


Calcul MATLAB sur GPU

Réunion des utilisateurs du mésocentre MUST - USMB

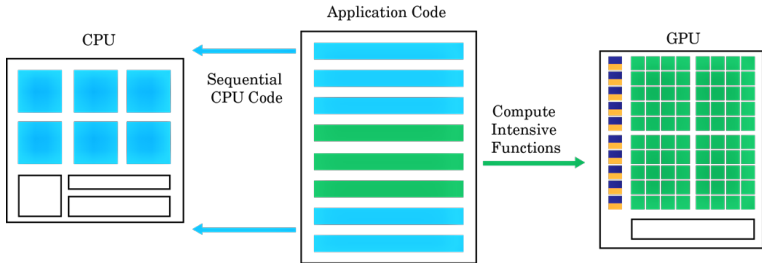
Abdourrahmane M. ATTO
Université Savoie Mont Blanc, LISTIC, France

Le 26-06-2019
à l'Auditorium du Laboratoire d'Annecy de Physique des Particules



Calcul à Hautes Performances

Architectures duales (CPU, GPU)



Matlab@Programmeurs

Calculs CUDA@GPU "via" Matlab :

- Conversion (moindre effort) de code Matlab;
- Distribution d'une partie conséquente de ce code sur GPU;
- Accélération des calculs massifs (4x à 20x).

CUDA@Programmeurs

Exploiter l'ergonomie Matlab pour:

- Évaluer (tests) des noyaux CUDA;
- Explorer les paramètres d'une combinaison de noyaux CUDA.


GPU coder

Optimisation sur les architectures duales (CPU, GPU)

\mathcal{T} = Partitionnement d'un programme séquentiel et génération d'un code CUDA

$$\mathcal{S} = \left\{ \begin{array}{l} \text{Exécuter les opérations massivement parallèle sur GPU.} \\ \text{Exécuter les sections séquentielles de codes sur CPU.} \\ \text{Minimiser les coûts de communication/transfert entre CPU et GPU.} \end{array} \right.$$

GPU Coder = $\mathcal{T}_{\text{Optimal}}$ / Sous contraintes : \mathcal{S}



Exemple1@CPU

```
M3D = rand(NbSamples, NbSamples,
NbIteration);
M2D = rand(NbSamples,NbSamples);
parfor/for idx = 1:NbIteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

Exemple1@(CPU, GPU)

Adaptation (1)

```
M3D = rand(NbSamples, NbSamples, NbIteration, 'gpuArray');
M2D = rand(NbSamples,NbSamples, 'gpuArray');
parfor/for idx = 1:NbIteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

Adaptation (2)

```
@() pagefun(@mtimes, rand(NbSamples, NbSamples,
NbIteration, 'gpuArray'), rand(NbSamples, 'gpuArray'));
```

Exemple2@CPU

```
function u=WaveEquationCPU(u, un, h, b, NbIteration)

for i=1:NbIteration
    v=conv2(u, h, 'same');
    utemp = 2u - un + v - b(u - un) ;
    un = u;
    u = utemp;
    :
    :
end

end
```

Exemple2@(CPU, GPU)

```
function u=WaveEquationGPU(u, un, h, b, NbIteration)
    u=gpuArray(u); un=gpuArray(un); h=gpuArray(h);
    for i=1:NbIteration
        v=conv2(u, h, 'same');
        utemp = 2u - un + v - b(u - un) ;
        un = u;
        u = utemp;
        :
        :
    end
    u=gather(u);
end
```

Informations sur les cartes GPU *via* Matlab

```

%% Nombre de cartes GPU
>> N = gpuDeviceCount()
    1

%% Informations sur la carte GPU n
>> ContenuCarteGPU = gpuDevice(n)

%% Selection/Imposer - carte GPU
if N > 1
    gpuDevice(2)
else
    gpuDevice(1)
end
  
```

MUST-1

@localGPU*

gpuDeviceCountM1	2
Name	'Tesla K80'
ComputeCapability	'3.7'
SupportsDouble	1
DriverVersion	9.1000
ToolkitVersion	8
MaxThreadsPerBlock	1024
MaxShmemPerBlock	49152
MaxThreadBlockSize	[1024 1024 64]
MaxGridSize	[2.1475e+09 65535 65535]
SIMDWidth	32
TotalMemory	1.1997e+10
AvailableMemory	1.1862e+10
MultiprocessorCount	13
ClockRateKHz	82 3500

MUST-2

@localGPU-Cécile

gpuDeviceCountM2	1
Name	'Tesla V100-PCIE-16GB'
ComputeCapability	'7.0'
SupportsDouble	1
DriverVersion	9.2000
ToolkitVersion	9.1000
MaxThreadsPerBlock	1024
MaxShmemPerBlock	49152
MaxThreadBlockSize	[1024 1024 64]
MaxGridSize	[2.1475e+09 65535 65535]
SIMDWidth	32
TotalMemory	1.6946e+10
AvailableMemory	9.2931e+09
MultiprocessorCount	80
ClockRateKHz	1 380 000



```
NbSamples = 10 (liluputiens) / 5000 (titans)
Nblteration = 10 (liluputiens) / 1000 (titans)
```

Exemple1@CPU

```
M3D = rand(NbSamples, NbSamples,
Nblteration);
M2D = rand(NbSamples,NbSamples);
parfor/for idx = 1:Nblteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

Exemple1@GPU

CPU2GPU Adaptation 1

```
M3D = rand(NbSamples, NbSamples, Nblteration, 'gpuArray');
M2D = rand(NbSamples,NbSamples, 'gpuArray');
parfor/for idx = 1:Nblteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

end

CPU2GPU Adaptation 2

```
@() pagefun(@mtimes, rand(NbSamples, NbSamples,
Nblteration, 'gpuArray'), rand(NbSamples, 'gpuArray'));
```

Exemple2@CPU

```
function u=WaveEquationCPU(u, un, h, Nblteration)
    % Size of wave field = 3000x3000 pixels
    for i=1:Nblteration
        v=conv2(u, h, 'same');
        u_temp = 2un - u + v ;
        un = u;
        u = u_temp;
        :
        :
    end
    %%%%%%%%%%
end
```

Exemple2@GPU

```
function u=WaveEquationGPU(u, un, h, Nblteration)
    u=gpuArray(u); un=gpuArray(un); h=gpuArray(h);
    for i=1:Nblteration
        v=conv2(u, h, 'same');
        u_temp = 2un - u + v ;
        un = u;
        u = u_temp;
        :
        :
    end
    u=gather(u);
end
```

NbSamples = 10 (liluputiens) / 5000 (titans)

Nbliteration = 10 (liluputiens) / 1000 (titans)

Exemple1@CPU 59.6" [titans]

```
M3D = rand(NbSamples, NbSamples,
Nbliteration);
M2D = rand(NbSamples,NbSamples);
parfor/for idx = 1:Nbliteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

Exemple1@GPU

0.8" [titans]

CPU2GPU Adaptation 1

```
M3D = rand(NbSamples, NbSamples, Nbliteration, 'gpuArray');
M2D = rand(NbSamples,NbSamples, 'gpuArray');
parfor/for idx = 1:Nbliteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

end

CPU2GPU Adaptation 2

```
@() pagefun(@mtimes, rand(NbSamples, NbSamples,
Nbliteration, 'gpuArray'), rand(NbSamples, 'gpuArray'));
```

Exemple2@CPU 53.0" [titans]

```
function u=WaveEquationCPU(u, un, h, Nbliteration)
% Size of wave field = 3000x3000 pixels
for i=1:Nbliteration
    v=conv2(u, h, 'same');
    u_temp = 2u_n - u + v ;
    u_n = u;
    u = u_temp;
    :
end
%%%%%%%%%%
end
```

Exemple2@GPU 1.6" [titans]

```
function u=WaveEquationGPU(u, un, h, Nbliteration)
u=gpuArray(u); un=gpuArray(un); h=gpuArray(h);
for i=1:Nbliteration
    v=conv2(u, h, 'same');
    u_temp = 2u_n - u + v ;
    u_n = u;
    u = u_temp;
    :
end
u=gather(u);
end
```

NbSamples = 10 (liluptiens) / 5000 (titans)

Nblteration = 10 (liluptiens) / 1000 (titans)

Exemple1@CPU 5.56" [liluput!s]

```
M3D = rand(NbSamples, NbSamples,
Nblteration);
M2D = rand(NbSamples,NbSamples);
parfor/for idx = 1:Nblteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

Exemple1@GPU

0.01" [liluput!s]

CPU2GPU Adaptation 1

```
M3D = rand(NbSamples, NbSamples, Nblteration, 'gpuArray');
M2D = rand(NbSamples,NbSamples, 'gpuArray');
parfor/for idx = 1:Nblteration
    mtimes(squeeze(M3D(:,:,idx)),M2D);
end
```

end

CPU2GPU Adaptation 2

```
@() pagefun(@mtimes, rand(NbSamples, NbSamples,
Nblteration, 'gpuArray'), rand(NbSamples, 'gpuArray'));
```

Exemple2@CPU

0.1" [liluput!s]

```
function u=WaveEquationCPU(u, un, h, Nblteration)
% Size of wave field = 3000x3000 pixels
for i=1:Nblteration
    v=conv2(u, h, 'same');
    u_temp = 2u_n - u + v ;
    u_n = u;
    u = u_temp;

    :
end
%%%%%%%%%%
end
```

Exemple2@GPU

0.5" [liluput!s]

```
function u=WaveEquationGPU(u, un, h, Nblteration)
u=gpuArray(u); un=gpuArray(un); h=gpuArray(h);
for i=1:Nblteration
    v=conv2(u, h, 'same');
    u_temp = 2u_n - u + v ;
    u_n = u;
    u = u_temp;

    :
end
u=gather(u);
end
```


MUST-1

@localGPU*

gpuDeviceCountM1	2
Name	'Tesla K80'
ComputeCapability	'3.7'
SupportsDouble	1
DriverVersion	9.1000
ToolkitVersion	8
MaxThreadsPerBlock	1024
MaxShmemPerBlock	49152
MaxThreadBlockSize	[1024 1024 64]
MaxGridSize	[2.1475e+09 65535 65535]
SIMDWidth	32
TotalMemory	1.1997e+10
AvailableMemory	1.1862e+10
MultiprocessorCount	13
ClockRateKHz	82 3500

MUST-2

@localGPU-Cécile

gpuDeviceCountM2	1
Name	'Tesla V100-PCIE-16GB'
ComputeCapability	'7.0'
SupportsDouble	1
DriverVersion	9.2000
ToolkitVersion	9.1000
MaxThreadsPerBlock	1024
MaxShmemPerBlock	49152
MaxThreadBlockSize	[1024 1024 64]
MaxGridSize	[2.1475e+09 65535 65535]
SIMDWidth	32
TotalMemory	1.6946e+10
AvailableMemory	9.2931e+09
MultiprocessorCount	80
ClockRateKHz	1 380 000

MUST-1@K80

@localGPU*

MTimes-Rand-3Dx2D:	9.460067
Wave-Propagation-Equation:	4.364735

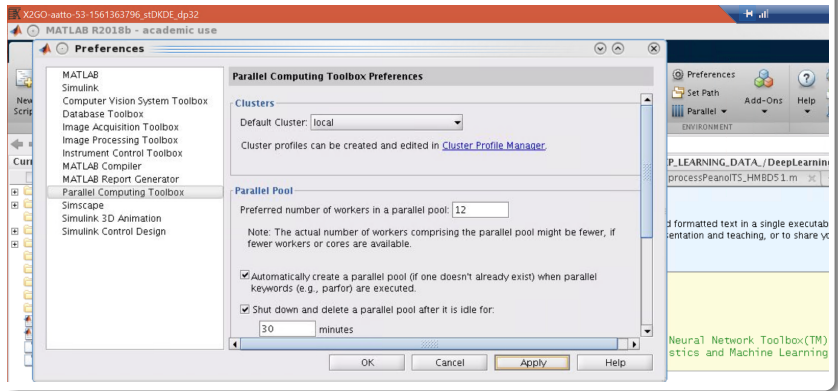
MUST-2@V100

@localGPU-Cécile

MTimes-Rand-3Dx2D:	0.833975
Wave-Propagation-Equation:	1.632387

Si pool CPU plus conséquent

Matlab@GPU ↗



The screenshot shows the MATLAB R2018b Preferences dialog box, specifically the Parallel Computing Toolbox Preferences section. The window title is "MATLAB R2018b - academic use". The left sidebar shows the "Parallel Computing Toolbox" selected. The main area is titled "Parallel Computing Toolbox Preferences" and contains the following settings:

- Clusters:** Default Cluster: local
- Parallel Pool:** Preferred number of workers in a parallel pool: 12
- Note: The actual number of workers comprising the parallel pool might be fewer, if fewer workers or cores are available.
- Automatically create a parallel pool (if one doesn't already exist) when parallel keywords (e.g., parfor) are executed.
- Shut down and delete a parallel pool after it is idle for: 30 minutes

Buttons at the bottom: OK, Cancel, Apply, Help.

Versions

Matlab / CUDA Toolkit / \exists_{ance} Matlab-Built-In@GPU

Error In DeepSegmentLearningGPU (line 176) \rightarrow MultivariateKernelDensity

There is a problem with the graphics driver or with this GPU device.

Be sure that you have a supported GPU and that the latest driver is installed.

Documentation Home

Parallel Computing Toolbox

GPU Computing

GPU Computing in MATLAB

Parallel Computing Toolbox

GPU Computing

Run Built-In Functions on a GPU

ON THIS PAGE

MATLAB Functions with gpuArray Arguments

Example: Functions with gpuArray Input and Output

Sparse Arrays on a GPU

Considerations for Complex Numbers

Acknowledgments

See Also

MATLAB Functions with gpuArray Arguments

Many MATLAB[®] built-in functions support `gpuArray` input arguments. Whenever any of these functions is called with at least one `gpuArray` as an input argument, the function executes on the GPU and generates a same function call; the MATLAB arrays are transferred to the GPU for the function execution. Supporting functions include the discrete Fourier transform (`fft`), matrix multiplication (`mtimes`), and left matrix division.

The following functions and their symbol operators are enhanced to accept `gpuArray` input arguments so that they execute on the GPU:

<code>abs</code>	<code>copan</code>	<code>flip</code>	<code>isnan</code>
<code>acos</code>	<code>complex</code>	<code>fliplr</code>	<code>isnumeric</code>
<code>acosd</code>	<code>cond</code>	<code>flipud</code>	<code>isreal</code>
<code>acosh</code>	<code>conj</code>	<code>floor</code>	<code>isrow</code>
<code>acot</code>	<code>conv</code>	<code>fprintf</code>	<code>issorted</code>
<code>acotd</code>	<code>conv2</code>	<code>full</code>	<code>issparse</code>
<code>acoth</code>	<code>convn</code>	<code>gamma</code>	<code>issymmetric</code>
<code>acsc</code>	<code>corrcoef</code>	<code>gamaInc</code>	<code>istril</code>
<code>acscd</code>	<code>cos</code>	<code>gamaIncInv</code>	<code>istriu</code>
<code>acsch</code>	<code>cosd</code>	<code>gamaIn</code>	<code>isvector</code>
<code>accumarray</code>	<code>cosh</code>	<code>gather</code>	<code>kron</code>
<code>all</code>	<code>cot</code>	<code>ge</code>	<code>ldivide</code>
<code>and</code>	<code>cotd</code>	<code>gmes</code>	<code>le</code>
<code>angle</code>	<code>coth</code>	<code>gradient</code>	<code>legendre</code>
<code>any</code>	<code>cov</code>	<code>gt</code>	<code>length</code>
<code>arrayfun</code>	<code>cross</code>	<code>hankel</code>	<code>log</code>
<code>asec</code>	<code>csc</code>	<code>head</code>	<code>log10</code>
<code>asecd</code>	<code>cscd</code>	<code>histcounts</code>	<code>log1p</code>
<code>asech</code>	<code>cscch</code>	<code>horzcat</code>	<code>log2</code>
<code>asin</code>	<code>ctranspose</code>	<code>hsv2rgb</code>	<code>logical</code>
<code>asind</code>	<code>cummax</code>	<code>hypot</code>	<code>lsqr</code>
<code>asinh</code>	<code>cummin</code>	<code>idivide</code>	<code>lt</code>
<code>assert</code>	<code>cumprod</code>	<code>iffft</code>	<code>lu</code>
<code>atan</code>	<code>cumsum</code>	<code>iffft2</code>	<code>mat2str</code>
<code>atan2</code>	<code>deg2rad</code>	<code>ifftn</code>	<code>max</code>
<code>atan2d</code>	<code>delz</code>	<code>iffftshift</code>	<code>median</code>
<code>atand</code>	<code>det</code>	<code>imag</code>	<code>mean</code>
<code>atanh</code>	<code>detrend</code>	<code>ind2sub</code>	<code>meshgrid</code>
<code>bandwidth</code>	<code>diag</code>	<code>Inf</code>	<code>min</code>
<code>besselj</code>	<code>diff</code>	<code>inpolygon</code>	<code>minus</code>
<code>bessely</code>	<code>discretize</code>	<code>int16</code>	<code>nldivide</code>
<code>beta</code>	<code>disp</code>	<code>int2str</code>	<code>mod</code>
<code>betainc</code>	<code>display</code>	<code>int32</code>	<code>mode</code>
<code>betaincinv</code>	<code>dot</code>	<code>int64</code>	<code>nommean</code>
<code>betaIn</code>	<code>double</code>	<code>int8</code>	<code>novstd</code>
<code>big</code>	<code>eig</code>	<code>interp1</code>	<code>novsum</code>
<code>bigstab</code>	<code>eps</code>	<code>interp2</code>	<code>novar</code>
<code>bitand</code>	<code>eq</code>	<code>interp3</code>	<code>npower</code>
<code>bitcep</code>	<code>erf</code>	<code>interp</code>	<code>nrdivide</code>
<code>bitget</code>	<code>erfc</code>	<code>intersect</code>	<code>ntimes</code>
<code>bitor</code>	<code>erfcinv</code>	<code>inv</code>	<code>NaN</code>

Set options for training




```
opts = trainingOptions('sgdm');
```

Train the network

```
net = trainNetwork ...
(XTrain, TTrain, layers, opts);
```

Make predictions

```
trainFeatures = ...
activations(net, XTrain, 6);
```



```
layers = [ imageInputLayer([512 512 3])
  convolution2dLayer([3 3],96);
  ...
  softmaxLayer()
  pixelClassificationLayer() ]
```

```
opts = trainingOptions('sgdm',...
  'ExecutionEnvironment', 'multi-gpu');
```

>> net = alexnet;

>> analyzeNetwork(net)

Deep Learning Network Analyzer

— □ ×

net

Analysis date: 25-Jun-2019 21:59:33

25 
layers0 
warnings0 
errors

data



ANALYSIS RESULT

#	NAME	TYPE	ACTIVATIONS	LEARNABLES
1	data 227x227x3 Images with 'zerocenter' normalization	Image Input	227x227x3	-
2	conv1 96 11x11x3 convolutions with stride [4 4] and padding [0 0 0]	Convolution	55x55x96	Weights 11x11x3x96 Bias 1x1x96
3	relu1 ReLU	ReLU	55x55x96	-
4	norm1 cross channel normalization with 5 channels per element	Cross Channel Normalization	55x55x96	-
5	pool1 3x3 max pooling with stride [2 2] and padding [0 0 0]	Max Pooling	27x27x96	-
6	conv2 256 5x5x48 convolutions with stride [1 1] and padding [2 2 2 2]	Convolution	27x27x256	Weights 5x5x48x256 Bias 1x1x256
7	relu2 ReLU	ReLU	27x27x256	-
8	norm2 cross channel normalization with 5 channels per element	Cross Channel Normalization	27x27x256	-
9	pool2 3x3 max pooling with stride [2 2] and padding [0 0 0]	Max Pooling	13x13x256	-
10	conv3 384 3x3x256 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	13x13x384	Weights 3x3x256x384 Bias 1x1x384
11	relu3 ReLU	ReLU	13x13x384	-
12	conv4 384 3x3x192 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	13x13x384	Weights 3x3x192x384 Bias 1x1x384
13	relu4 ReLU	ReLU	13x13x384	-
14	conv5 256 3x3x192 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	13x13x256	Weights 3x3x192x256 Bias 1x1x256
15	relu5 ReLU	ReLU	13x13x256	-

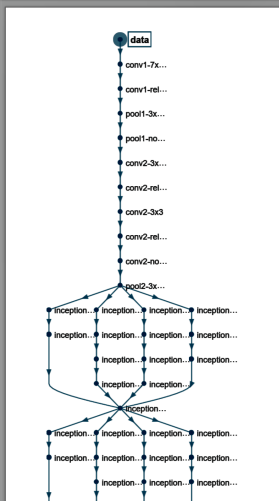
>> net = googlenet;

>> analyzeNetwork(net)

Deep Learning Network Analyzer

net

Analysis date: 25-Jun-2019 21:59:33

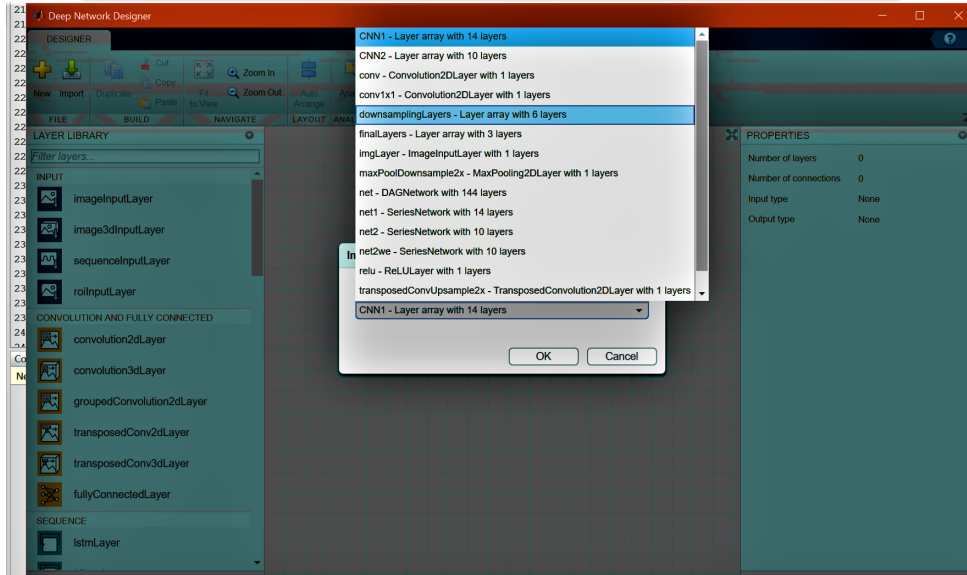
144 
layers

ANALYSIS RESULT

#	NAME	TYPE	ACTIVATIONS	LEARNABL
1	data 224x224x3 images with 'zero-center' normalization	Image Input	224x224x3	-
2	conv1-7x7_s2 64 7x7x3 convolutions with stride [2 2] and padding [3 3 3 3]	Convolution	112x112x64	Weights Bias
3	conv1-relu_7x7 ReLU	ReLU	112x112x64	-
4	pool1-3x3_s2 3x3 max pooling with stride [2 2] and padding [0 1 0 1]	Max Pooling	56x56x64	-
5	pool1-norm1 cross channel normalization with 5 channels per element	Cross Channel Normalization	56x56x64	-
6	conv2-3x3_reduce 64 1x1x64 convolutions with stride [1 1] and padding [0 0 0 0]	Convolution	56x56x64	Weights Bias
7	conv2-relu_3x3_reduce ReLU	ReLU	56x56x64	-
8	conv2-3x3 192 3x3x64 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	56x56x192	Weights Bias
9	conv2-relu_3x3 ReLU	ReLU	56x56x192	-
10	conv2-norm2 cross channel normalization with 5 channels per element	Cross Channel Normalization	56x56x192	-
11	pool2-3x3_s2 3x3 max pooling with stride [2 2] and padding [0 1 0 1]	Max Pooling	28x28x192	-
12	inception_3a-1x1 64 1x1x192 convolutions with stride [1 1] and padding [0 0 0 0]	Convolution	28x28x64	Weights Bias
13	inception_3a-relu_1x1 ReLU	ReLU	28x28x64	-
14	inception_3a-3x3_reduce 96 1x1x192 convolutions with stride [1 1] and padding [0 0 0 0]	Convolution	28x28x96	Weights Bias
15	inception_3a-relu_3x3_reduce ReLU	ReLU	28x28x96	-
16	inception_3a-3x3 128 3x3x96 convolutions with stride [1 1] and padding [1 1 1 1]	Convolution	28x28x128	Weights Bias
17	inception_3a-relu_3x3 ReLU	ReLU	28x28x128	-

Outil pour DL Designer

>> deepNetworkDesigner



The screenshot displays the 'Deep Network Designer' application. A central dialog box lists various layer types for selection. The 'CNN1 - Layer array with 14 layers' option is highlighted. The background interface includes a 'DESIGNER' toolbar, a 'LAYER LIBRARY' on the left, and a 'PROPERTIES' panel on the right.

LAYER LIBRARY

Filter layers...

INPUT

- imageInputLayer
- image3dInputLayer
- sequenceInputLayer
- roiInputLayer

CONVOLUTION AND FULLY CONNECTED

- convolution2dLayer
- convolution3dLayer
- groupedConvolution2dLayer
- transposedConv2dLayer
- transposedConv3dLayer
- fullyConnectedLayer

SEQUENCE

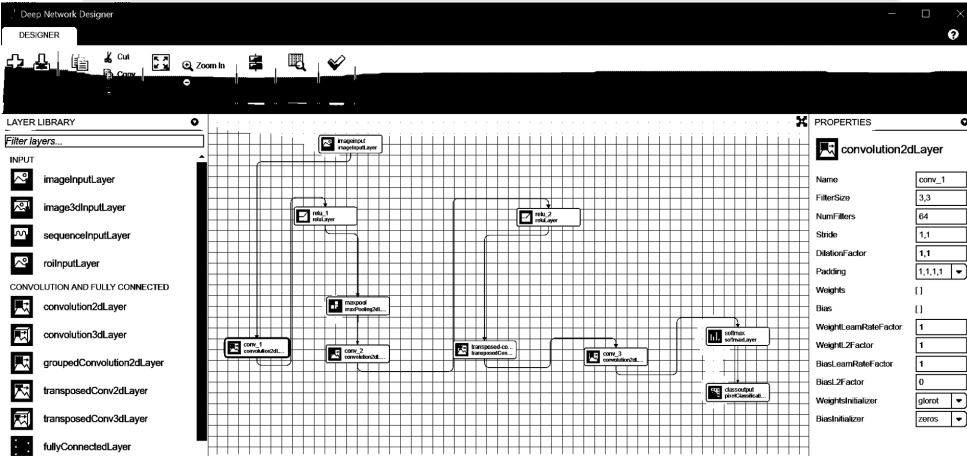
- lstmLayer

PROPERTIES

Number of layers	0
Number of connections	0
Input type	None
Output type	None

>> deepNetworkDesigner

Une histoire de chemins



Explorer les chemins le jour chasse la peur de se promener la nuit.

Aller là où il n'y a pas de chemin et laissez une trace (RWE).

Le chemin se construit en marchant (AM). La bonne volonté raccourci le chemin (BR).

Si vous ne savez pas où vous allez, n'importe quel chemin vous y mènera (LC).

La route fût longue mais le chemin est beau (DPY).


Si tous les chemins mènent à l'IA, les Krees sont une facette de l'homo-3.0 au sens du Suprémor.

Application DL pour la détection de violence visuelle par analyse de vidéos

- [1] A. M. Atto, A. Benoit, P. Lambert,
Hilbert Based Video To Timed-Image Filling and Learning To Recognize Visual Violence
<https://dx.doi.org/10.13140/RG.2.2.12648.11526/1>

Observables	Examples
Violence from interactions (external induction)	
1 Human/Human	Battle, slap, punch.
2 Human/Object	Surgery, mutilation, projectile throw.
3 Human/Fluid	Drowning, gazing.
4 Human/Animal	Human attack by animals, animal hunting and shoot.
5 Object/Object	Car <i>versus</i> car accidents, crash of an airplane (ground or sea).
6 Animal/Animal	Animals fight clubs, predator <i>versus</i> prey showdowns.
Suspicious motions (self-induction)	
7 Living body abnormal motion	Terror or aggressive faces, person falling down
8 Inert structure abnormal motion	Conflagrations, explosion, smoke, flames and ashes, flowing blood.
Sensitive objects and symbols	
9 Sensitive objects	Guns, weapons, bombs, broken glasses, stripped electrical socket, blood stain frightful masks and veils.
10 Sensitive symbols	PEGI "-18" / "-16" / "-12" / "-10", flammable material signs, swastika injury and hatred posters

Name	Categories and number of video samples		
VSD-L2	Violence 1 137		Non-Violence 10398
VSD-L3	Moderate Violence 406	Extreme Violence 418	Non-Violence 10398



 Deep-learning / MUST VSD-L2

 3D CNN

Category	<i>Non – Violence</i>	<i>Violence</i>
<i>Non – Violence</i>	94 %	06 %
<i>Violence</i>	78 %	22 %
Mean accuracy	58 %	

 2D-Timed-Image CNN

Category	<i>Non – Violence</i>	<i>Violence</i>
<i>Non – Violence</i>	94 %	6 %
<i>Violence</i>	36 %	64 %
Mean accuracy	79 %	

Information détaillées disponibles dans [1] :

Hilbert Based Video To Timed-Image Filling and Learning To Recognize Visual Violence