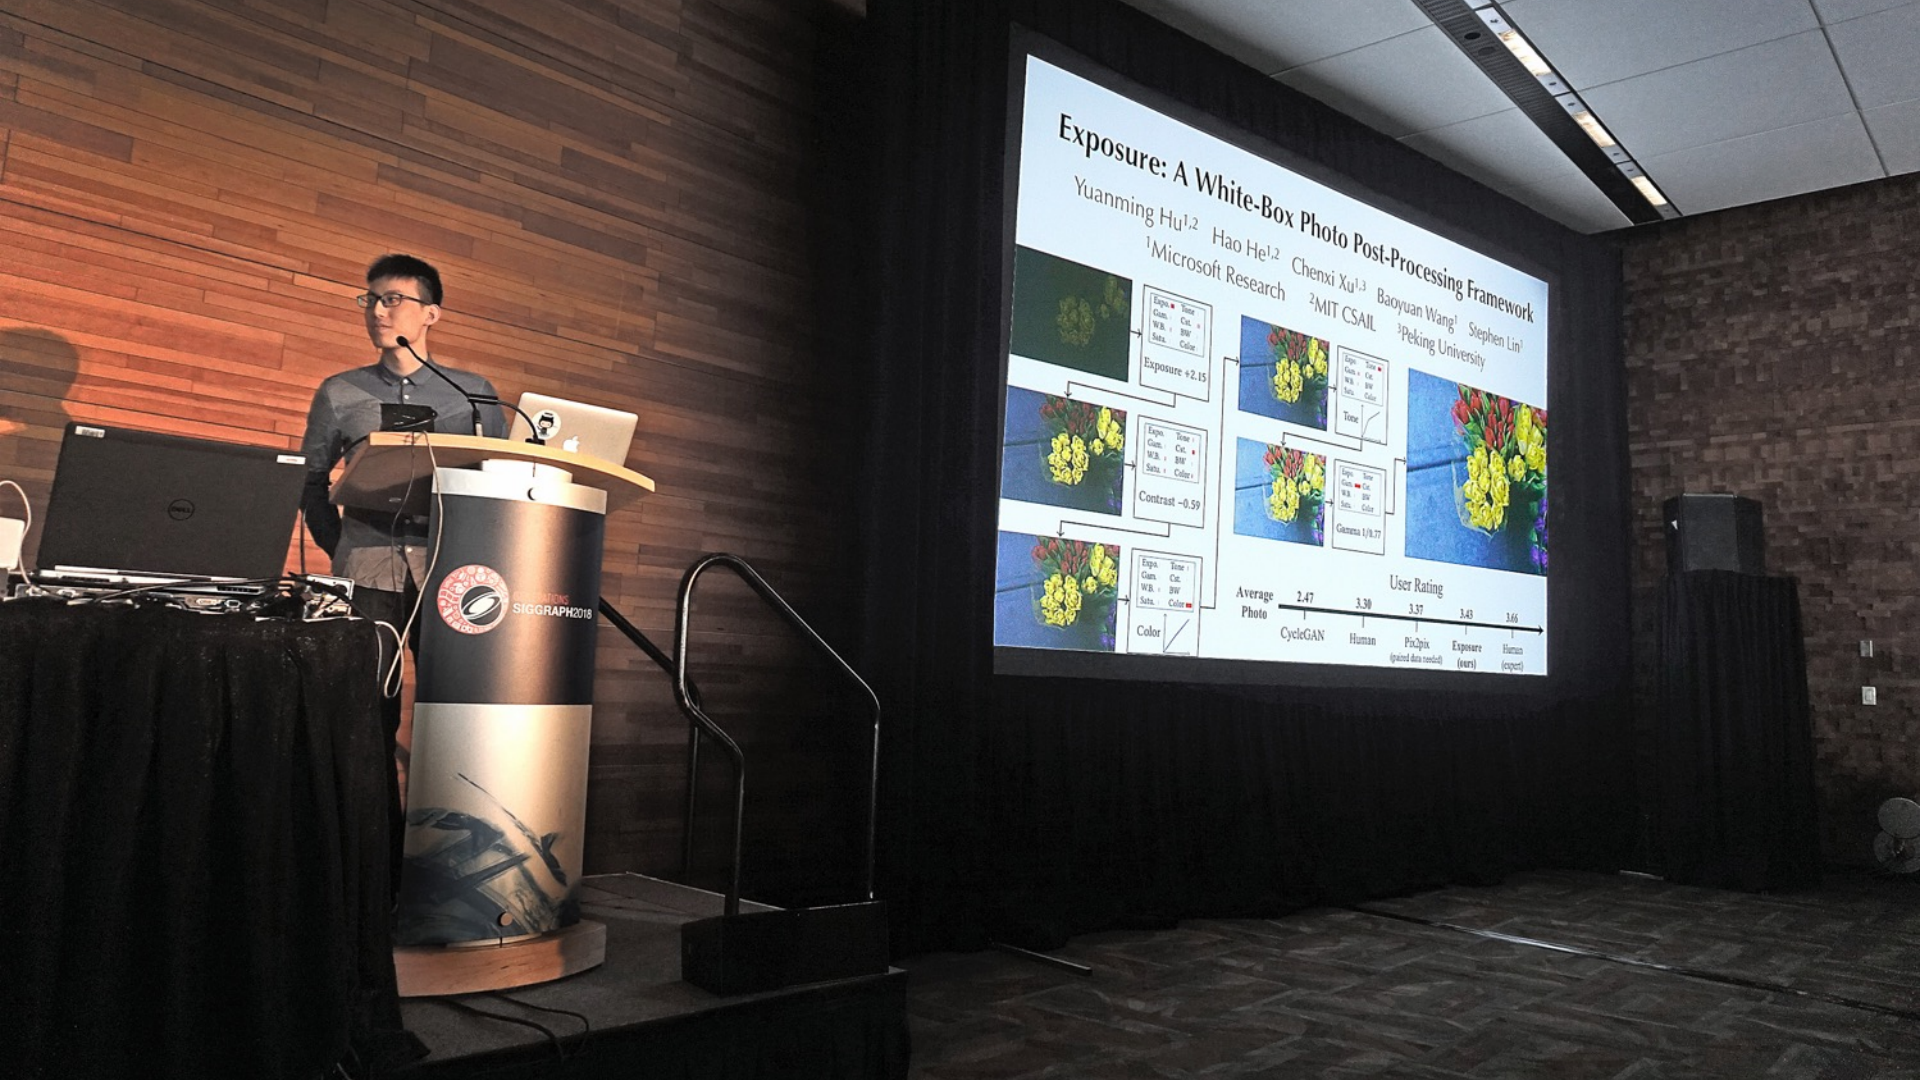
Exposure: A White-Box Photo Post-Processing Framework
Yuanming Hu[1,2] Hao He[1,2] Chenxi Xu[1,3] Baoyuan Wang[1] Stephen Lin[1]
[1]Microsoft Research [2]MIT CSAIL [3]Peking University
Exposure +2.15
Contrast −0.59
Color
Tone
Gamma 1/0.77
User Rating
Average Photo
2.47
3.30
3.37
3.43
3.66
CycleGAN
Human
Pix2pix (paired data needed)
Exposure (ours)
Human (expert)
SIGGRAPH2018

HALFTONE

"Reprographic technique that **simulates continuous tone** imagery through the use of **dots**, varying either in **size** or in **spacing**, thus generating a gradient-like effect."

- wikipedia.org

HALFTONE IMAGE EDITING

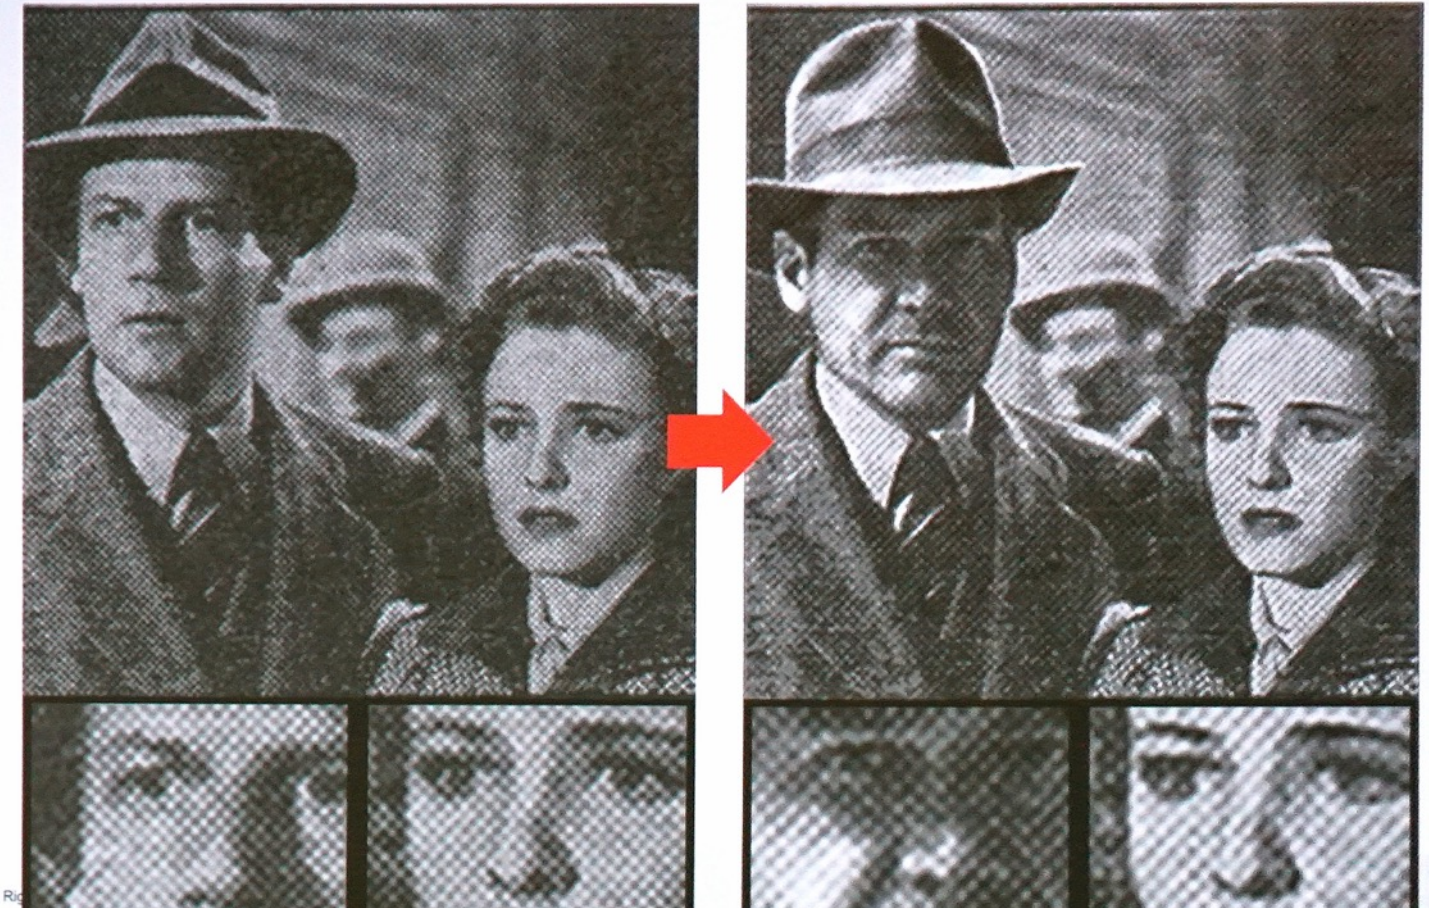

Texture transfer

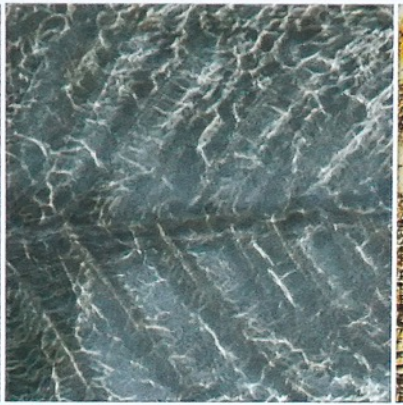

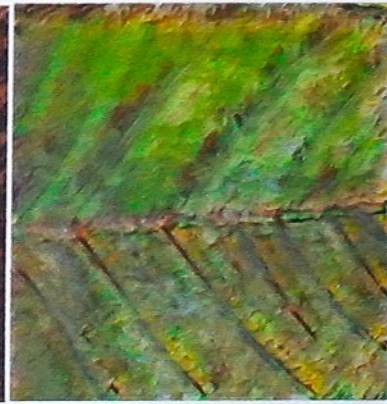

Synthetic Depth-of-Field with a Single-Camera Mobile Phone

NEAL WADHWA, RAHUL GARG, DAVID E. JACOBS, BRYAN E. FELDMAN, NORI KANAZAWA, ROBERT CARROLL, YAIR MOVSHOVITZ-ATTIAS, JONATHAN T. BARRON, YAEL PRITCH, and MARC LEVOY,
Google Research

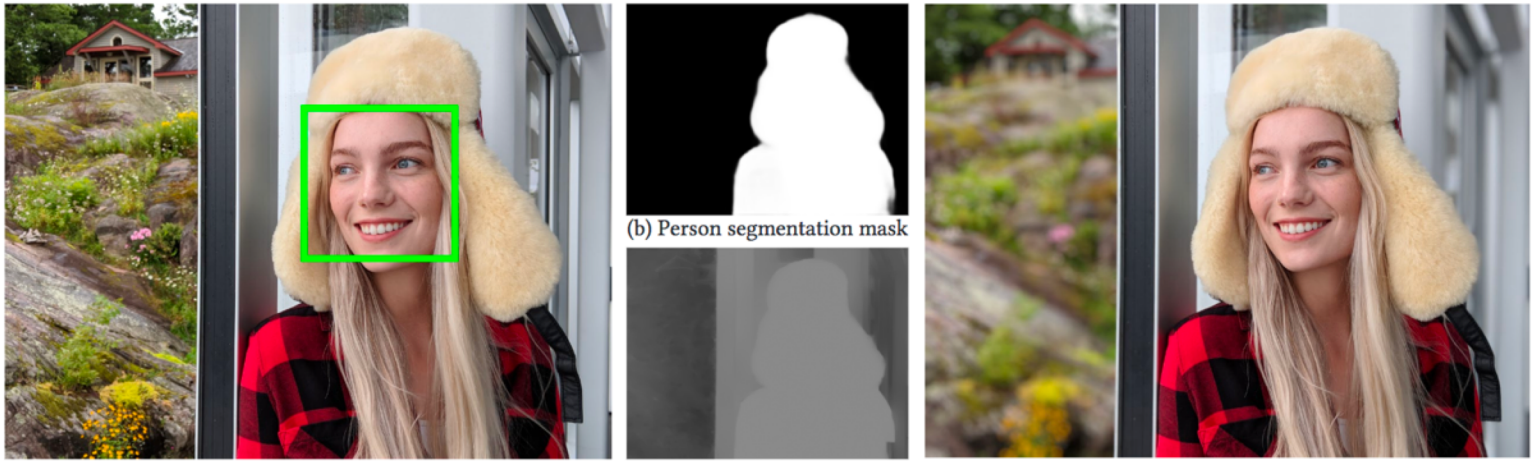

(a) Input image with detected face

(b) Person segmentation mask

(c) Mask + disparity from DP

(d) Our output synthetic shallow depth-of-field image

Fig. 1. We present a system that uses a person segmentation mask (b) and a noisy depth map computed using the camera's dual-pixel (DP) auto-focus hardware (c) to produce a synthetic shallow depth-of-field image (d) with a depth-dependent blur on a mobile phone. Our system is marketed as "Portrait Mode" on several Google-branded phones.

Overview

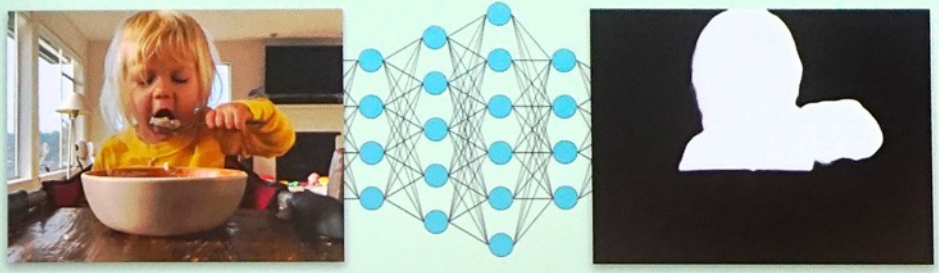

Neural Networks for Person Segmentation

Synthetic Shallow Depth of Field Imag

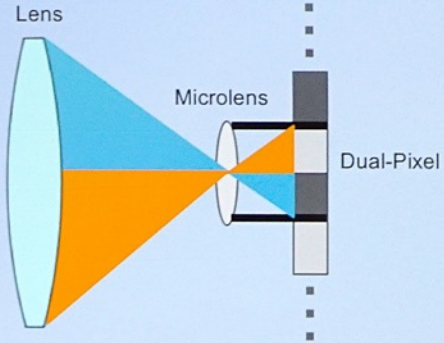

Depth from Dual-Pixel Sensor

Issue: Narrow Baseline

~6.5 cm

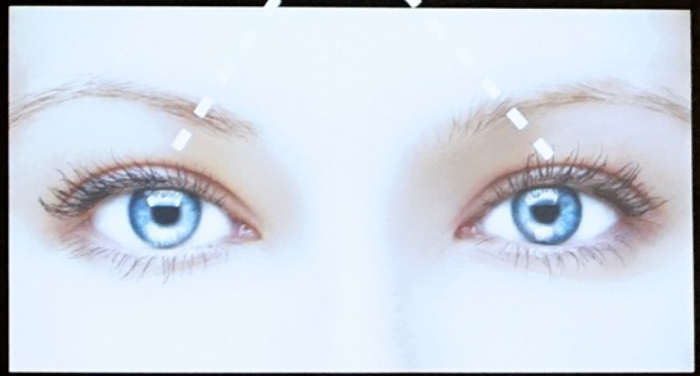

~1.5 cm

Problem

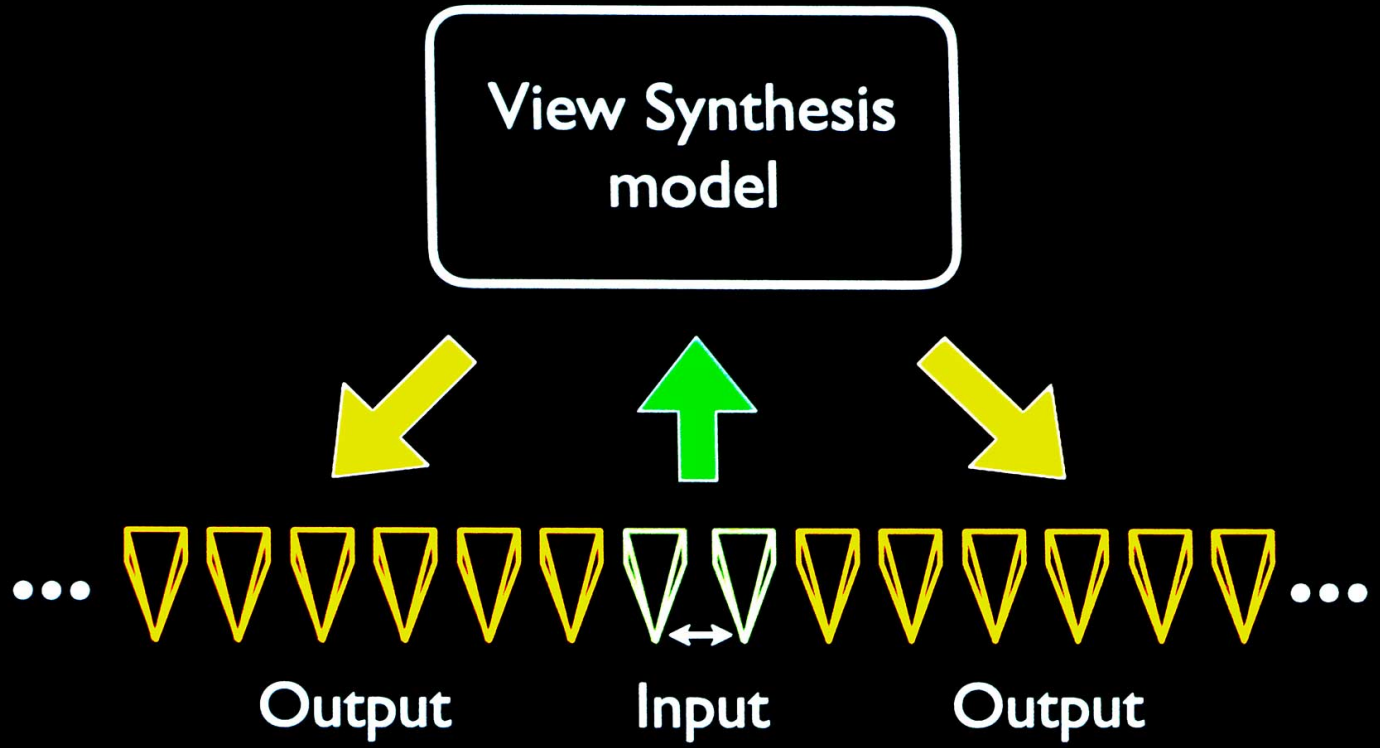

A New Dataset

>7,000 video clips of high-quality real estate scenes from YouTube

CRIS

CRIS CPU architecture

General Registers:

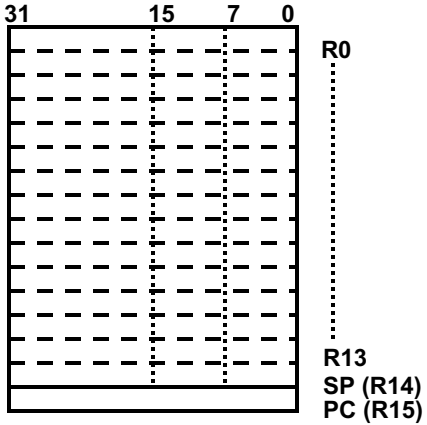

Special Registers:

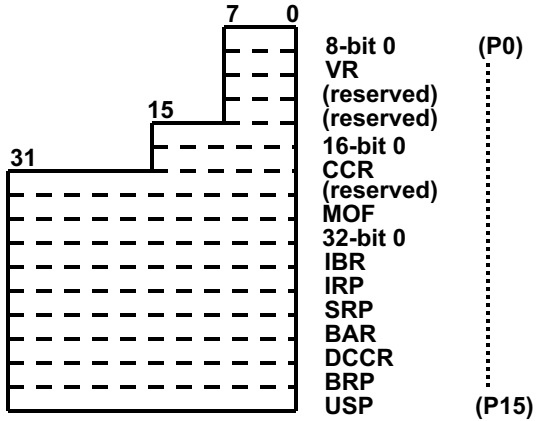

Pipelining scheme:

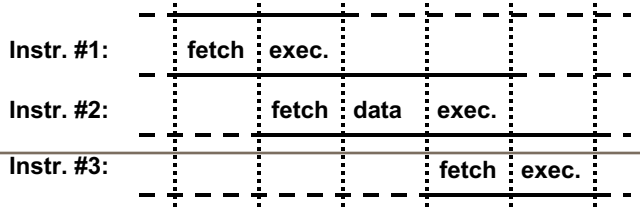

16-bit Instruction Format:

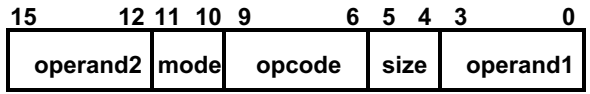

CRIS architecture history

- Development started in 1991.
- First implementation in silicon in ETRAX 1 (1993).
- GCC backend
- C/C Cache added in ETRAX 100 (1998).
- MMU and multiply added in ETRAX 100LX (2000).
- Part of the official Linux distribution
- Will be used in AXIS new camera controller chip (2001).

CRIS Block Diagram

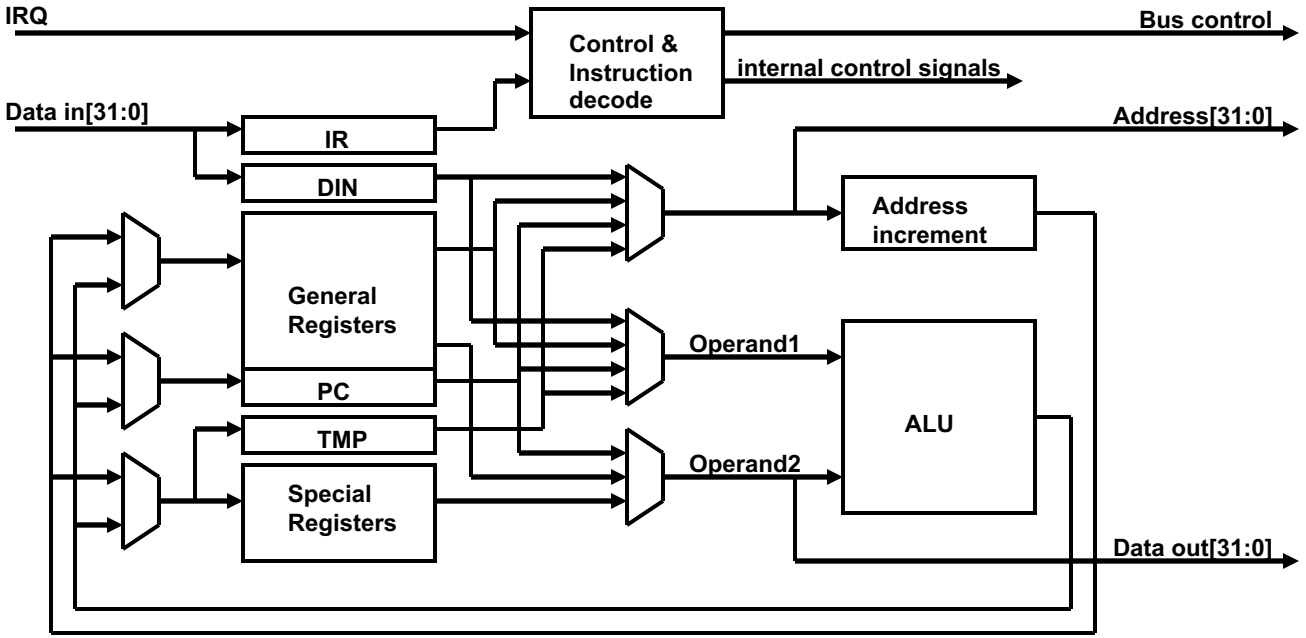

Future of the CRIS architecture

- ASIC technology improvements will finally make the original architecture obsolete.
- Alternatives for future high end products:
  - Select a commercial core. (high cost, low flexibility).
  - Design a new architecture.
  - Improve the existing architecture. (Add more pipelining, multiple issue etc.)
  - Multi-processor approach. (Will have large impact on software design.)
- CRIS will still have a long life in low end applications. (e.g. man CPU offloading).

What to do if you start today

- Embedded processor market has matured.
  - You can find suitable and well supported cores for most applications today.
  - License and royalty fees are still rather expensive.
- Gate count is no longer critical.
  - The high gate count of the commercial alternatives can be accepted today, because the total area will still be small.
- Embedded memories change the scenario.
  - Large memories on-chip open up new architecture possibilities.

Security electronics

Example of security electronics

- Alarm systems
  - Property protection
  - Loss prevention
- Fire
  - Detection
- Gates
  - Automatic gates
  - Toll systems
- Communications equipment
  - Sound
  - Radio
  - Datacom
- Law enforcement
- Lighting
  - Visible
  - IR
- Camera systems
  - CCTV
  - IP Cameras
- Access control
  - ID’s, badges and readers
  - Door ctrl
- Vehicles
  - Equipment
  - Protection
- Locks
  - Doors
  - Safe
- IT security
  - Computer security
  - Network security
- Home automation
  - IoT
- Personal emergency
- Public safety
- Special equipment
  - Gas detectors
  - Radar
- Services

Trade shows

- IFSEC London
  - http://www.ifsec.events/international/
- ISCWest LasVegas
  - http://www.iscwest.com
- Security Essen Germany
  - http://www.security-essen.de/
- Conference/Seminars
  - Education
  - Certifications
- Exhibit
  - Meet the vendors
  - See the equipment

Alarm systems

- > Technology
- > Equipment
  - – Detectors
  - – Panels
  - – Keypads
  - – Wireless modules
  - – GSM/3G callers
  - – Lights/Sirens
  - – Smoke generators
  - – etc
- > Vendors
  - – 1000+ far east
  - – Large western
    - – Like www.dsc.com

New requirements

- > Wireless
- > App controlled
- > email, whatsapp notifications
- > Facebook/Google login
- > alarm.com compatible
- > ifttt.com/recipes (If this, then that)

Developing electronics

Development challenges

- Distance to production
- China
- Open source
- Free information
- Crowd founding (Kickstarter etc)
- Time to market
- Quality

Mandatory certifications

- UL
  - Certification for professional US market
  - The “Underwriters laboratories”
  - American safety and certification company
- CE-marking
  - Mandatory conformity marking for the European economic area
  - ***Conformité Européenne***, meaning **European Conformity**
- FCC compliance statement
  - Mandatory marking for all electronics
  - Federal communications Commission

FCC

Other certifications

- Consumer technology
  - HDMI, USB
  - WiFi
  - SD-Card
- 2G/3G/4G/5G
  - Carrier certification
- Vendor specific requirements
  - Apple Lightning
  - Apple App Review
  - Windows hardware certification
- Non mandatory testing
  - Technischer Überwachungsverein – TÜV
  - Technical Research Institute of Sweden - SP
- Technology license
  - Dolby
  - H.265
- Known difficult areas:
  - Onboard equipment for train and aircrafts
  - Vehicles

Innovative product portfolio

AXIS D2050-VE

Axis' Zipstream technology – More video, less storage

- Reduce storage and bandwidth by an average 50% or more
  - Optimized for video surveillance
  - Fully compatible with H.264
  - New unique method
    - Acts on motion, details and noise
  - Radically lowering bandwidth and storag
    - Keep the essence

How much do I gain?

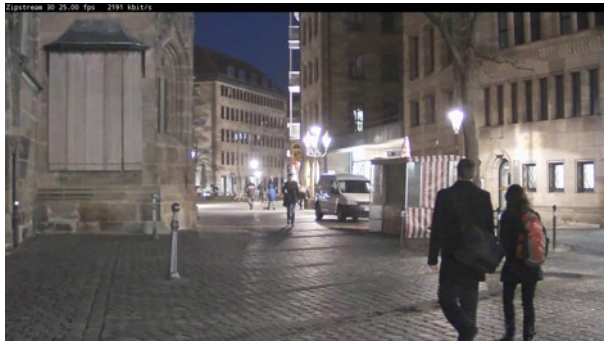

City surveillance:
Street level recording with small movements most of the time

Zipstream strength:	Dynamic GOP:	
High	On	60%

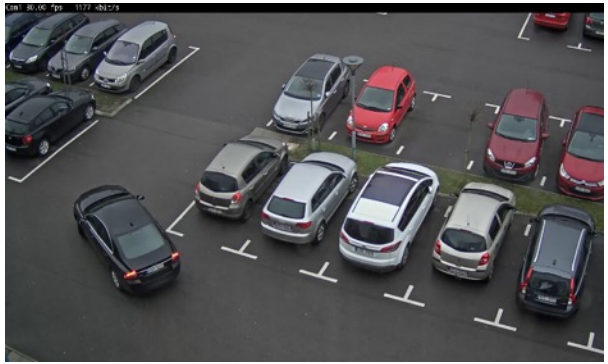

Outdoor VMD triggered recording:
Night time, average reduction for 12h surveillance.

Zipstream strength:	Dynamic GOP:	
High	On	75%

Lightfinder technology – an Axis innovation

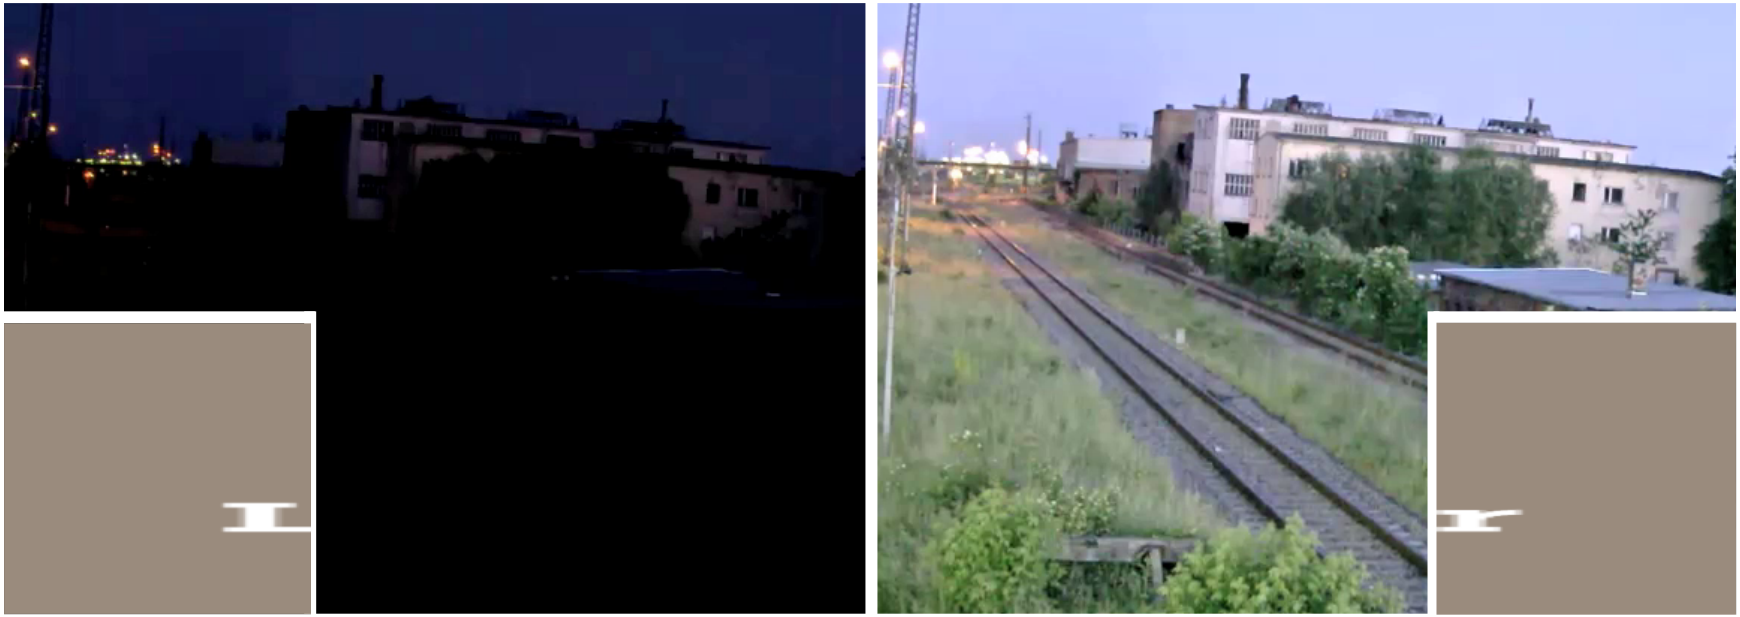

Extreme light sensitivity

Thermal network cameras – an Axis innovation

24/7 detection in tough conditions

> Bright lights

> Deep shadows

> Rain, snow and fog

> Smoke

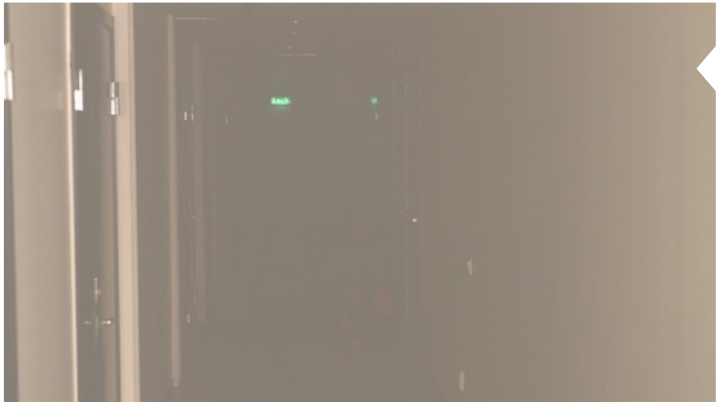

Thermal camera OFF

Thermal camera ON


AXIS


Thank you!
Megapixel
Quality
Ease of use
Network video
Reliability
open
Innovation
Easy installation
Protect
Service
all the way
support
integration
Image
usability
Thermal
Convergence
HDTV
Access
control
Partners
World
Leader
Intelligent
Competence
Quality
Focus
Video
encoder
Safe
Cameras
Outdoor
Secure
Global