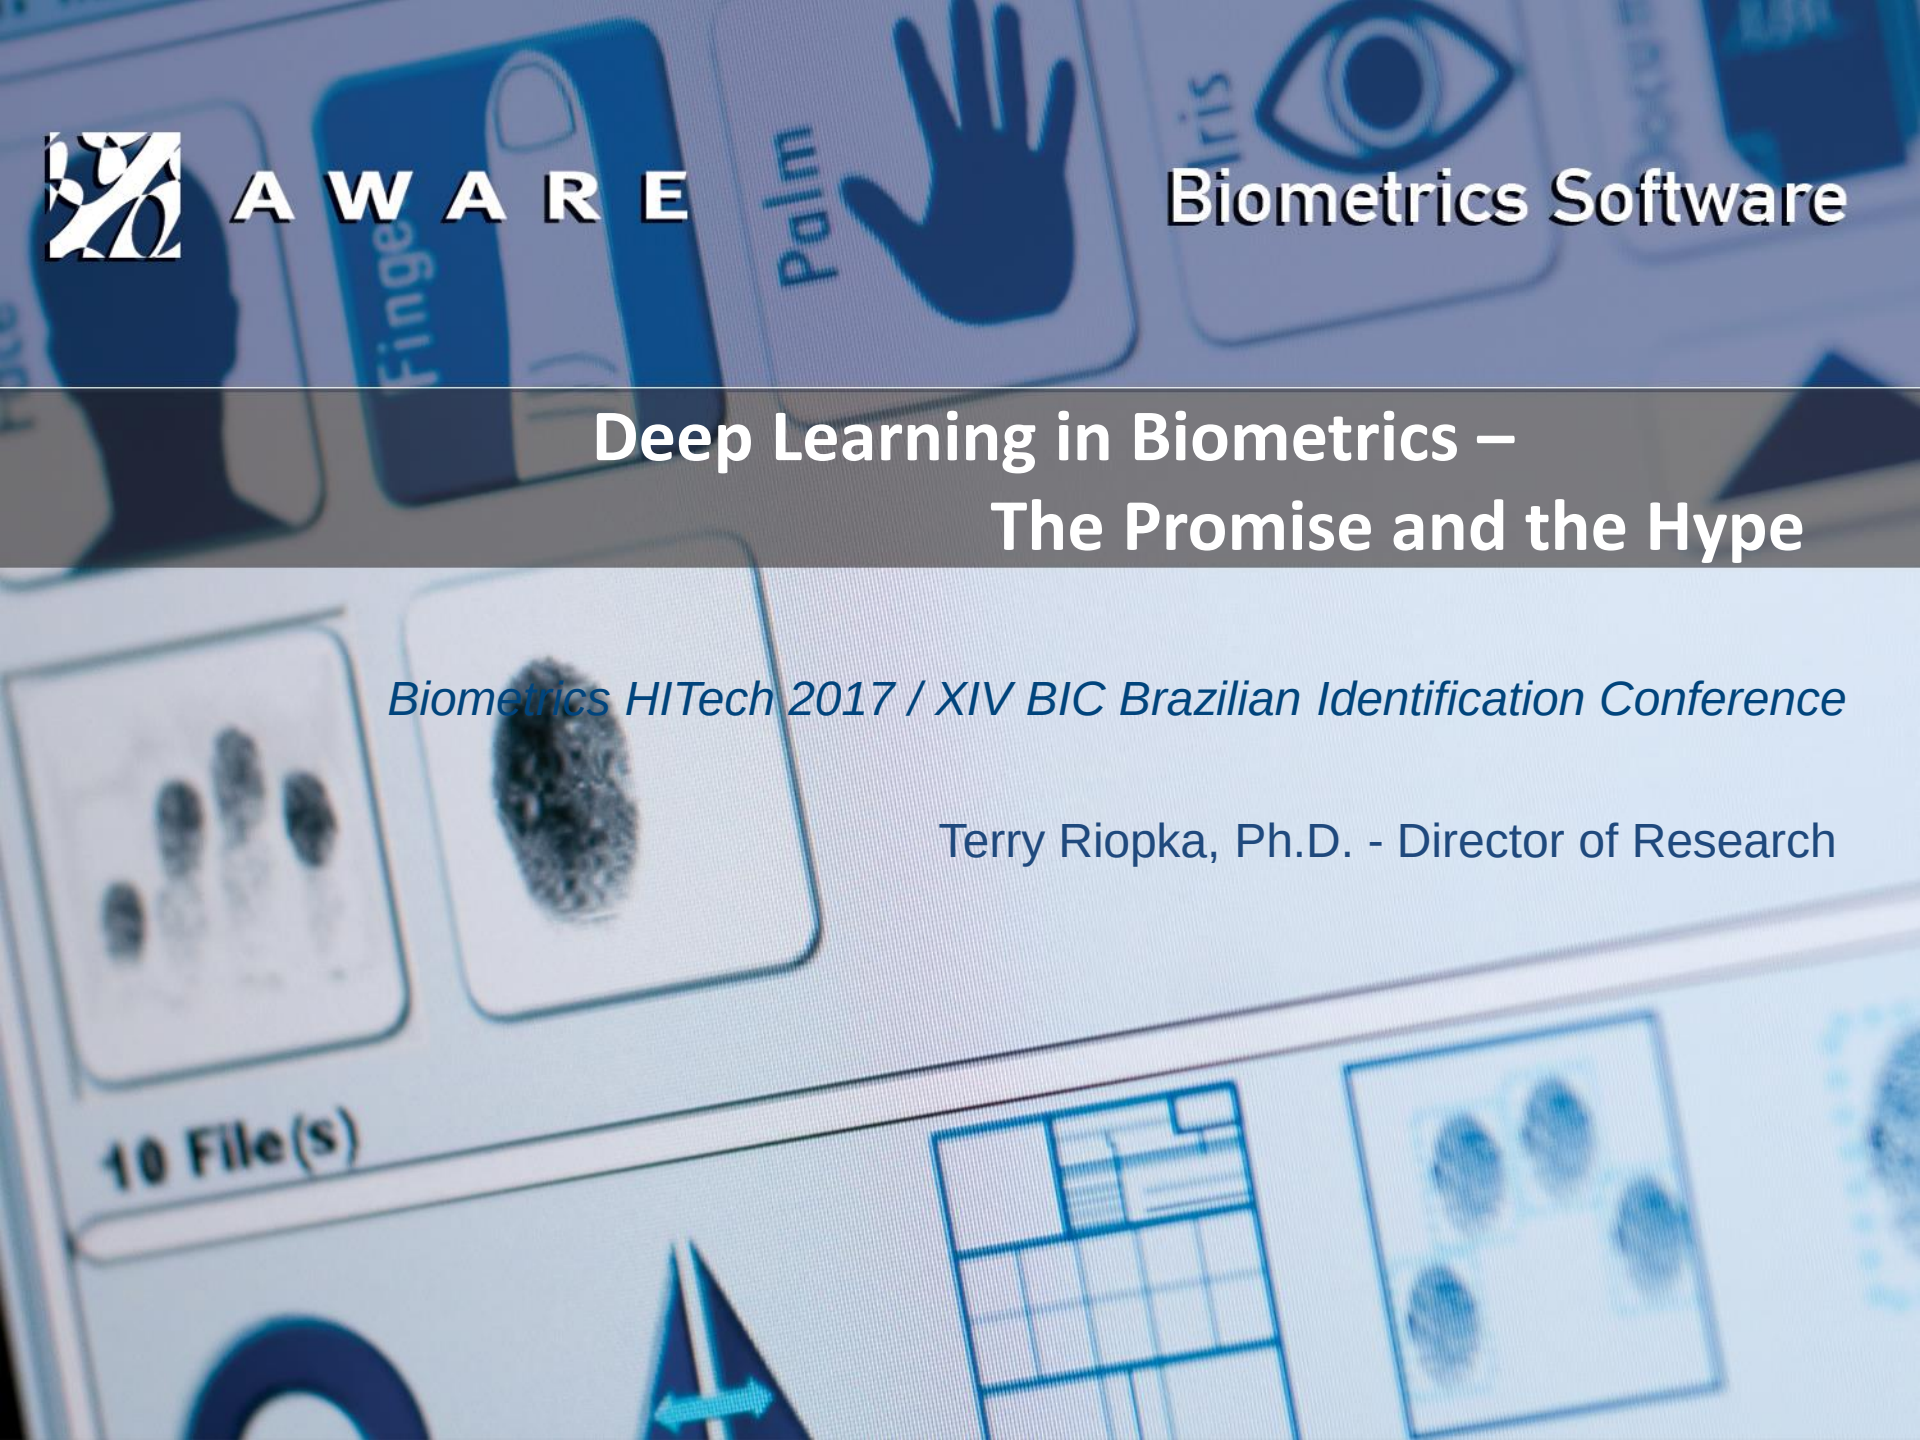

Deep Learning in Biometrics – The Promise and the Hype

Biometrics HITech 2017 / XIV BIC Brazilian Identification Conference

Terry Riopka, Ph.D. - Director of Research

Overview

- Traditional Neural Nets: Backpropagation
- From Backpropagation to Deep Learning
- Is Deep Learning Useful for All Biometrics?
- Is Deep Learning Really Something New?
- Challenges and Weaknesses of Deep Learning

©Terry Riopka - Aware, Inc. - 2017

Traditional Neural Net - Backpropagation

Example: Number Recognition

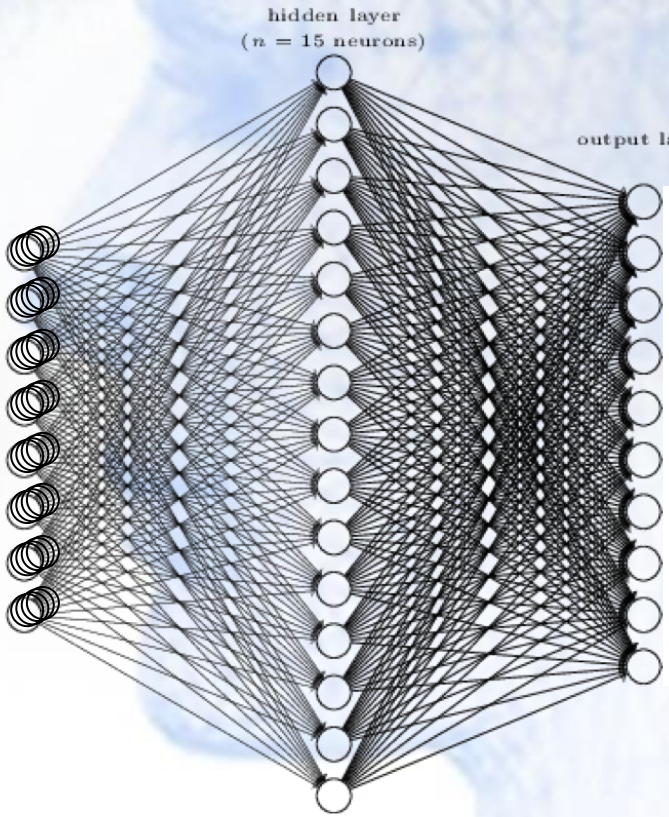

©Terry Riopka - Aware, Inc. - 2017

Traditional Neural Net - Backpropagation

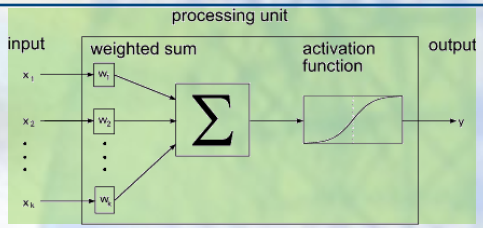

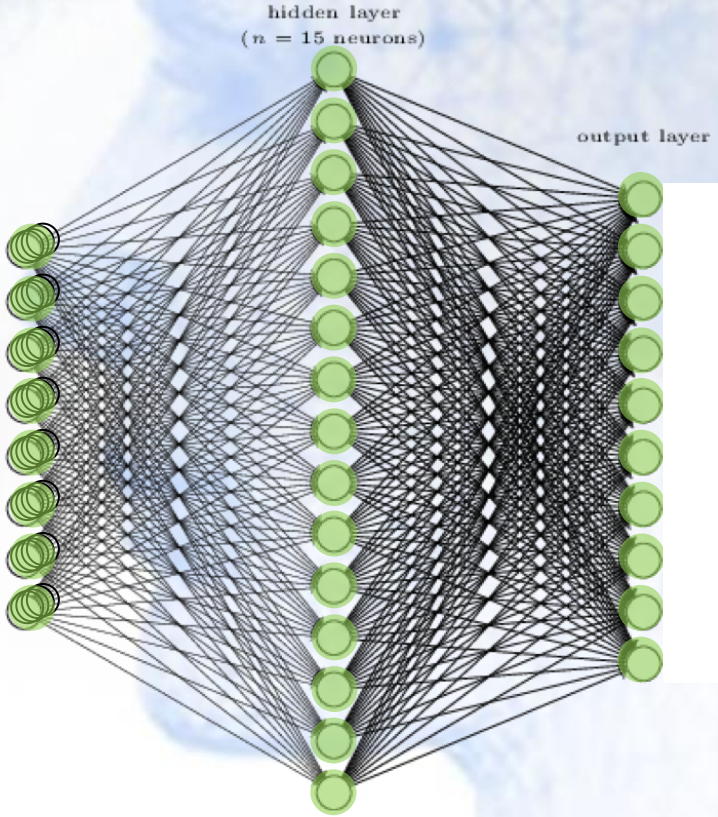

©Terry Riopka - Aware, Inc. - 2017

Traditional Neural Net - Backpropagation

Example: Number Recognition

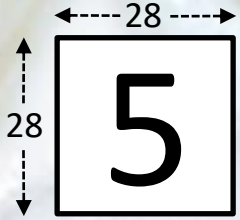

28x28 = 784
Input Pixel
Values

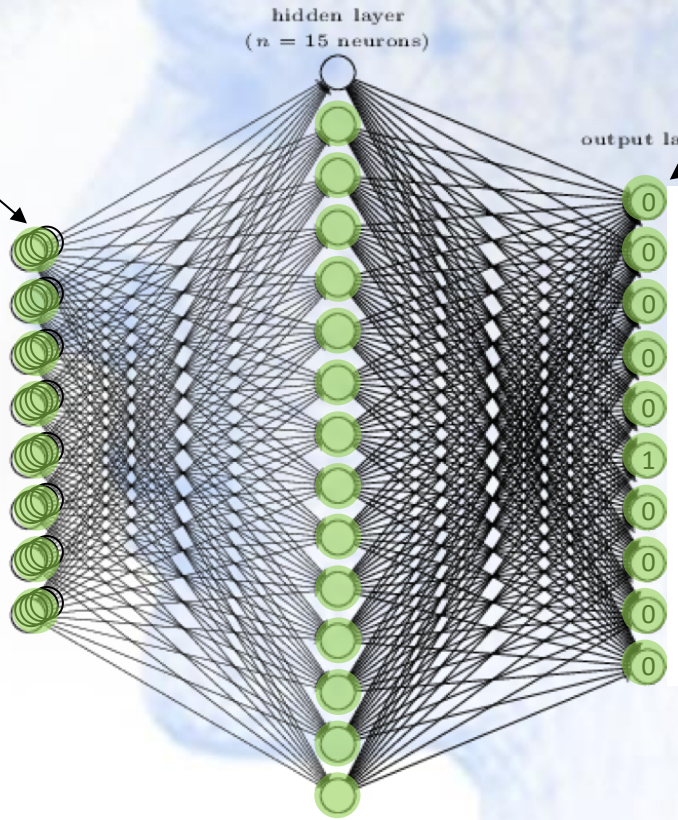

Train for
this output

©Terry Riopka - Aware, Inc. - 2017

Traditional Neural Net - Backpropagation

Example: Number Recognition

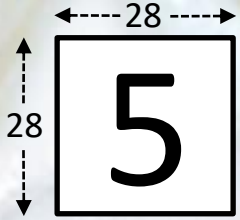

28x28 = 784
Input Pixel
Values

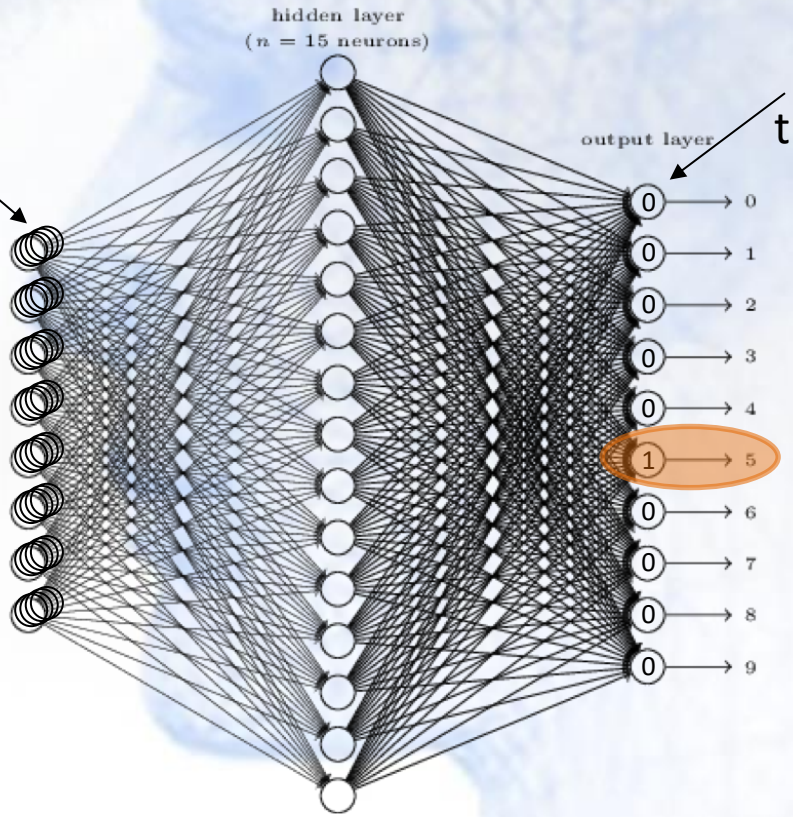

Train for
this output

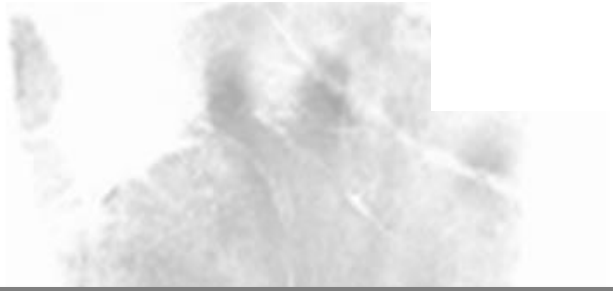

©Terry Riopka - Aware, Inc. - 2017

Traditional Neural Net - Backpropagation

Example: Number Recognition

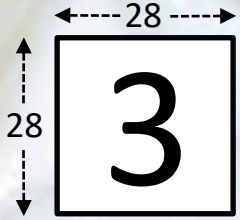

28x28 = 784
Input Pixel
Values

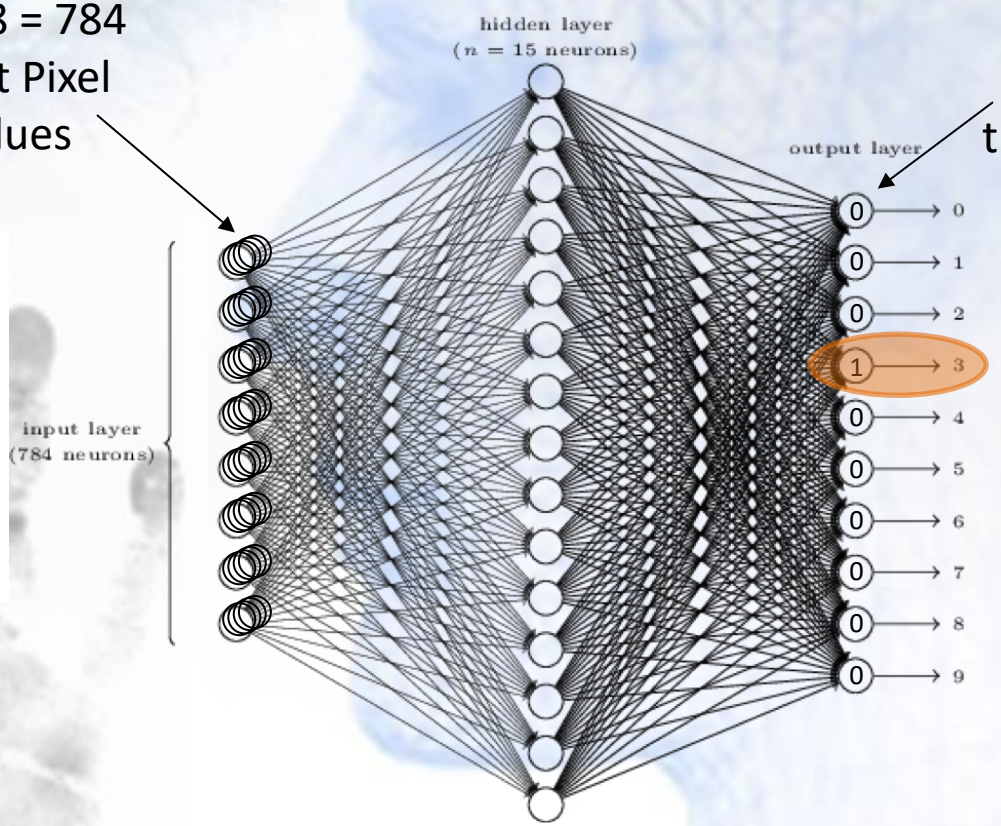

Train for
this output

©Terry Riopka - Aware, Inc. - 2017

AWARE

Traditional Neural Net - Backpropagation

Example: Number Recognition

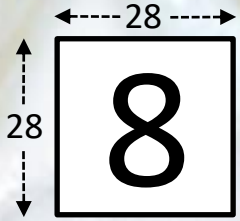

28x28 = 784
Input Pixel
Values

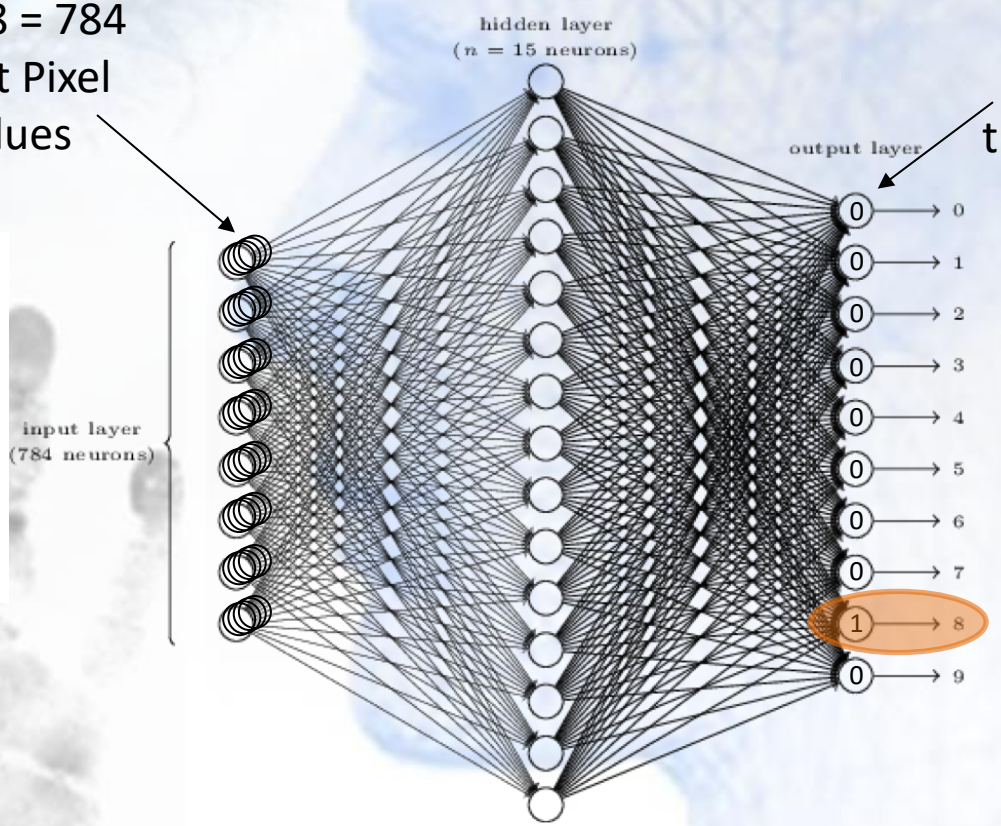

Train for
this output

©Terry Riopka - Aware, Inc. - 2017

AWARE

Traditional Neural Net - Backpropagation

Example: Number Recognition

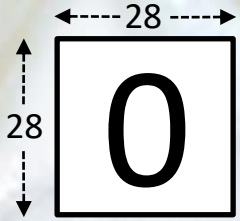

28x28 = 784
Input Pixel
Values

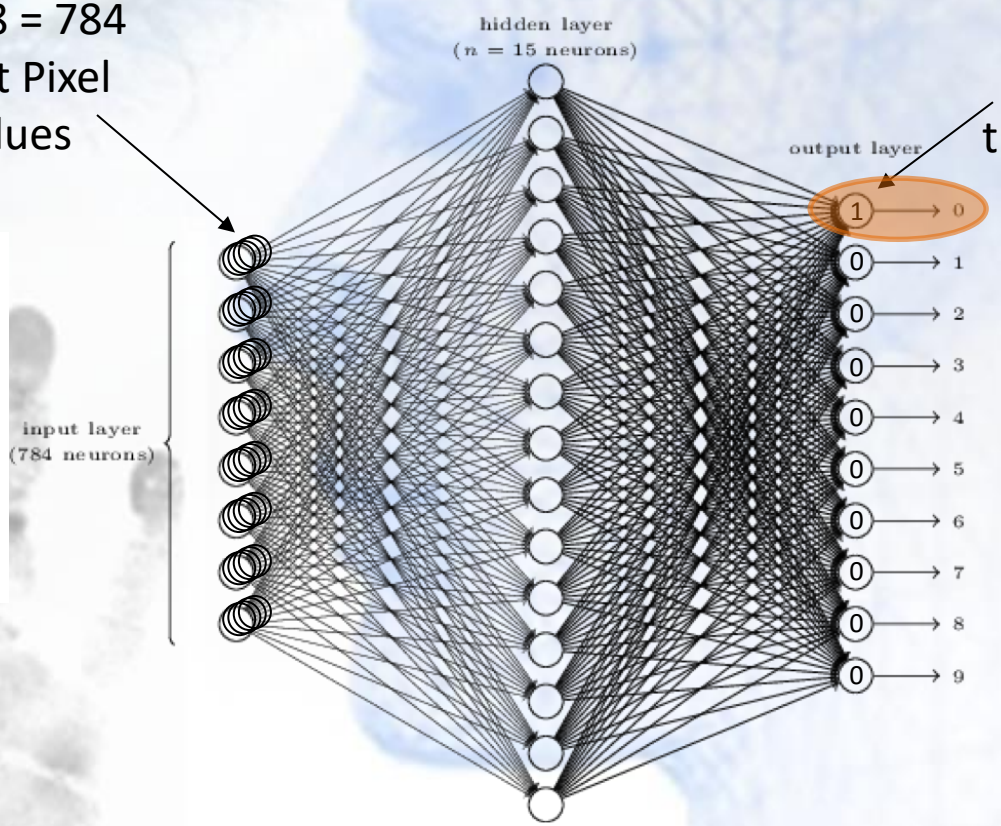

Train for
this output

ETC.

©Terry Bianka, Aware, Inc., 2017

Traditional Neural Net - Backpropagation

Example: Number Recognition

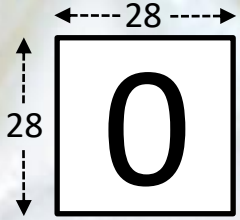

28x28 = 784 Input Pixel Values

ETC.

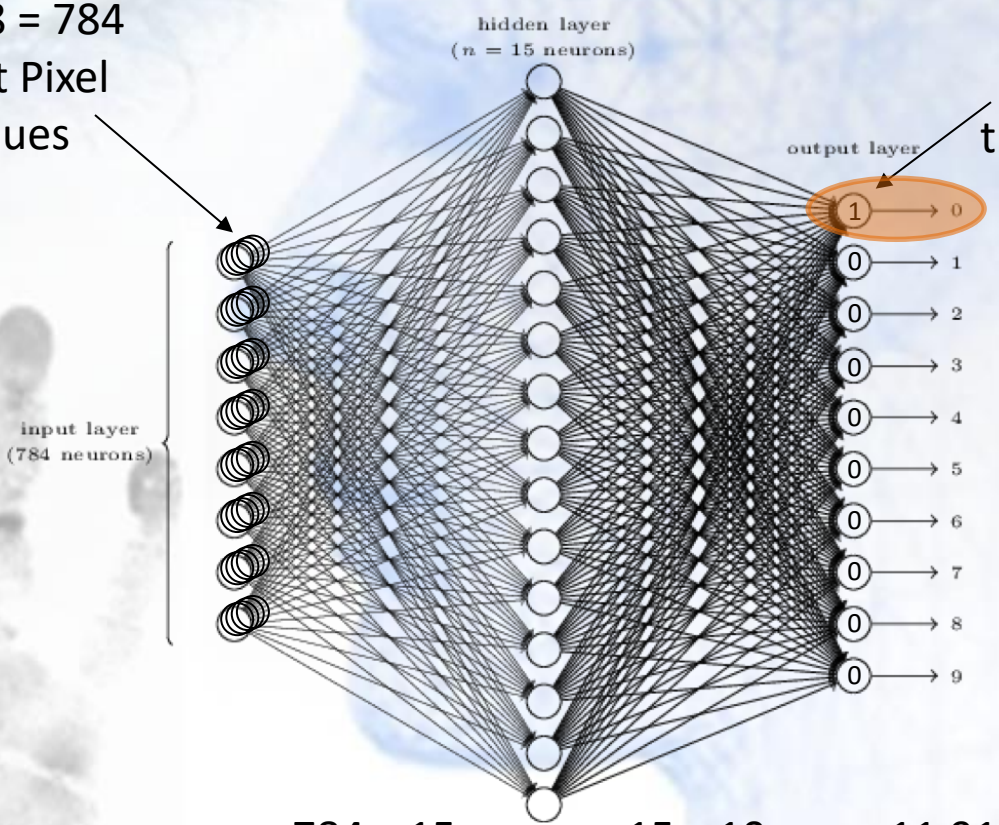

Train for this output

Number of adjustable (trainable) parameters: 784 x 15 + 15 x 10 = 11,910

©Aware, Inc. 2017

Traditional Neural Net - Backpropagation

Example: Naïve Face Recognition (1980s)

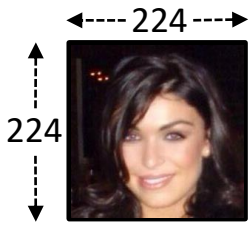

224x224 = 50,176 Input Pixel Values

input layer (50,176 neurons)

hidden layer (n = 1000 neurons)

output layer

Train for this output

1 0 0 0 0 0 0 0 0 0

Subject ID: 1011000000

Number of adjustable (trainable) parameters: 50,176 x 1000 + 1000 x 10 = 50,186,000 !

From Backpropagation to Deep Learning

Sophisticated Non-Linear Mapping Function !

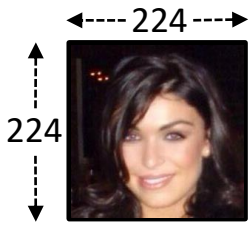

224x224 = 50,176 Input Pixel Values

input layer (50,176 neurons)

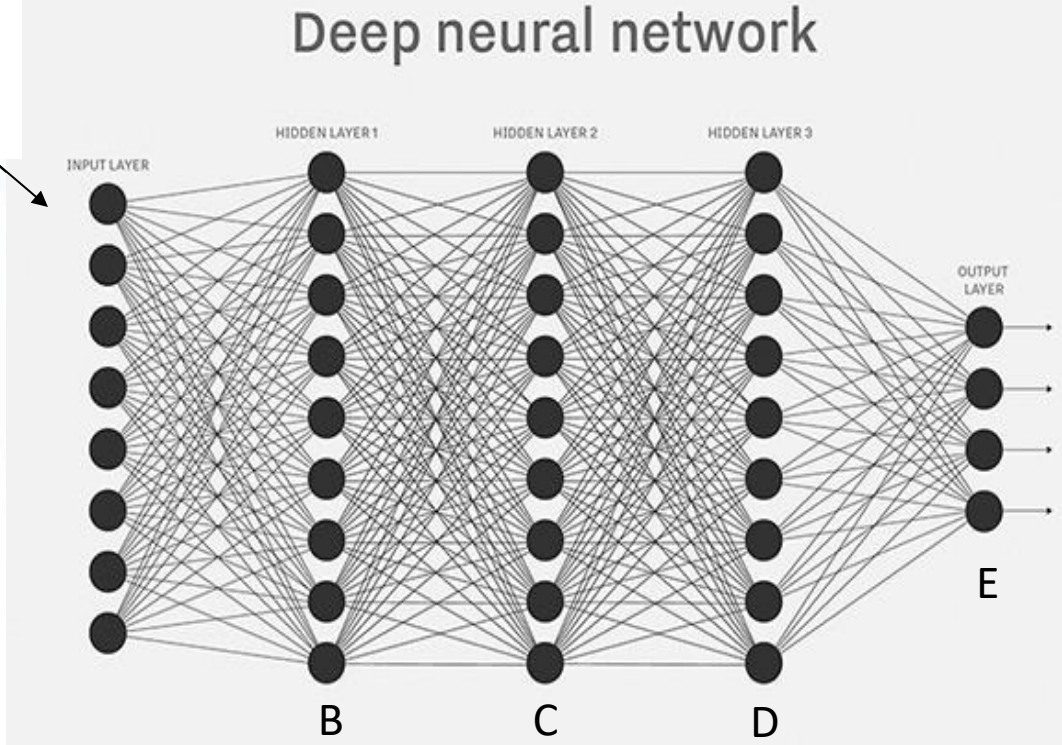

Number of adjustable (trainable) parameters: 50,176 x B + B x C + C x D + D x E = !!!!

From Backpropagation to Deep Learning – Why Now?

1. Advanced Hardware

- Fast, cheap, and efficient vector processing hardware (GPUs)
  - top-of-the-line CPU: 114 GFLOPS vs. 9000 GFLOPS per GPU!!

2. More Data

- Internet / Public Databases / Availability of Labeled Data

3. Advanced Theory

- New activation functions for faster training and deeper nets

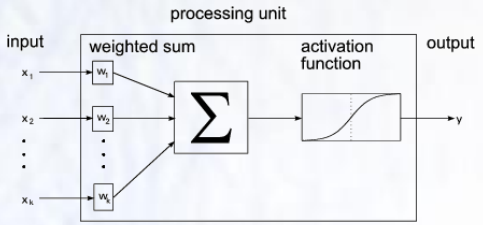

- Willingness to adopt sparse models → Convolutional nets

©Terry Riopka - Aware, Inc. - 2017

From Backpropagation to Deep Learning

Adoption of Sparse Models

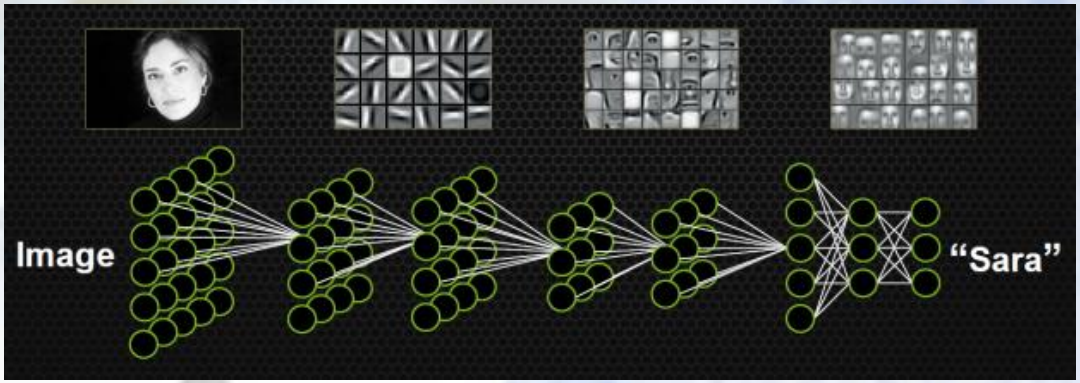

Convolutional Nets

Traditional Backpropagation

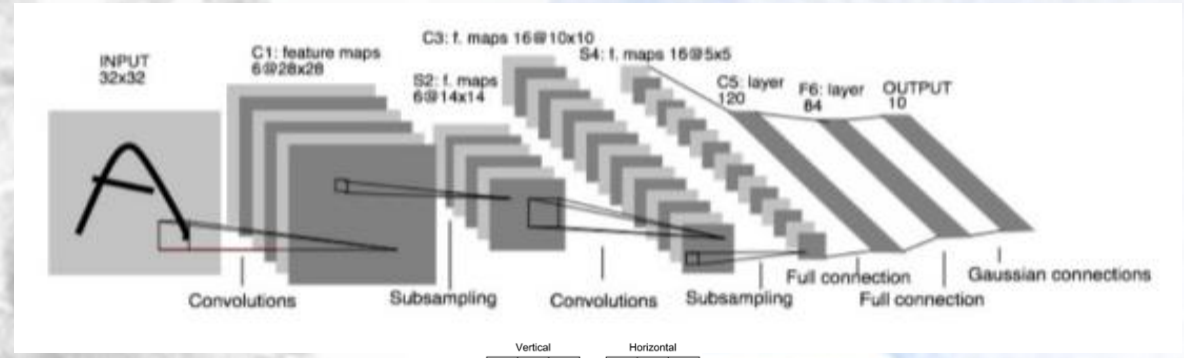

Learning 3x3 Image Filter Coefficients:

Vertical		
-1	0	1
-2	0	2
-1	0	1

Horizontal		
-1	-2	-1
0	0	0
1	2	1

Etc.

©Terry Riopka - Aware, Inc. - 2017

From Backpropagation to Deep Learning

Adoption of Sparse Models

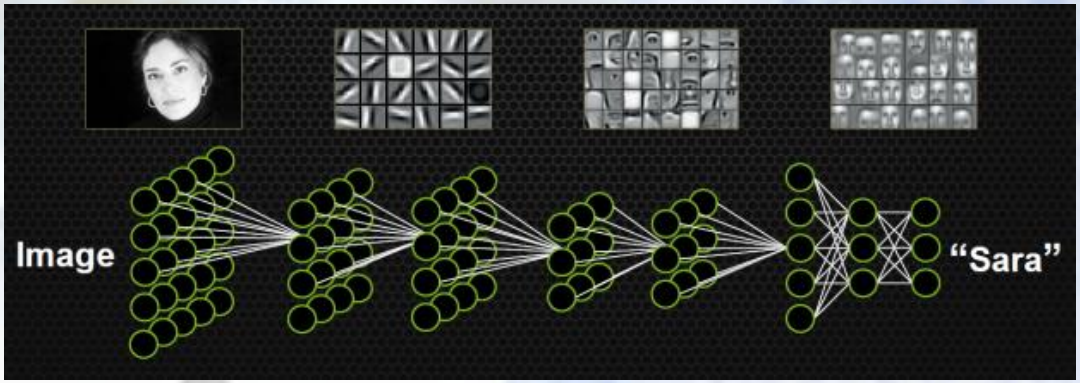

Better Features

=

Better Classification!

Feature Extraction | Classification

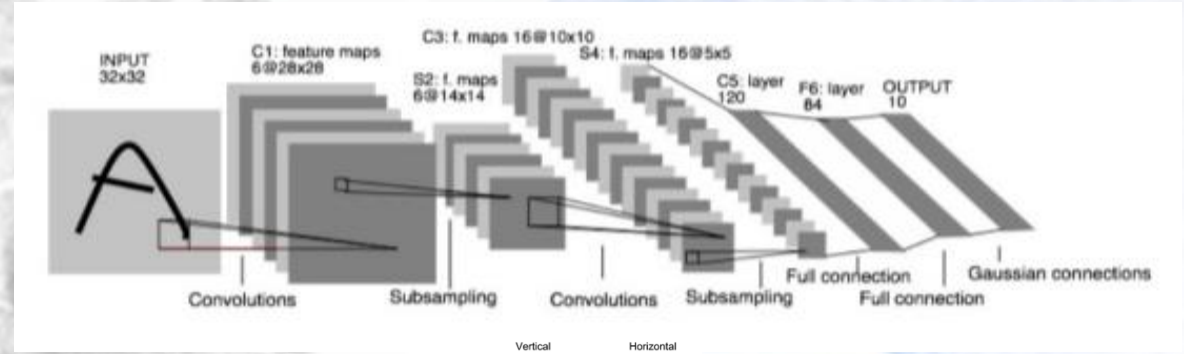

Learning 3x3 Image Filter Coefficients:

Vertical

-1	0	1
-2	0	2
-1	0	1

Horizontal

-1	-2	-1
0	0	0
1	2	1

Etc.

©Terry Riopka - Aware, Inc. - 2017

Deep Learning For Face Recognition

No consensus on preprocessing or feature set !

Face Detection

Preprocessing
?

Feature Extraction
?

engineered features

? ad-hoc features

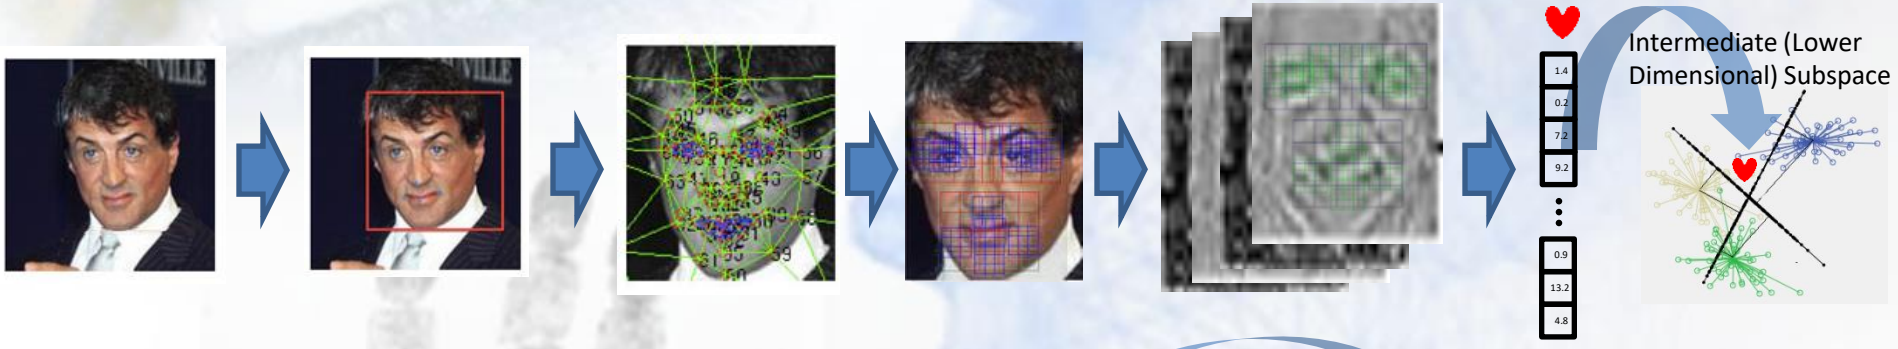

©Terry Riopka - Aware, Inc. - 2017

Deep Learning For Face Recognition

Adoption of Sparse Models

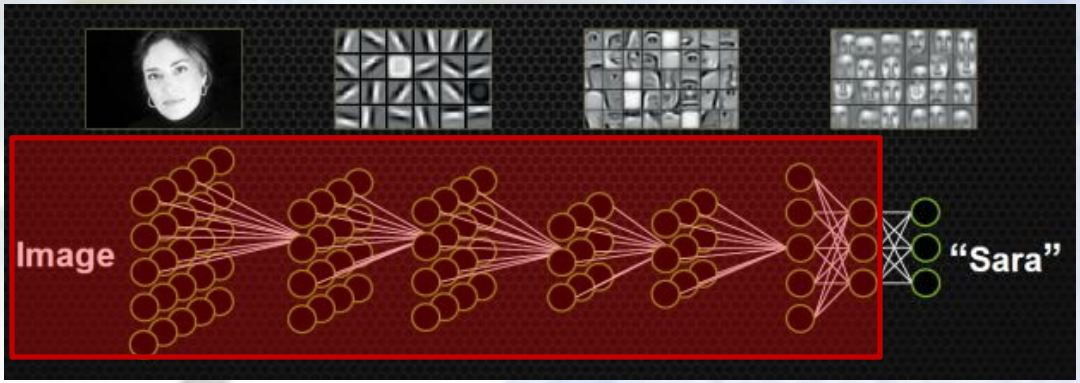

Better Features

=

Better Classification!

Feature Extraction | Classification

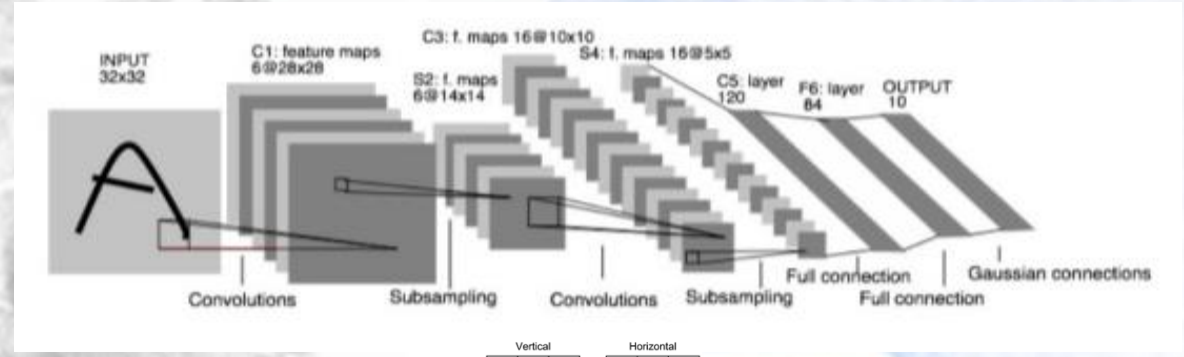

Learning 3x3 Image Filter Coefficients:

Vertical		
-1	0	1
-2	0	2
-1	0	1

Horizontal		
-1	-2	-1
0	0	0
1	2	1

Etc.

©Terry Riopka - Aware, Inc. - 2017

Deep Learning For Face Recognition

Face Detection

Preprocessing

Feature Extraction

- engineered features
- learned features
- ad-hoc features

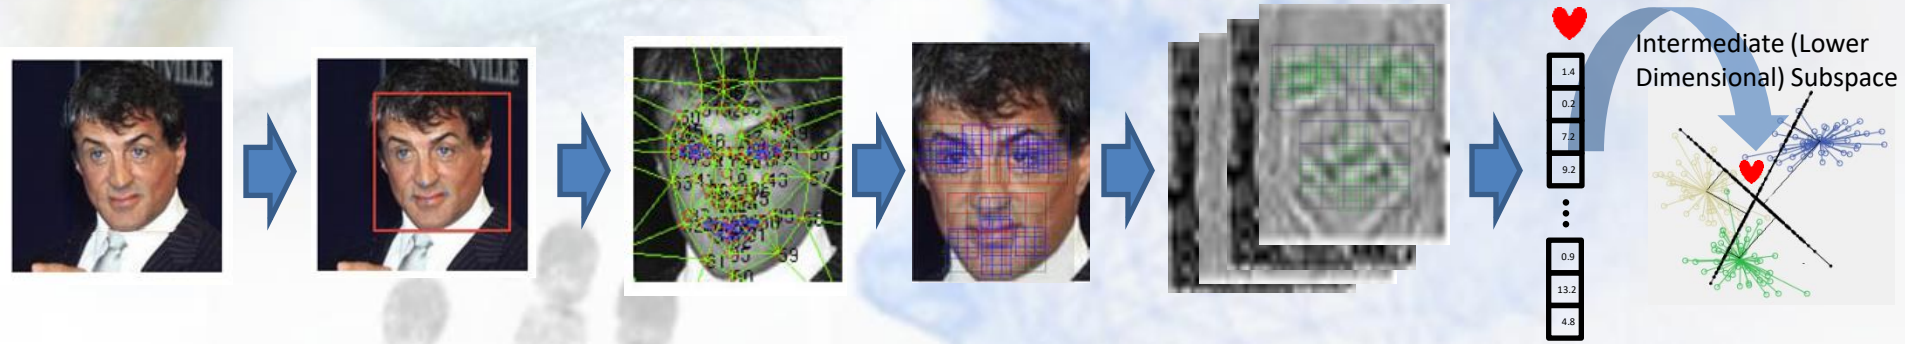

©Terry Riopka - Aware, Inc. - 2017

Deep Learning For Face Recognition

Face Detection

Preprocessing and Feature Extraction

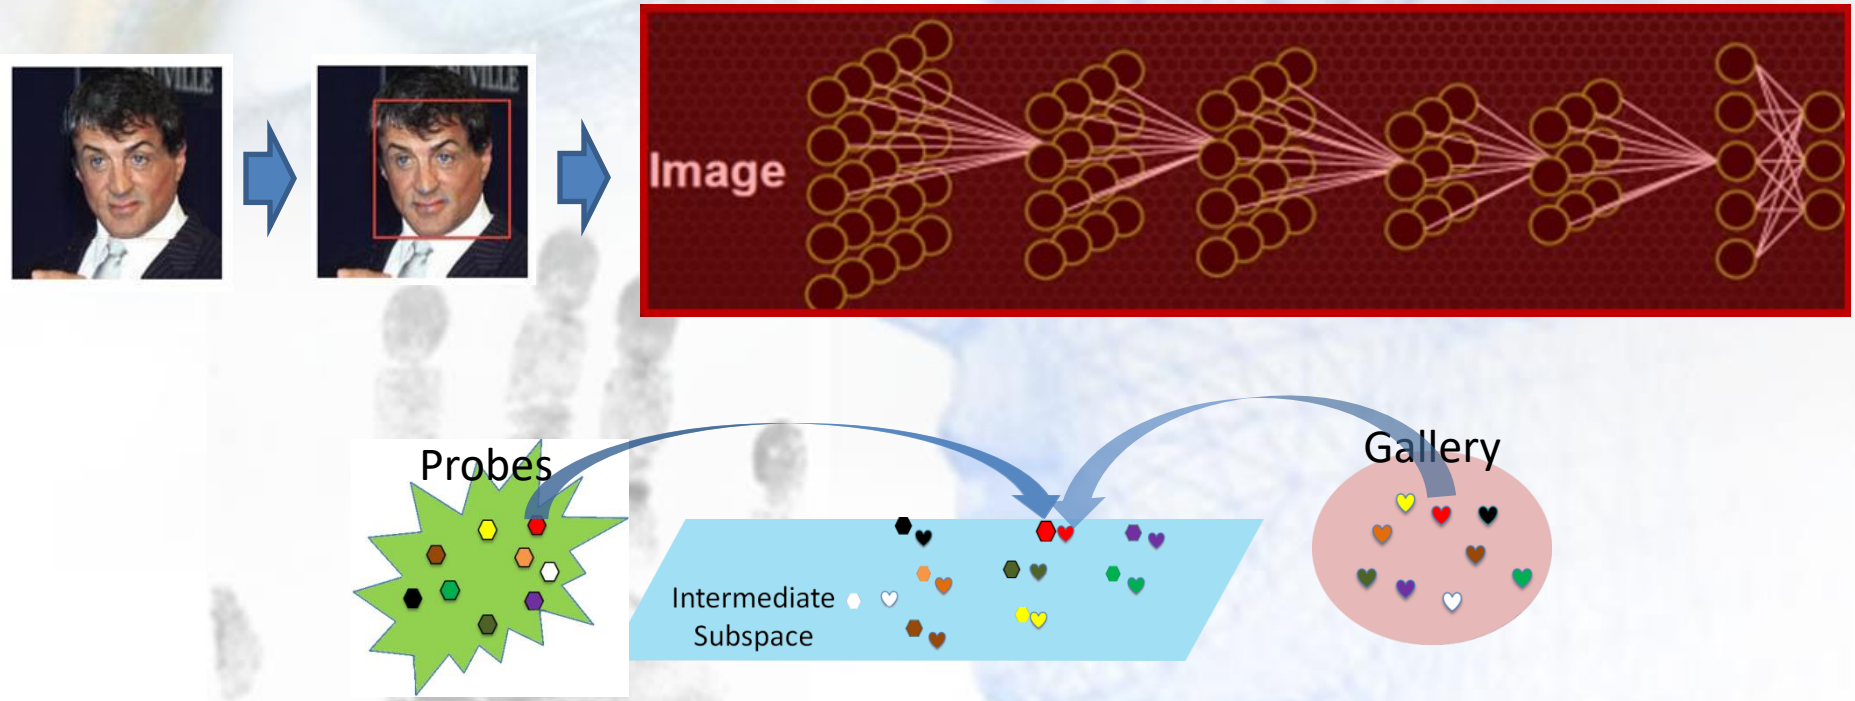

©Terry Riopka - Aware, Inc. - 2017

Deep Learning for Other Biometrics?

Fingerprint Matching - maybe

Known, good features

- bifurcations, ridge endings, ridges, core, *etc.*

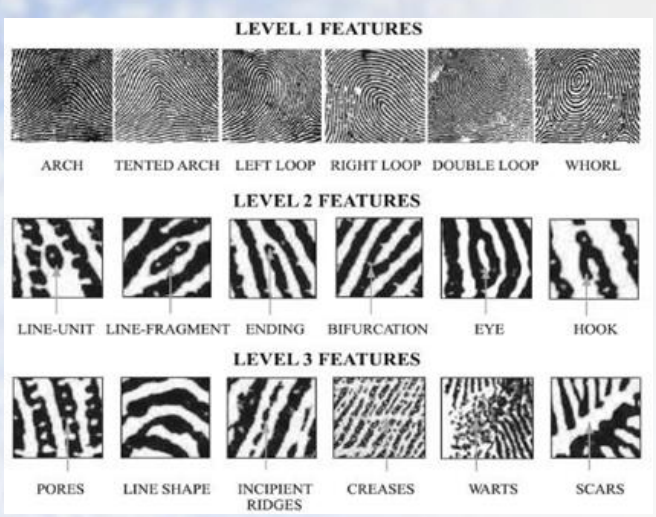

Iris Matching - maybe

Known, good features

- texture/frequency filter responses

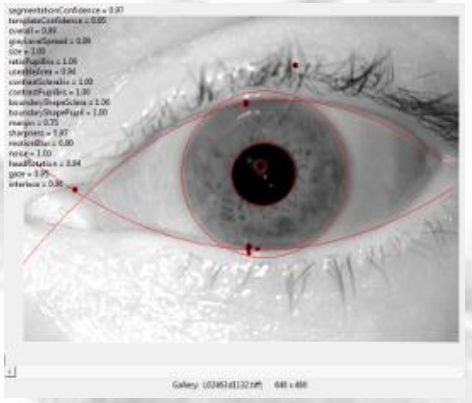

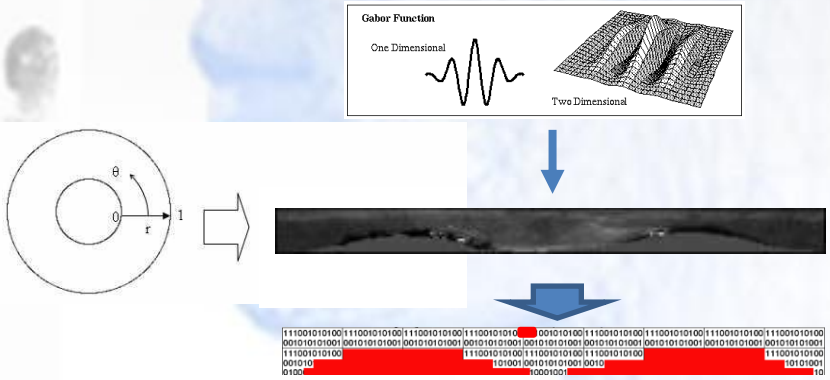

©Terry Riopka - Aware, Inc. - 2017

Deep Learning for Other Biometrics?

Periocular Matching - yes

- relevant parts of the eye can be labeled
- however, they are difficult to quantify or characterize
- can use same training data as face

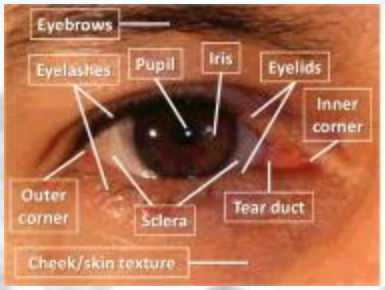

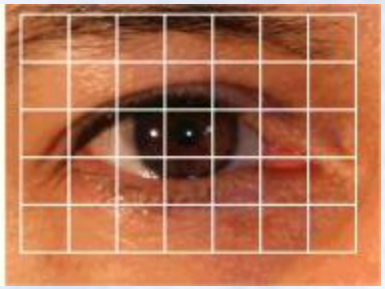

Keystroke Dynamics – no – features can be relatively well defined and engineered

Signature Verification – maybe - training data difficult to acquire

Gait – maybe – training data difficult to acquire

©Terry Riopka - Aware, Inc. - 2017

Deep Learning for Other Biometrics?

Voice/Speech/Language

Speech Processing - yes

- voice synthesis/digitized speech (Siri OS11, 2017)
- machine translation
- speech recognition

Voice Biometrics - maybe

- speaker identification
- speaker verification

}

- known very good features
- robust mathematical, statistical methods

©Terry Riopka - Aware, Inc. - 2017

Deep Learning – The Promise

What is it, really? ➡ Machine learning algorithm for pattern classification and recognition

Will result in more interesting and complex interactions between humans and machines due to combination of:

- Powerful, compact hardware for training and implementation
- Large scale availability of data
- Availability of cheap, diverse sensors enabling large, rapid data acquisition
- Low power WIFI connectivity
- Internet of things

©Terry Riopka - Aware, Inc. - 2017

Deep Learning – The Hype : Challenges and Weaknesses of DL

What is it not? ➡ AI / Reasoning System / Explainable

- Relies on large amounts of *labeled* training data
  - data generation is booming – to be useful, must be → collected AND labeled
  - Need efficient → data storage, organization, accessibility
  - GIGO ➡ Garbage In, Garbage Out

Example: Face Recognition

- robust to → distortions, color variance
- less robust to → blur, occlusion, noise

Why? Trained that way! – still hard to be good at everything!

©Terry Riopka - Aware, Inc. - 2017

Deep Learning – The Hype : Challenges and Weaknesses of DL

What is it not? ➡ AI / Reasoning System / Explainable

- Training time is proportional to amount of data and network size
  - adding additional data requires re-training
- Size of model is proportional to complexity ➡ Challenging for mobile apps
- Can be misled! ➡ Adversarial examples

57.7%
PANDA

©Terry Riopka - Aware, Inc. - 2017

Deep Learning – The Hype : Challenges and Weaknesses of DL

What is it not? ➡ AI / Reasoning System / Explainable

- Training time is proportional to amount of data and network size
  - adding additional data requires re-training
- Size of model is proportional to complexity ➡ Challenging for mobile apps
- Can be misled! ➡ Adversarial examples

57.7%
PANDA

99.3%
MONKEY

GoogleBrain

NIPS 2017

Targeted Adversarial Attack

Develop an adversarial attack that causes image classifiers to predict a specific target class

©Terry Riopka - Aware, Inc. - 2017

Obrigado!

Terry P. Riopka, PhD.
Director R&D

Aware, Inc.
40 Middlesex Tnpke
Bedford, MA 01730
+1.781.276.4000

AWARE