



Artificial Intelligence!

The Future of Forecasting and S&OP?

Dr. Sven F. Crone

Assistant Professor Lancaster University
Co-Director Research Centre for Marketing & Forecasting
CEO iqast



Q&A via Slido.com #AI-tutorial

slido

How to Introduce Slido to Your Audience

At this event, we want to make sure we address your most burning questions. Therefore, we'll be using a simple audience interaction platform called **Slido**.

Slido allows you to submit your questions as well as upvote the questions of other participants. Questions with the highest number of votes will stand a better chance to get answered by speakers.

Throughout the event, you will also be able to express your opinion by voting on live polls.

It's really easy to join.

1. Please take out your smartphones and connect to the WiFi
2. Open the web browser
3. Go to www.slido.com and enter the event code, which is...

www.slido.com → #AI-workshop



Agenda

AI in Forecasting

- **Capabilities of AI**
 - The hype
 - How AI can “see”
 - AI Demos
- **Using AI in Forecasting**
 - AI in Forecasting
 - Case studies
- **The Gap**
 - Forecasting?
 - State of Forecasting
 - Sizing the Gap


Deep Learning



Picking up on the most commonly occurring

The New York Times

AI Art at Christie's Sells for \$432,500



Artificial intelligence (AI)
October 2018

"Tech tax" necessary to avoid dystopia, says leading economist

Jeffrey Sachs warns AI could lead to wealth being concentrated in the hands of a few thousand people

23 Oct 2018

Facebook and by today's robbers

Toby Walsh

23 Oct 2018

The five / Five scientific predictions by Professor Stephen Hawking

How do we stop technology falling hands?

Hamah Fry

ns&i

artificial intelligence VOICES

Pepper the AI robot

She stared at the ceiling, and it was a perfect politician's

TECH How art can be viola

AMERICAS Houston robot


HOME NEW Many v "trust is



"Edmond de Belamy, from La Famille de Belamy," by the French art collective Obvious, was sold on Thursday at Christie's New York. Christie's

By Gabe Cohn

The Hype on AI, ML & Deep Learning?



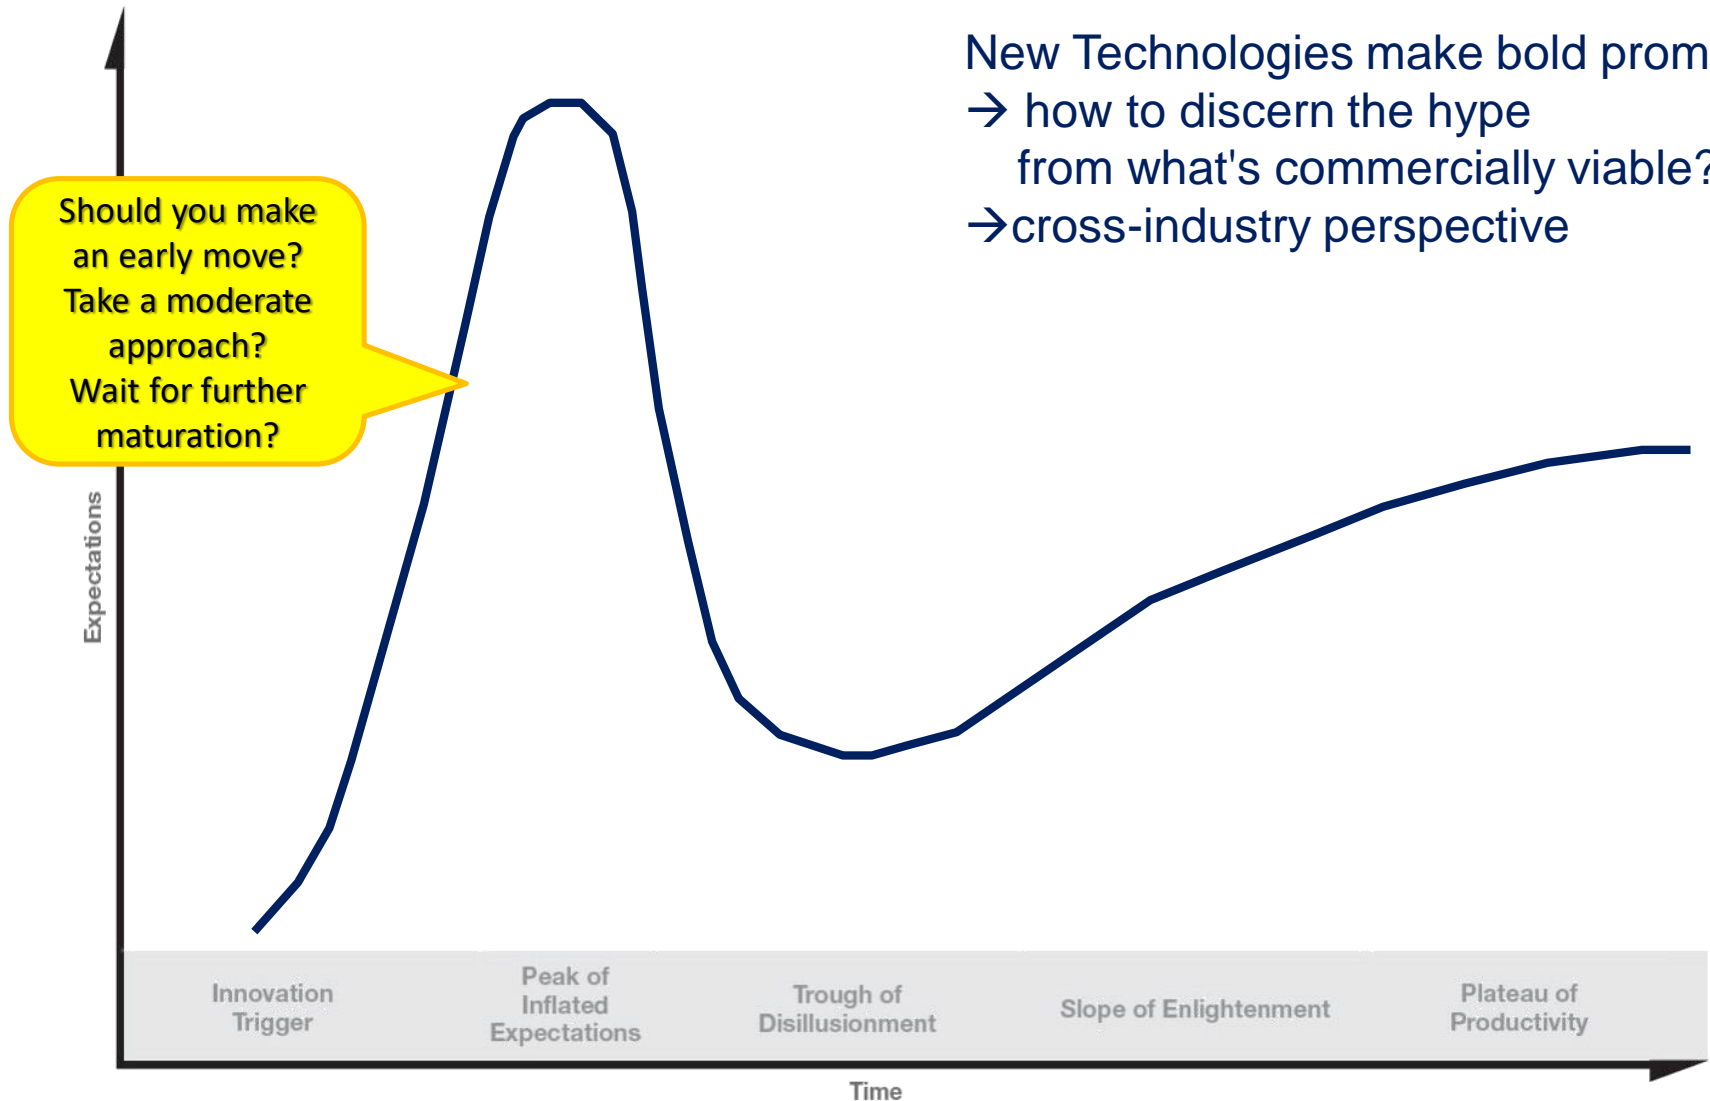
The Hype on Big Data?



The Hype on Forecasting?

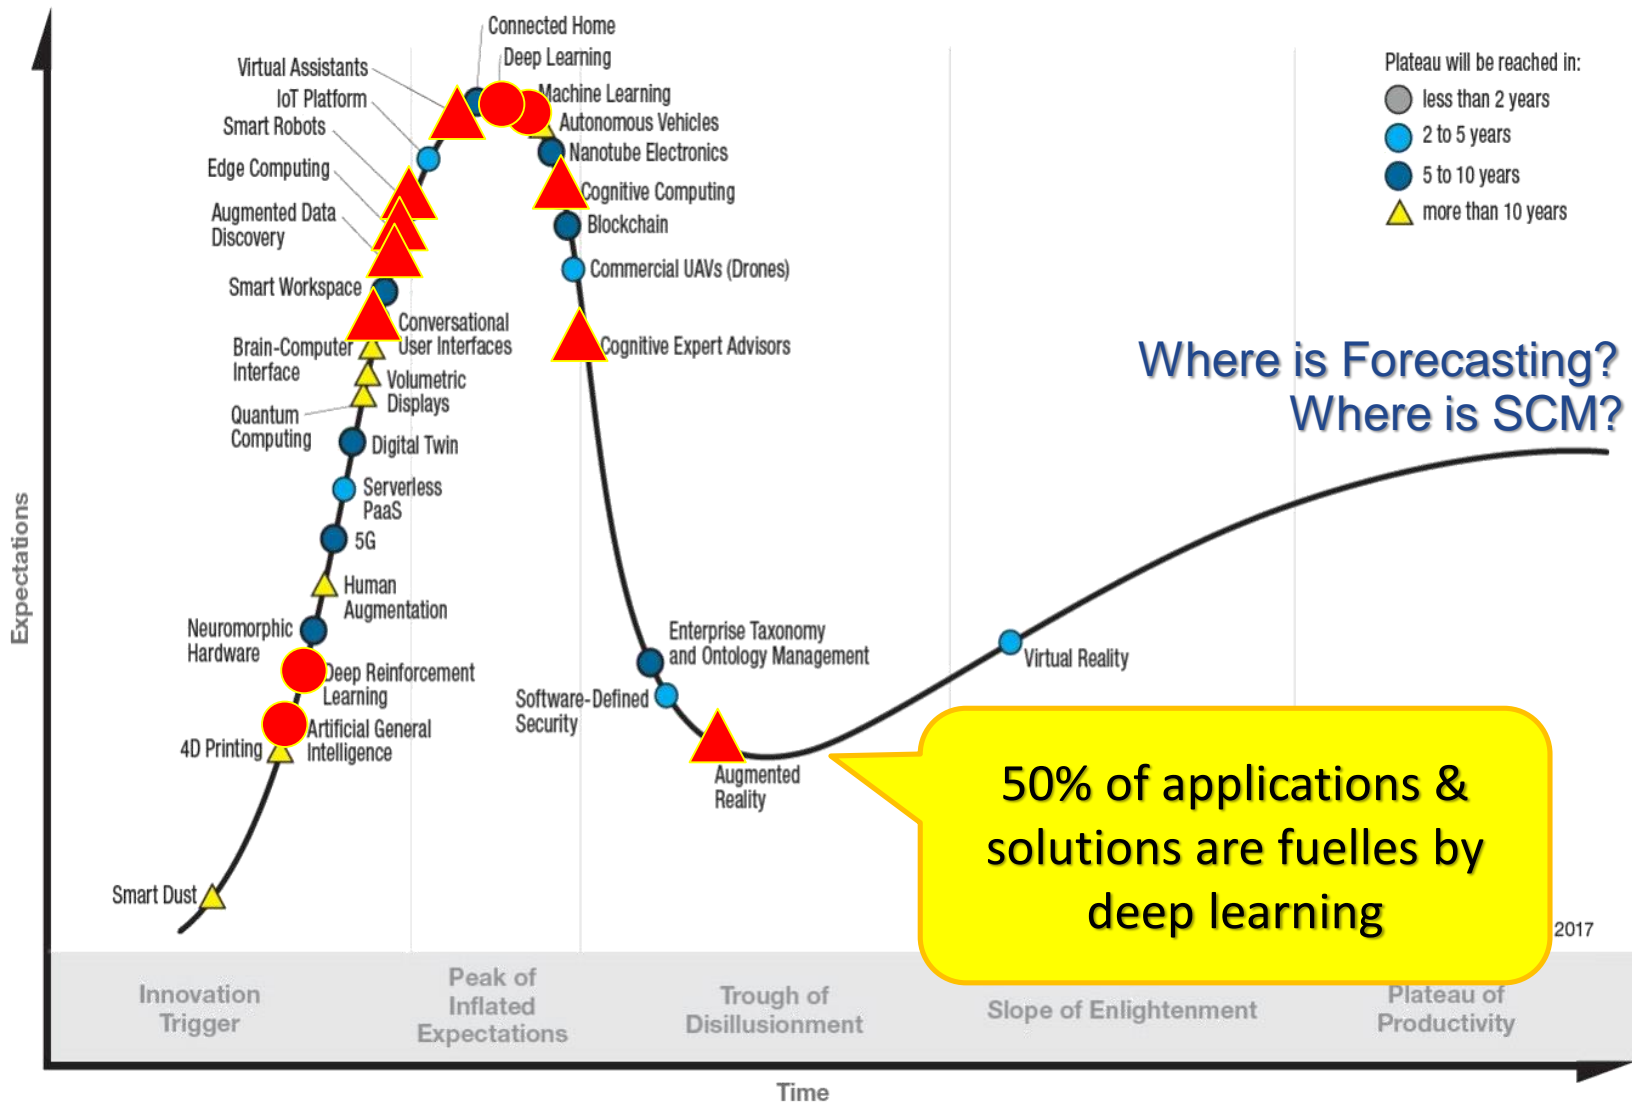


Emerging Technologies Hype Cycle



New Technologies make bold promises
→ how to discern the hype from what's commercially viable?
→ cross-industry perspective

Emerging Technologies Hype Cycle



By JOHN MARKOFF JUNE 25, 2012



An image of a cat that a neural network taught itself to recognize. Jim Wilson/The New York Times

MOUNTAIN VIEW, Calif. — Inside Google's secretive X laboratory, known for inventing self-driving cars and augmented reality glasses, a small group of researchers began working several years ago on a similar human task.



How Computers Can Teach Themselves to Recognize Cats

By Tanya Lewis, Staff Writer | May 28, 2015 07:03am ET

- f
- 🐦
- g+
- 👤
- 🔗



Credit: Michelangelus | Shutterstock.com

In June 2012, a network of 16,000 computers trained itself to recognize a cat by looking at 10 million images from YouTube videos. Today, the technique is used in everything from Google image searches to Facebook's newsfeed algorithms.

The feline recognition feat was accomplished using "deep



an approach to machine learning that works by exposing a

Picking up on the most commonly oc images featured on YouTube, the sys achieved 81.7 percent accuracy in de human faces, 76.7 percent accuracy identifying human body parts and 74 accuracy when identifying cats.

REVIEW

Deep learning

Yann LeCun^{1,2}, Yoshua Bengio³ & Geoffrey Hinton^{1,4}

Deep learning allows computers to learn from data with multiple levels of abstraction, similar to human cognition. Visual object recognition, for example, should change its interpretation in each layer from the previous layer. Deep learning discovers internal features automatically from raw data, such as audio, whereas recurrent neural networks learn to process sequential data, such as text and video.

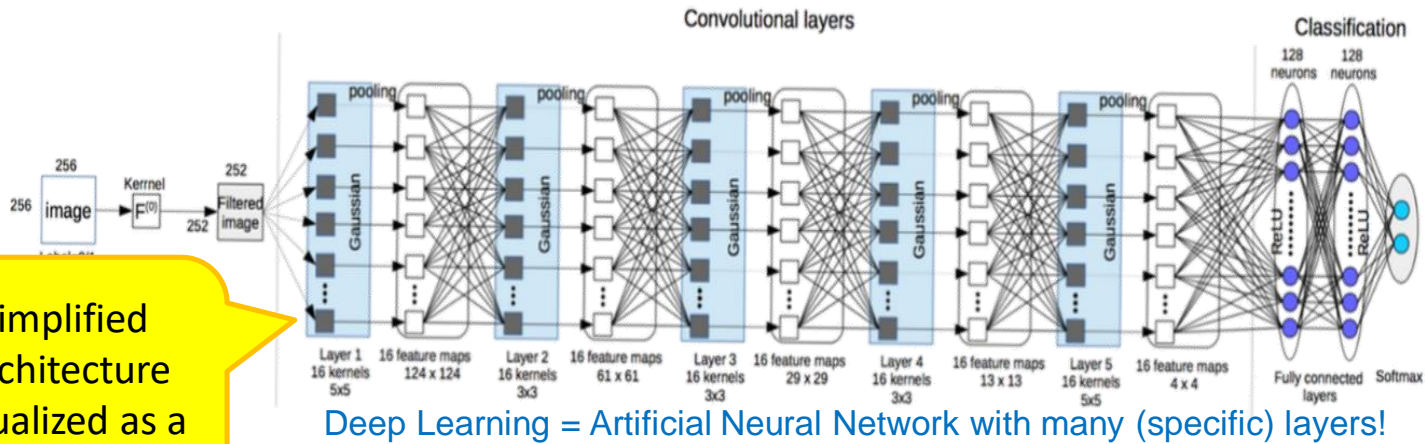
Machine-learning techniques have revolutionized many domains of science, business and government. In addition to their use in recommendation systems, such as those that suggest records in image recognition¹ and speech recognition², it is increasingly present in computer vision, such as in the analysis of satellite images, transcribing spoken words³ and predicting the effects of mutations on gene expression and disease⁴. Perhaps more recently, deep learning has produced extremely promising results in natural language understanding⁵, particularly in tasks such as sentiment analysis, question answering⁶ and language translation⁷.

Source: Nature 521, no. 7555 (2015): 436-444

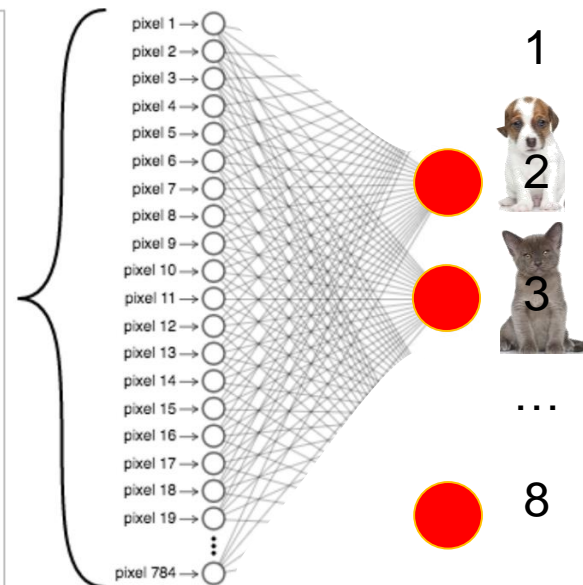


How does a Neural Network „see“?

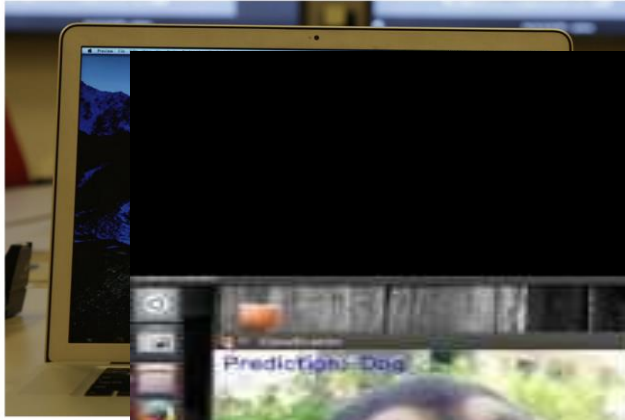
simplified Architecture visualized as a directed graph



28 x 28
= 784 pixels



How Computers Can Teach Themselves to Recognize Cats



An image of a cat that a

MOUNT
for inven
of resear
human'

```

[ 0.02630092]]
[ am 97.37% sure this is a Cat
[ 0.99980155]]
[ am 99.97% sure this is a Dog
[ 0.68480487]]
[ am 99.52% sure this is a Cat
[ 0.00738149]]
[ am 99.27% sure this is a Cat
[ 0.00810282]]
[ am 99.99% sure this is a Cat
[ 0.03820846]]
[ am 85.63% sure this is a Dog
Wrongly classified
[ 0.99999751]]
[ am 100.00% sure this is a Dog

```

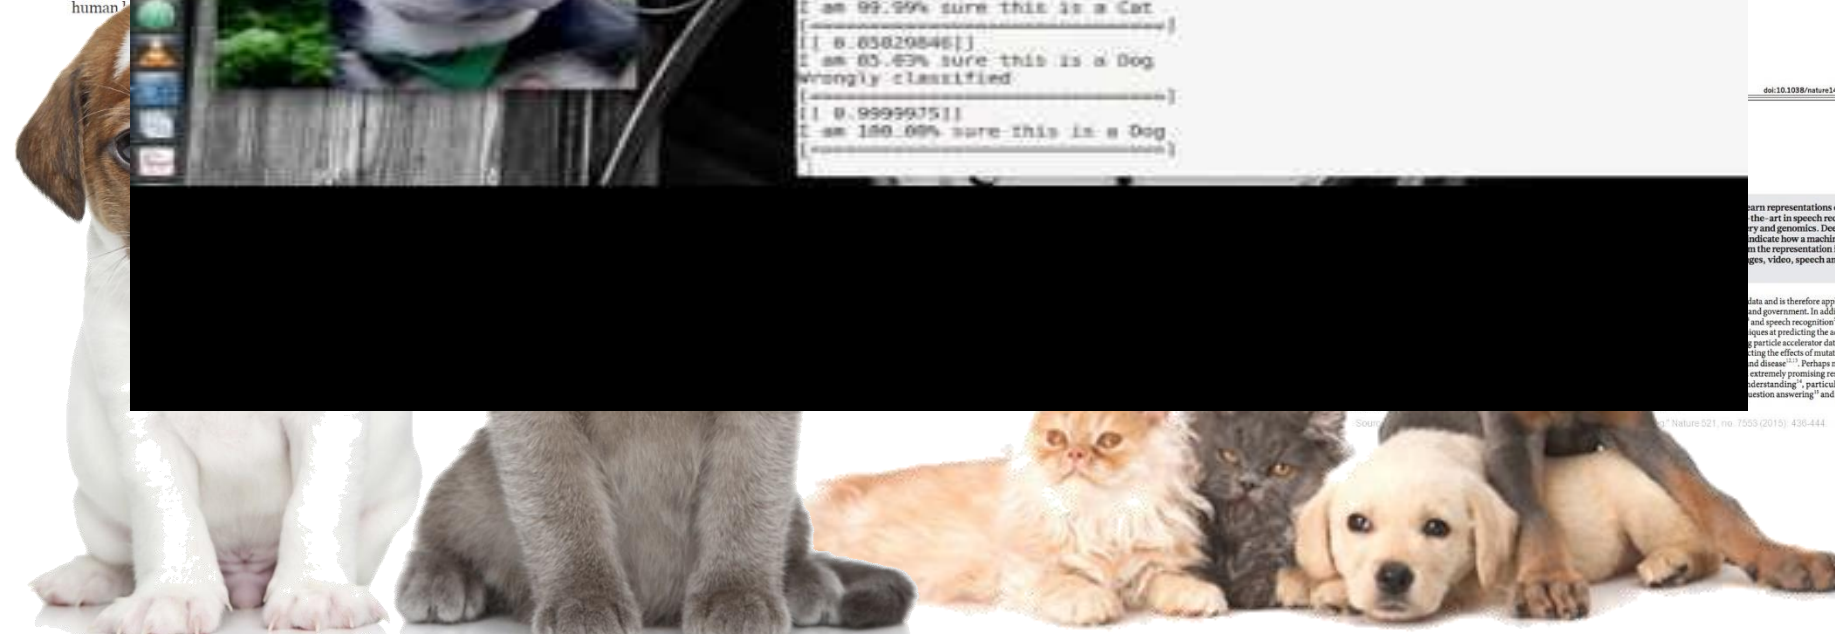
ork of
ained itself
looking at
m
ay, the
everything
earches to
l

n feat was
'deep
sing a

doi:10.1038/nature.14539

arn representations of
the-art in speech recy
and genomics. Deep
indicate how a machine
in the representation in
ges, video, speech and

data and is therefore applica-
and government. In addition
and speech recognition", it
igns at predicting the activi-
particle accelerator data"
ting the effects of mutations
and disease". Perhaps more
extremely promising results
nderstanding", particularly
question answering" and lan-



SUPERHUMAN VISUAL PATTERN RECOGNITION

JÜRGEN SCHMIDHUBER 2013

2011: First Superhuman Visual Pattern Recognition

IJCNN 2011 competition in Silicon Valley:
twice better than humans
three times better than the closest artificial competitor
six times better than the best non-neural method

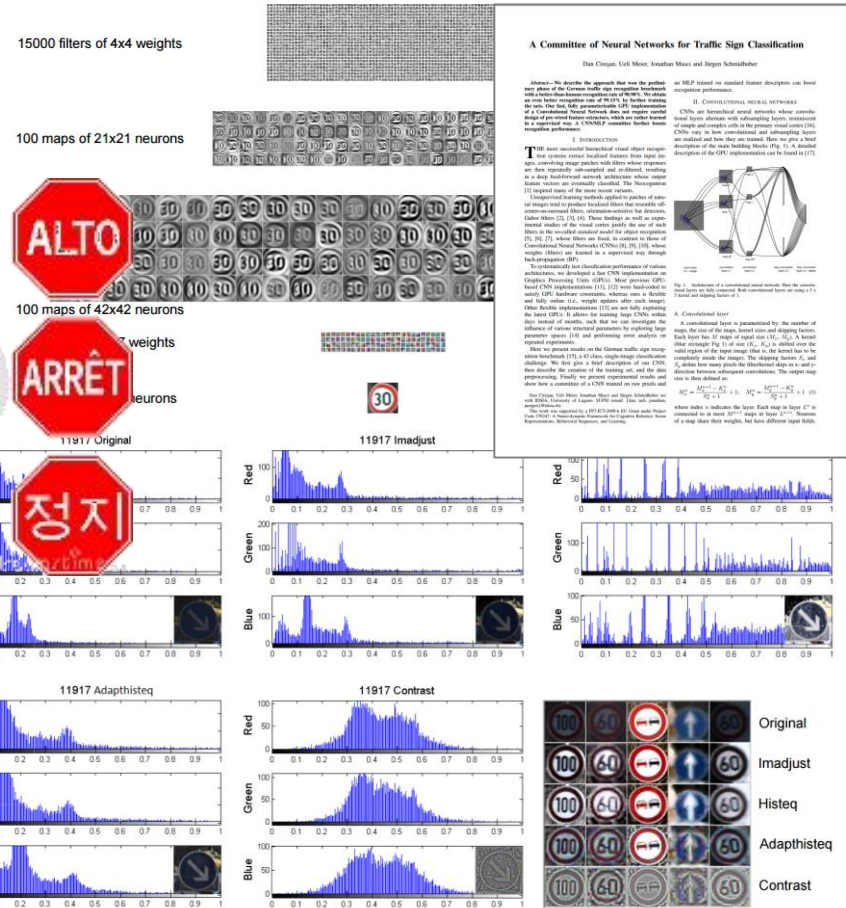
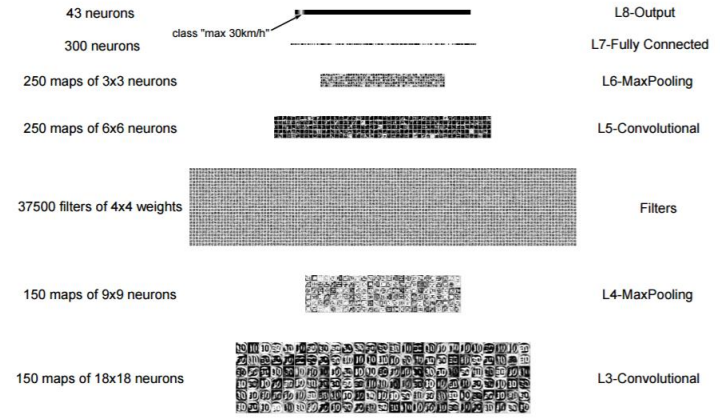
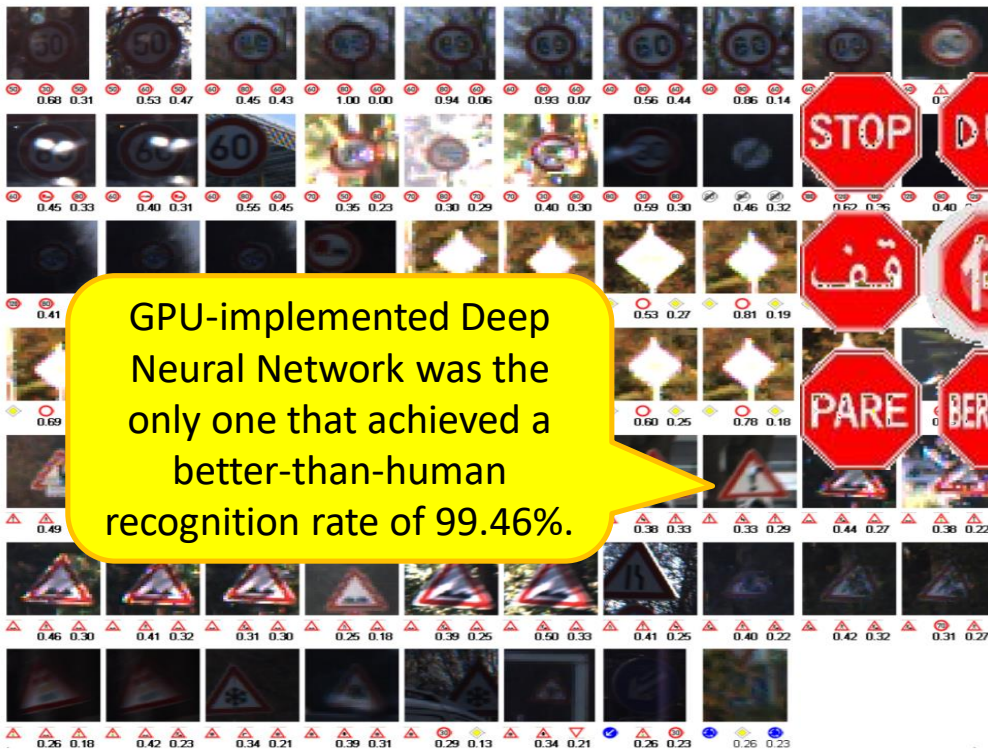
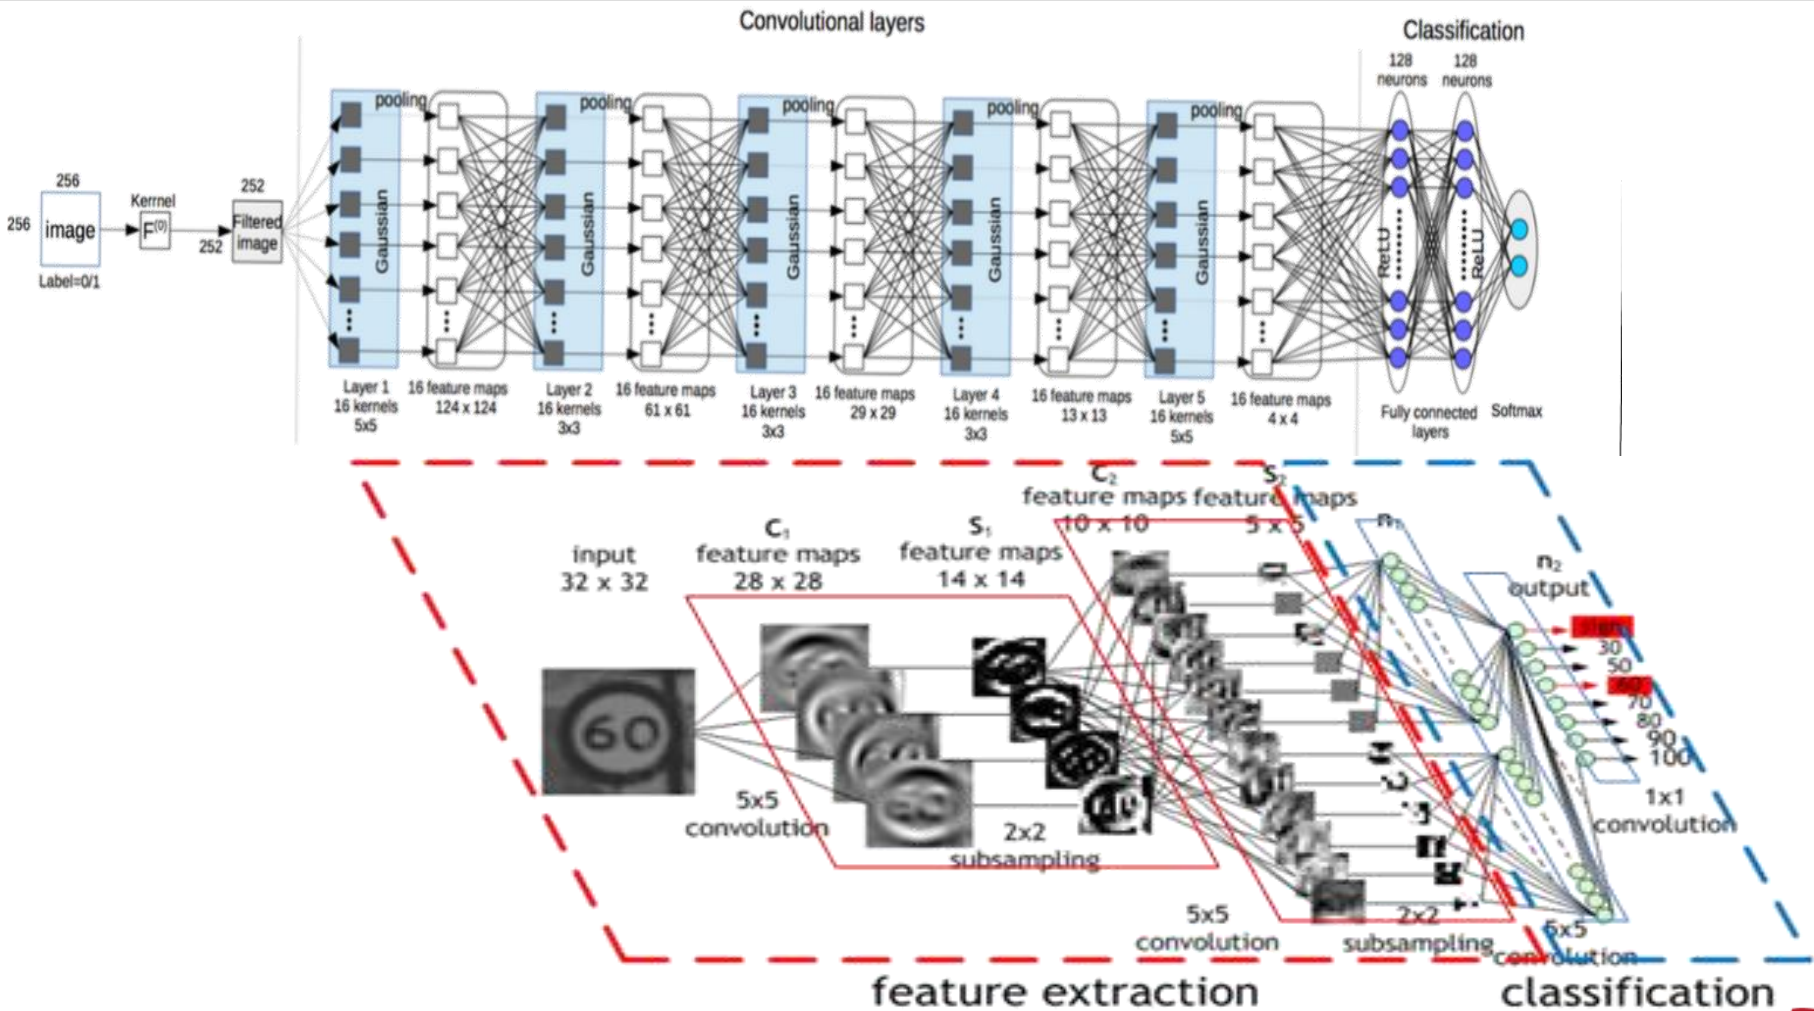



Figure 5: The 68 errors of the MCDNN, with correct label (left) and first (middle) and second best (right) predictions. Best seen in color.

How does a ANN „see“?



Why not before? Needed to overcome vanishing gradient problem
→ parameterisation (Optimisation!) like before with backprop!

SUPERHUMAN VISUAL



Recognition View

Road Sign Recognition

0.747872	0.031060	0.025785	
0.902956	0.012108	0.012063	

Status:

Display Delay : 100 ms
 Image Loading : 3.87 ms
 Processing Time : 167.67 ms

Stop


Normal Environments

NUM

Histeq
 Adapthisteq
 Contrast

Figure 5: The 68 errors of the MCDNN, with correct label (left) and first (middle) and second best (right) predictions. Best seen in color.

AI & Deep Learning in Autonomous cars, Speech Recognition, Image recognition ...



Narrow AI

PERFORMS TASKS THAT NORMALLY REQUIRE HUMAN INTELLIGENCE, BUT CAN ONLY PERFORM TASKS IN A VERY SPECIFIC AND NARROWLY DEFINED DOMAIN.



General AI

"HUMAN-LEVEL AI".
HAS A *GENERAL PROBLEM SOLVING ABILITY* THAT
ENABLES IT TO LEARN NEW TASKS ACROSS SEVERAL
DOMAINS.



... **Superintelligent AI**

Unsupervised Deep Learning in Games

The
Economist

- in 43 out of 49 cases DeepMind outperformed programs designed to play that particular game
- in $\frac{3}{4}$ it defeated professional human players.
- *Super-human* performance ?!

Economist.com

<https://www.youtube.com/watch?v=V1eYniJ0Rnk>

Agenda

AI in Forecasting



- Capabilities of AI
 - The hype
 - How AI can “see”
 - AI Demos
- **Using AI in Forecasting**
 - AI in Forecasting
 - Case studies
- **The Gap**
 - Forecasting?
 - State of Forecasting
 - Sizing the Gap

What is Forecasting?

This is a time series ...

(a collection of observations made sequentially in time)

Jan 2016	91.36
Feb 2016	100.77
Mar 2016	87.85
Apr 2016	88.86
May 2016	99.93
Jun 2016	115.32
July 2016	92.30
Aug 2016	103.71
Sep 2016	97.74
Oct 2016	111.17
...	...
Dec 2016	89.26





... how will it continue?

”Estimation how the sequence of observations observed in the past will continue into the future”









[Makridakis, Wheelwright, Hyndman, 1998]

What is Forecasting?



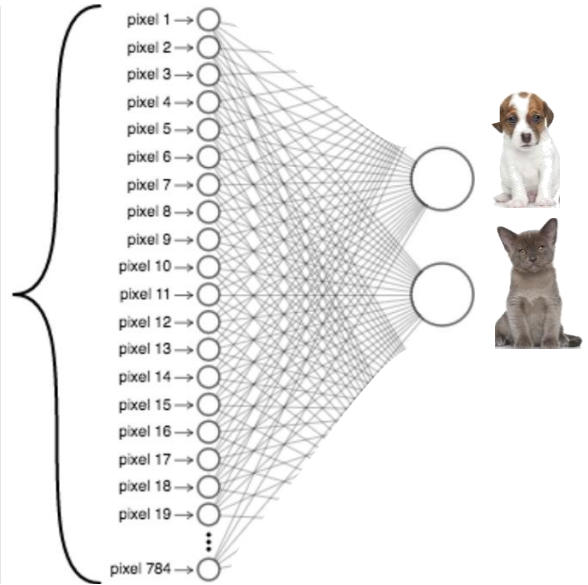
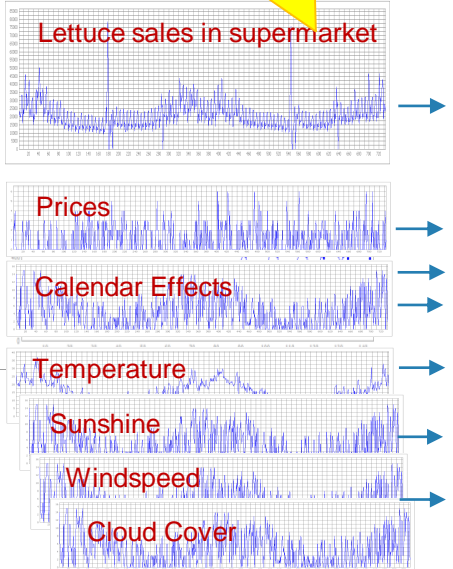
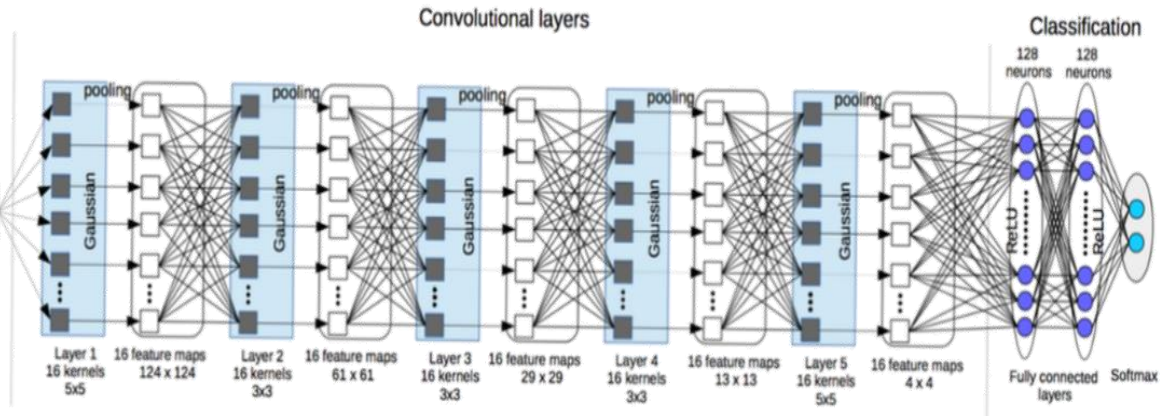
Why Forecast with AI?

Oh, pls repeat
for 40000 items
every day
in 1000 locations
...

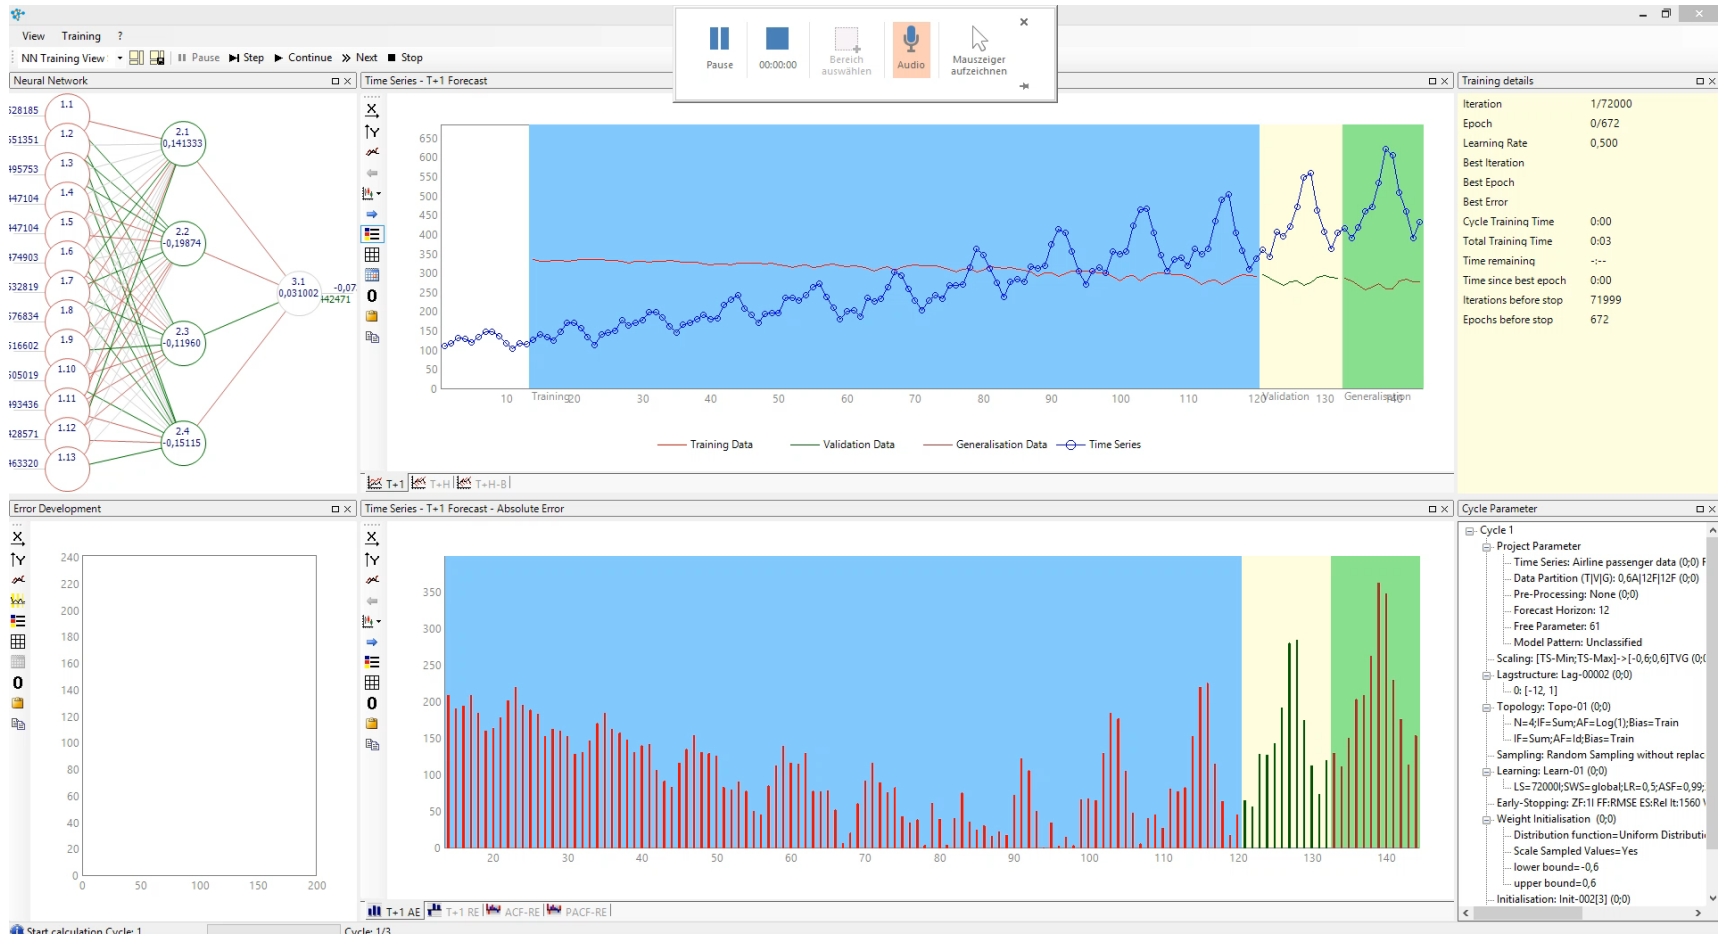


How does a Neural Network „see“?

Needs compute power to repeat for 40000 items in 1000 locations daily ...




Neural Network Demo



FMCG Forecasting for manufacturing @

Time Series Horizon + Forecast-Confidence



Median sMAPE	Test
Seasonal Linear Regression (35B)	18.20
SLR + Judgment	18.02
MLP AR, SinCos + Selection	9.00
Improvement	-9.02
Improvement in %	-49.83%

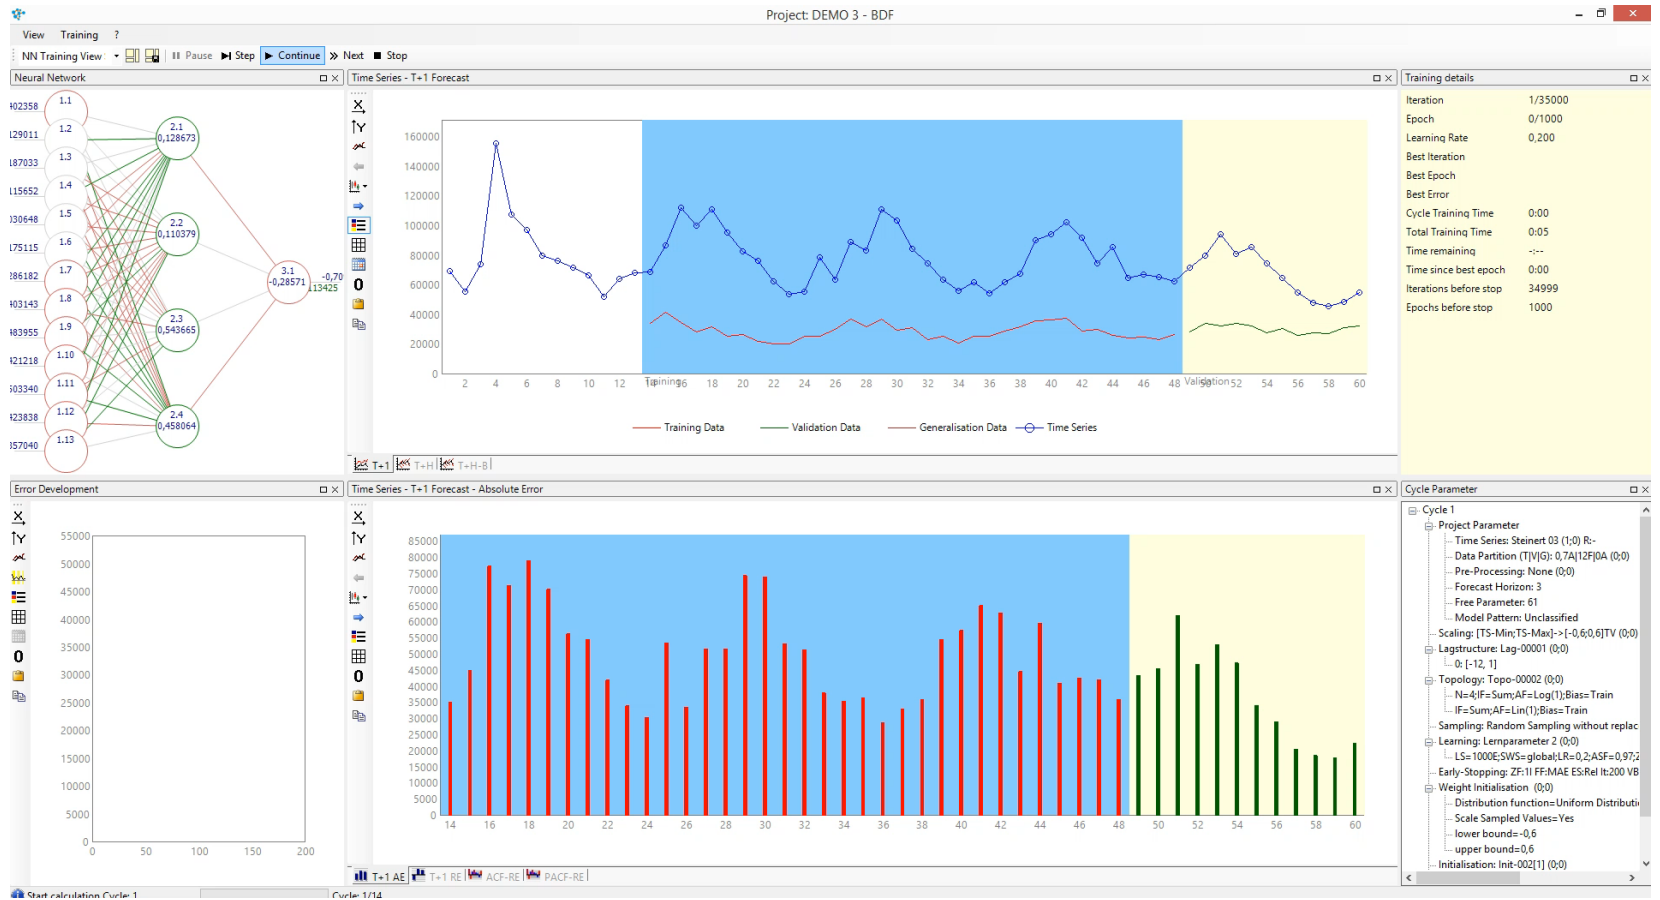
NNET achieves super-human forecast accuracy

sMAPE test	SAP error	iqast error	Δ error	Δ error %	# items
Canada	40,7	33,8	-6,9	-16,9%	47
Germany	55,4	51,7	-3,7	-6.8%	155
France	43,7	42,6	-1,2	-2.7%	262
Greece	50,0	48,4	-1,7	-3.3%	196
Italy	47,0	44,7	-2,3	-4.9%	175
Netherlands	47,0	45,5	-1,5	-3.2%	154
Poland	47,0	39,8	-7,2	-15.3%	78
South Africa	37,3	35,9	-1,4	-3.7%	36
Average				-7.5%	

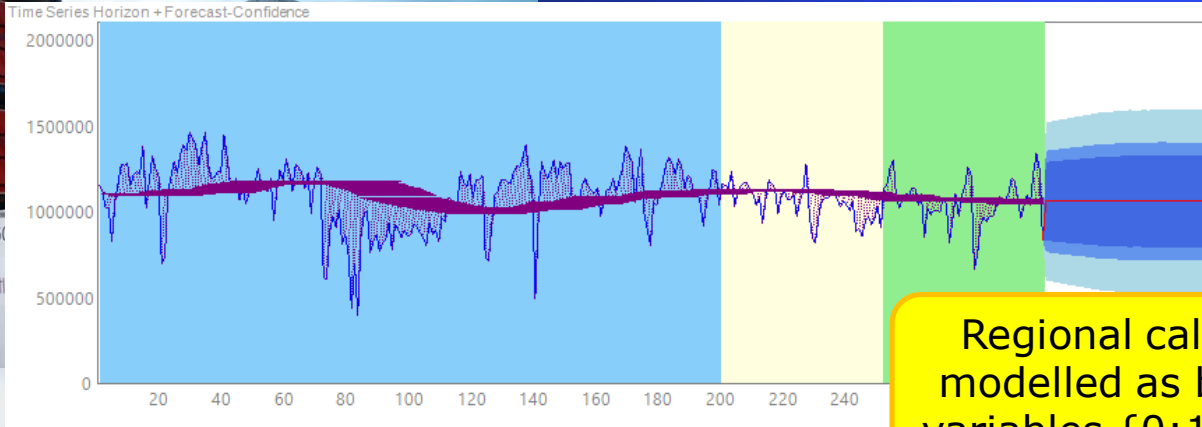
NNET achieves super-human selection accuracy




Neural Network Demo



Container forecasts @














Regional calendar events modelled as binary dummy variables {0;1} in occurrence



Container forecasts @ virgin atlantic cargo

Expand Models with more variables?

- Internal factors, eg. Capacity (different hullspace)
- External factors, eg. GDP growth, industry leading indicators, Global PMI, OECD Confidence index etc.



Beer Forecasting for manufacturing @



Beer Forecasting for manufacturing @



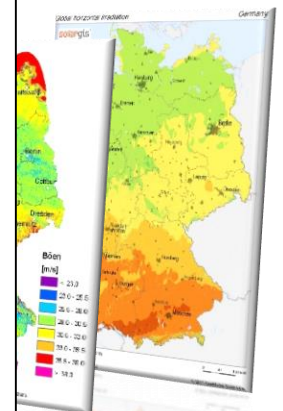
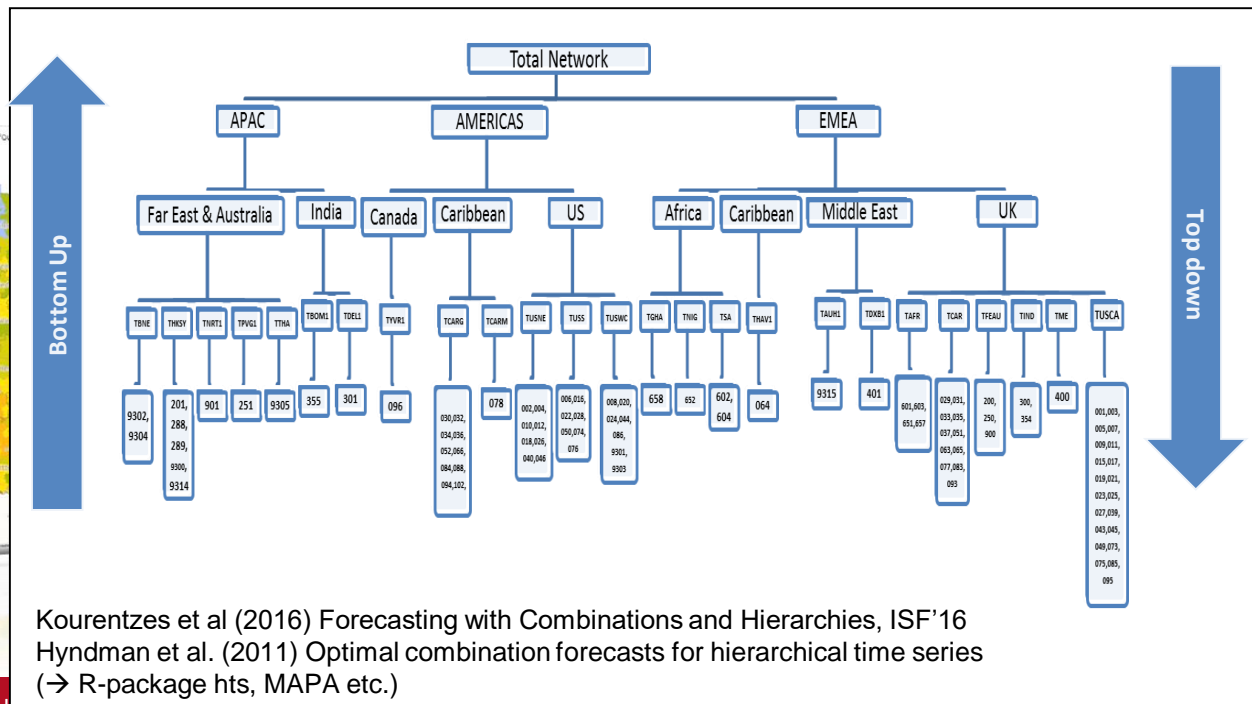
Drivers for Consumer Demand?

- Weather
- Price & Price promotions
- Marketing
- Politics & regulation
- Changes in Retail structures
- Manufacturer consolidation
- Regional Preferences
- etc.

+5 - 8%
 +10 - 700%
 +2 - 3%
 -0 - 100%



- **only few are regional**
- **these have lesser impact!**



Beer Forecasting for manufacturing @



Facebook
Heute ist Gyde Schmidts Geburtstag.
(98 weitere neue Benachrichtigungen)
www.facebook.com







Beer Forecasting for manufacturing @



Forecast from MLP including bank holidays, sales price changes & weather & football events

Time Series Horizon + Forecast-Confidence



Forecast Algorithm	Test sMAPE Error
Naive Level	23.91%
Naive Season	21.84%
Single Exponential Smoothing	
Seasonal ARIMA	
Additive Exponential Smoothing	
Dynamic Regression	13.37%
Neural Networks	13.01%
Human Expert Demand Planner	14.42%

achieves „super-human“ forecast accuracy




Neural Nets in Retail & Promos

Fit appropriate forecasting models to:

- Predict future value
- Understand effect (elasticity) of each exogenous factor → affect forecasts


Time Series: 5 | Product: 1080215 (Trend: 0, Season13: 0, Season52: 1)

$$f(X) = +125.37 * \text{Constant} + 0.71 * \text{Lag1} + 0.07 * \text{Lag2} + 610.76 * \text{Adcode: 401} + 617.24 * \text{Discount} - 1569.46 * \text{Xmas} - 696.48 * \text{Xmas+1} + 1209.29 * \text{Xmas-2} - 1581.59 * \text{Easter-1} - 894.91 * \text{Labour}$$



CIF 2016

International Time Series Forecasting Competition Computational Intelligence in Forecasting




#	Method		SMAPE	SD #
1	LSTM-Deseasonalized	Microsoft	0.105 +- 0.107	1
2	LSTMs and ETS	Microsoft	0.108 +- 0.116	2
3	*ETS		0.119 +- 0.142	9
4	MLP	iqast	0.121 +- 0.135	7
5	REST		0.124 +- 0.133	6
6	*FRBE		0.129 +- 0.162	13
7	HEM		0.130 +- 0.147	11
8	*Avg		0.131 +- 0.133	5
9	*BaggedETS		0.131 +- 0.176	17
10	LSTM	Microsoft	0.133 +- 0.155	12
11	Fuzzy c-regression		0.137 +- 0.127	4
12	PB-GRNN		0.145 +- 0.166	14
13	PB-RF		0.145 +- 0.166	15
14	*ARIMA		0.146 +- 0.218	21
15	*RW		0.146 +- 0.137	8
16	*Theta		0.148 +- 0.122	3
17	PB-MLP		0.149 +- 0.172	16
18	TSFIS		0.151 +- 0.147	10
19	*Boot.EXPOS		0.153 +- 0.206	20
20	MTSFA		0.165 +- 0.180	18
21	FCDNN		0.166 +- 0.194	19
22	MSAKAF		0.204 +- 0.225	22
23	HFM		0.224 +- 0.251	23
24	CORN		0.288 +- 0.263	24






pioneers in forecasting with artificial intelligence




New AI forecasting engine for SAP APO DP

- intelligent model selection
- intelligent forecast algorithms
- intelligent safety stocks



How to get started?



NDA

Demo on
your data
(probono)

POC
on 2-3 data
subsets

Project / Pilot study
with contract

RFI for
software

1-2 man days
skype call results

need historical sales
explanatory variables
(price, promo, ...)

5 man-days analysis
Kick-off on site
Final results on site
<€10k

10-20 man-days analysis
scalability to other countries
product categories



Take aways

- AI is hyped, but here to stay
- AI (aka Deep Learning Neural Networks) is a driver for Digitalisation and Industry 4.0
- Forecast accuracy is a core obstacle in SCM
- AI is largely ignored in Forecasting for SCM
- Innovators are leading the way
 - High opportunities from low-cost pilot studies
 - Try new algorithms!
 - Neural Networks
 - Support Vector Regression
 - Decision Trees
 - K-Nearest Neighbours
 - ...



© 2017 Sven F. Crone. ALL RIGHTS RESERVED

Unless otherwise indicated, all material in this presentation is copyrighted by Sven F. Crone. No part of these pages, either text or images may be used for any purpose other than personal use, unless explicit authorisation is given by the author. Reproduction, modification, storage in a retrieval system or retransmission, in any form or by any means – electronic or otherwise, for reasons other than personal use, is strictly prohibited without prior permission.

This restriction explicitly includes the use of this training material for company internal or external trainings, train the trainer setups or any other training activities other than a one-to-one personal job handover.

1. The content of the pages of this presentation is for your general information and use only. It is subject to change without notice.
2. Neither we nor any third parties provide any warranty or guarantee as to the accuracy, timeliness, performance, completeness or suitability of the information and materials found or offered in this presentation for any particular purpose. You acknowledge that such information and materials may contain inaccuracies or errors and we expressly exclude liability for any such inaccuracies or errors to the fullest extent permitted by law.
3. Your use of any information or materials in this presentation is entirely at your own risk, for which we shall not be liable. It shall be your own responsibility to ensure that any products, services or information identified in this presentation meet your specific requirements.
4. This presentation contains material which is owned by or licensed to us. This material includes, but is not limited to, the design, layout, look, appearance and graphics. Reproduction is prohibited other than in accordance with the copyright notice, which forms part of these terms and conditions.
5. All trade marks reproduced in this website which are not the property of, or licensed to, the operator are acknowledged in the presentation.
6. Unauthorised use of this presentation may give rise to a claim for damages and/or be a criminal offence.
7. From time to time this presentation may also include links to other presentations or websites. These links are provided for your convenience to provide further information. They do not signify that we endorse the other presentations or website(s). We have no responsibility for the content of the linked presentations or website(s).
8. Your use of this presentation and any dispute arising out of such use of the presentation is subject to the laws of Hamburg, Germany.
9. No images rendered can be used for mechanical, electronic or printed reproduction without prior written permission by the author. Images generated are for previewing purposes only.

