



The poster for the 'Workshop on TinyML for Sustainable Development' is primarily blue. At the top left, it features the logos for ICTP (International Centre for Theoretical Physics), UNESCO, and the United Nations. The main title is in a dark blue box. Below the title, a white box contains the dates '22 - 26 July 2024', the location 'São Paulo, Brazil', and the deadline '6 May 2024'. A central white box provides contact information: 'FURTHER INFORMATION:' with a QR code, 'E-mail: smr3961@ictp.it', 'Web: https://indico.ictp.it/event/10499/', and a note that 'Female scientists are encouraged to apply.'. At the bottom, there are logos for the University of Padua, IBM, UNIFEI, and the event's sponsors.

ICTP  
International Centre  
for Theoretical Physics

UNESCO

22 - 26 July 2024

São Paulo, Brazil

Deadline:  
6 May 2024

FURTHER INFORMATION:

E-mail: smr3961@ictp.it

Web: https://indico.ictp.it/event/10499/

Female scientists are encouraged to apply.

UNIFEI

# From Convolutional Filters to Computer Vision at the Extreme Edge

---



*Brian Plancher*  
*Barnard College, Columbia University*  
[brianplancher.com](http://brianplancher.com)




# Quick Disclaimer:

Today will be **both too fast**  
and **too slow!**

# Computer Vision is Hard

# Computer Vision is Hard

**What color are the pants and the shirt?**



Slide Credit: Hamilton Chong

# Computer Vision is Hard




Slide Credit: Hamilton Chong

# Computer Vision is Hard





Slide Credit: Hamilton Chong

# Computer Vision is Hard



**Is square  
A or B  
darker in  
color?**

# Computer Vision is Hard





What **Features** of the image might be important for self driving cars?



What **Features** of the image might be important for self driving cars?




**Maybe  
straight  
lines to  
see the  
lanes  
of the  
road?**


# How might we find these features?



# How might we find these features?




# How might we find these features?



**Black: 0**

**White: 255**

# How might we find these features?




**Black: 0**  
**White: 255**

0	255	255	255	255
0	0	255	200	255
0	0	0	255	255
0	0	0	0	255
0	0	0	0	0

# How might we find these features?

Look for a Big Change!



**Black: 0**  
**White: 255**

0	255	255	255	255
0	0	255	200	255
0	0	0	255	255
0	0	0	0	255
0	0	0	0	0

# How might we find these features?

## Convolutions



# How might we find these features?

## Convolutions

Original Image

0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255

# How might we find these features?

## Convolutions

Original Image

0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255

Filter

-1	0	1
-1	0	1
-1	0	1

# How might we find these features?

## Convolutions

Original Image

0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255

Filter

-1	0	1
-1	0	1
-1	0	1

# How might we find these features?

## Convolutions

Original Image

0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255

Filter

-1	0	1
-1	0	1
-1	0	1

Output  
Feature Map

765
-----

# How might we find these features?

## Convolutions

Original Image

0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255
0	0	0	255	255	255

Filter



-1	0	1
-1	0	1
-1	0	1

Output  
Feature Map


0	765	765	0
0	765	765	0
0	765	765	0
0	765	765	0

# How might we find these features?




## Convolutions



# How might we combine these features to **classify an object**?



# The ImageNet Challenge and the birth of CNNs



The ImageNet Challenge provided 1.2 million examples of 1,000 **labeled** items and challenged algorithms to learn from the data and then was tested on another 100,000 images



# The ImageNet Challenge and the birth of CNNs






**Traditional Machine Learning Flow**

**Regression, Clustering, etc.**

**Vertical Lines, Horizontal Lines,  
Changes in Color, Changes in  
Focus, etc.**




# The ImageNet Challenge and the birth of CNNs



In 2010 teams had  
**75-50%** error

In 2011 teams had  
**75-25%** error

# The ImageNet Challenge and the birth of CNNs



In 2012 still no team  
had less than 25%  
error barrier except  
**AlexNet at 15%**

# The ImageNet Challenge and the birth of CNNs



**Traditional Machine Learning Flow**




**Deep Learning Flow**

**Let the computer figure out its own features and how to combine them!**

# AlexNet

Use convolutions to find features and the summarize them into higher level features




Combine the features to classify the various objects in the dataset

# How might we find these features?

## Convolutions




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



# How might we find these features?



## Convolutions

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



# How might we find these features?


## Convolutions





# How might we find these features?



## Convolutions



# How might we find these features?



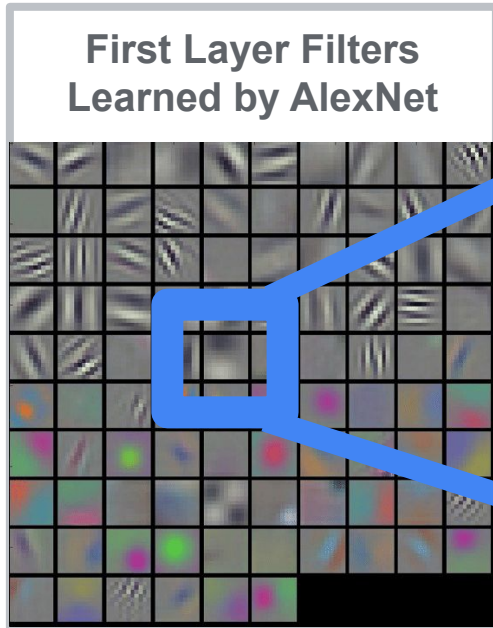
## Convolutions

First Layer Filters  
Learned by AlexNet




# How might we find these features?

## Convolutions




# AlexNet

Use convolutions to find features and the summarize them into higher level features



Combine the features to classify the various objects in the dataset

# The ImageNet Challenge and the birth of CNNs




[https://www.researchgate.net/figure/Historical-top5-error-rate-of-the-annual-winner-of-the-ImageNet-image-classification\\_fig7\\_303992986](https://www.researchgate.net/figure/Historical-top5-error-rate-of-the-annual-winner-of-the-ImageNet-image-classification_fig7_303992986)

# A word of caution...

Ackerman "Hacking the Brain With Adversarial Images"



+  $\epsilon$



=



"panda"

57.7% confidence

There is **no model** of  
the world semantically  
just mathematically

"gibbon"

99.3% confidence


<https://www.vox.com/future-perfect/2019/4/8/18297410/ai-tesla-self-driving-cars-adversarial-machine-learning>

# A word of caution...

Ackerman "Hacking the Brain With Adversarial Images"



+  $\epsilon$



=



"panda"

57.7% confidence

There is **no model** of  
the world semantically  
just mathematically

"gibbon"

99.3% confidence

So how do we bring this to the extreme edge?

**Well I'll let Marcelo  
tell you! :)**