



Analyzing 3D objects with power of Deep Learning and Cython

Alexandr Notchenko
avnotchenko@gmail.com



Skoltech

Skolkovo Institute of Science and Technology

Open
Data
Science

A 3D wireframe icon of a brain, composed of black lines connecting vertices to form a mesh structure.

3D Shape representations

- Meshes
- Point clouds
- Implicit surfaces / potentials
- Voxels
- Set of 2D projections

3D Shape representations

- Meshes
- Point clouds
- Implicit surfaces / potentials

Irregular size, not clear
how to use in NN

- Voxels
- Set of 2D projections

Regular size, good to go in CNN

3D Shape representations

- Meshes
 - Point clouds
 - Implicit surfaces / potentials
 - Voxels
 - Set of 2D projections
- Irregular size, not clear how to use in NN
- Regular size, good to go in CNN
- Not really 3D, self occlusion
-

Sparsity of voxel representation

* Mean sparsity for all classes of ModelNet40 train dataset at voxel resolution 40 equal to 5.5%.




Figure: Examples of some objects voxelizations at different resolutions **30**, **50**, **70**, **100** (from left to right), left-most objects are depicted using original meshes

Computer Tomography images of Brain

Image N=30




Image N=35




Image N=40




Image N=44




Image N=49




Image N=53




Image N=58




Image N=62





Image N=67



Computer Tomography images of prostate

Image N=0




Image N=1




Image N=2




Image N=3




Image N=4




Image N=5




Image N=6




Image N=7







Image N=8



Slices of MRI and CT Images with resolutions 110x110x110 (left), 384x384x19 (right)

SparseConvNet

Dr. Benjamin Graham
assoc. prof. at Warwick University
Facebook AI Research, Paris Lab



<> Code

! Issues 0

🔗 Pull requests 0

📁 Projects 0

📖 Wiki

⚡ Pulse

📊 Graphs

⚙️ Settings

Python Framework for sparse neural networks

Edit

[Add topics](#)

📄 200 commits

🌿 10 branches

📦 0 releases

👤 5 contributors

Branch: master ▾












New pull request

Create new file

Upload files

Find file

Clone or download ▾

 gangiman * Moved tests to their own directory ...	Latest commit a07b4ad 9 days ago
 PySCNutils	* Moved tests to their own directory 9 days ago
 SparseConvNet	* added some function for Voxel Pictures like 13 days ago
 tests	* Moved tests to their own directory 9 days ago
 .gitignore	Added a little bit of docstrings a year ago
 Makefile	* Moved tests to their own directory 9 days ago
 README.md	Updated README, added reqs and Installation instruction 4 months ago
 _SparseConvNet.pxd	* PyVoxelPicture - takes sparse matrix as input 20 days ago
 requirements.txt	* added multi-class classification routine 23 days ago
 setup.py	* PyVoxelPicture - takes sparse matrix as input 20 days ago
 sparseNetwork.pyx	* added some function for Voxel Pictures like 13 days ago

PySparseConvNet

Pros

C++ / CUDA kernels

Effective Memory usage

Can use any loss functions

Can access internal layer activations

Interactivity

Train on Sparse data, infer with Dense network

Cons

Not a general purpose Deep Learning Framework

Complicated code base

Non-standard loss functions are in python-land, overhead memory transfer from GPU (for now)

```

45 class SparseConvNetCUDA {
46 public:
47     std::vector<SpatiallySparseLayer*> layers;
48     std::vector<float> inputNormalizingConstants;
49     int dimension;
50     int nClasses;
51     int nTop;
52     int inputSpatialSize;
53     int nInputFeatures;
54     int nOutputFeatures;
55     int deviceID;
56     int nBatchProducerThreads;
57     cudaMemStream memStream;
58     cudaMemStream batchMemStreams [N_MAX_BATCH_PRODUCER_THREADS];
59     std::mutex batchLock [N_MAX_BATCH_PRODUCER_THREADS];
60     std::vector<SpatiallySparseBatch> batchPool;
61     std::vector<SpatiallySparseBatchSubInterface*> initialSubInterfaces;
62     std::vector<SpatiallySparseBatchSubInterface*> sharedSubInterfaces;
63     cublasHandle_t cublasHandle;
64
65     SparseConvNetCUDA(int dimension, int nInputFeatures, int nClasses,
66                     int pciBusID = -1, int nTop = 1,
67                     int nBatchProducerThreads = 1);
68     ~SparseConvNetCUDA();
69     void processBatch(SpatiallySparseBatch &batch, float learningRate,
70                    float momentum, std::ofstream &f, std::ofstream &g);
71     activation processBatchForward(SpatiallySparseBatch &batch);
72     void processBatchBackward(SpatiallySparseBatch &batch,
73                             float learningRate, float momentum,
74                             const std::vector<float> &dfeatures);
75     void processIndexLearnerBatch(SpatiallySparseBatch &batch, float learningRate,
76                                 float momentum, std::ofstream &f);
77
78     void addLearntLayer(int nFeatures, ActivationFunction activationFn = RELU,
79                      float dropout = 0.0f, float alpha = 1.0f);
80     void addNetworkInNetworkLayer(int nFeatures,
81                                  ActivationFunction activationFn = RELU,
82                                  float dropout = 0.0f);
83     void addConvolutionalLayer(int nFeatures, int filterSize, int filterStride,
84                               ActivationFunction activationFn = RELU,
85                               float dropout = 0.0f, int minActiveInputs = 1,
86                               float poolingToFollow = 1.0f);
87     void addNetLayerMP(int nFeatures, int filterSize, int filterStride

```

```

14 cdef extern from "SparseConvNet/SparseConvNetCUDA.h":
15     cdef cppclass SparseConvNetCUDA:
16         vector[SpatiallySparseLayer*] layers
17         int computeInputSpatialSize(int outputSpatialSize)
18         vector[vector[float]] predict(SpatiallySparseDataset &dataset)
19         vector[activation] layer_activations(SpatiallySparseDataset &dataset)
20         activation processBatchForward(SpatiallySparseBatch &batch)
21         void processBatchBackward(SpatiallySparseBatch &batch,
22                                 float learningRate,
23                                 float momentum,
24                                 vector[float] dfeatures)
25         void addConvolutionalLayer(int nFeatures,
26                                   int filterSize,
27                                   int filterStride,
28                                   ActivationFunction activationFn,
29                                   float dropout,
30                                   int minActiveInputs,
31                                   float poolingToFollow)

```


```

71 cdef class SparseNetwork:
72     """create a network object, configure layers, threads and
73     dimensionality of input
74     """
75     cdef SparseConvNet* net
76     cdef list layers
77     cdef int dimension
78     cdef int nInputFeatures
79     cdef int nClasses
80     cdef int input_spatial_size
81
82     def __cinit__(self, int dimension, int nInputFeatures, int nClasses,
83                 int cudaDevice=-1, int nTop=1, int nThreads=1):
84         """Initialzing Network.
85
86         dimension - number of input dimension
87         nInputFeatures number of features in one cell of the grid
88         """
89         self.layers = []
90         self.net = new SparseConvNet(dimension, nInputFeatures,
91                                     nClasses, cudaDevice, nTop, nThreads)
92         self.dimension = dimension
93         self.nInputFeatures = nInputFeatures
94         self.nClasses = nClasses
95         self.input_spatial_size = -1
96
97     def __dealloc__(self):
98         del self.net


```

Cython


SparseConvNet.pxd
sparseNetwork.pyx




sparseNetwork.cpp




sparseNetwork.o



SparseConvNet/*.cu
SparseConvNet/*.cpp
SparseConvNet/*.h



SparseConvNet/*.o



PySparseConvNet.[so / dylib]

```

421 cdef class PyVoxelPicture:
422     """wraps VoxelPicture
423     """
424     cdef VoxelPicture* pic
425     cdef SparseGrid grid
426     cdef vector[float] features
427     cdef int nSpatialSites
428
429     def __cinit__(self, np.ndarray[long, mode="c", ndim=2] indices,
430                  np.ndarray[double, mode="c", ndim=2] input_features,
431                  int renderSize, int label=-1, int n_features=1):
432
433         self.nSpatialSites = 0
434         self.pic = new VoxelPicture(indices, input_features, renderSize,
435                                     label, n_features)
436
437
438     def __dealloc__(self):
439         del self.pic
440         del self.grid
441         del self.features
442

```

```

59 class TestVoxelPicture(unittest.TestCase):
60
61     def test_constructor_from_row_matrix(self):
62         # indices - an array of shape (num_points, 3),
63         # its columns are indices x,y,z
64         indices = np.array([
65             [0, 0, 0],
66             [1, 0, 5],
67             [3, 4, 2],
68             [5, 5, 5]
69         ], dtype=np.int)
70         # size of 3-d tensor, all sides are equal
71         spatial_size = 6
72         n_features = 1
73         # features of size (num_points, num_features)
74         # in this case num_features=1
75         features = np.ones((indices.shape[0], 1), dtype=np.float)
76         # creating a picture object
77         pic = pscn.PyVoxelPicture(indices, features, spatial_size)
78         # extracting indices and features must be the same
79         # as ones it was created from
80         returned_indices, returned_features = pic.codifyInputData(spatial_size)
81         sparse_indices, sparse_features = convert_pairs_and_features_to_map(
82             returned_indices, returned_features, spatial_size, n_features)

```

Data type	Description
bool_	Boolean (True or False) stored as a byte
int_	Default integer type (same as C long; normally either int64 or int32)
intc	Identical to C int (normally int32 or int64)
intp	Integer used for indexing (same as C ssize_t; normally either int32 or int64)
int8	Byte (-128 to 127)
int16	Integer (-32768 to 32767)
int32	Integer (-2147483648 to 2147483647)
int64	Integer (-9223372036854775808 to 9223372036854775807)
uint8	Unsigned integer (0 to 255)
uint16	Unsigned integer (0 to 65535)
uint32	Unsigned integer (0 to 4294967295)
uint64	Unsigned integer (0 to 18446744073709551615)
float_	Shorthand for float64.
float16	Half precision float: sign bit, 5 bits exponent, 10 bits mantissa
float32	Single precision float: sign bit, 8 bits exponent, 23 bits mantissa
float64	Double precision float: sign bit, 11 bits exponent, 52 bits mantissa

C types	From Python types	To Python types
[unsigned] char	int, long	int
[unsigned] short		
int, long		
unsigned int	int, long	long
unsigned long		
[unsigned] long long		
float, double, long double	int, long, float	float
char *	str/bytes	str/bytes [1]
struct		dict


Shape Retrieval

Problem statement


Given a query object find several the most “similar” to the query objects from the given database.

The objects are considered to be similar if they belong to the same category of objects and have similar shapes.

Shape Retrieval



Triplet loss



The representation can be efficiently learned by minimizing triplet loss.

Triplet is a set (a, p, n) , where

- a - anchor object
- p - positive object that is similar to anchor object
- n - negative object that is not similar to anchor object

$$\lambda(\delta_+, \delta_-) = \max(\mu + \delta_+ - \delta_-)$$


where μ is a margin parameter, δ_+ and δ_- are distances between p and a and n and a .

Network description


layer type	size	stride	channels	spatial size	sparsity (%) ¹	MRI sparsity ²
Data input	-	-	1	126	0.18	9.76
Sparse Convolutional Layer	2	1	8	125	-	-
Leaky ReLU ($\alpha = 0.33$)	-	-	32	125	0.35	10.75
Sparse MaxPooling Layer	3	2	32	62	0.69	12.59
Sparse Convolutional Layer	2	1	256	61	-	-
Leaky ReLU ($\alpha = 0.33$)	-	-	64	61	1.07	14.94
Sparse MaxPooling Layer	3	2	64	30	1.93	19.62
Sparse Convolutional Layer	2	1	512	29	-	-
Leaky ReLU ($\alpha = 0.33$)	-	-	96	29	3.26	26.49
Sparse MaxPooling Layer	3	2	96	14	7.32	41.47
Sparse Convolutional Layer	2	1	768	13	-	-
Leaky ReLU ($\alpha = 0.33$)	-	-	128	13	15.14	64.55
Sparse MaxPooling Layer	3	2	128	6	46.30	95.21
Sparse Convolutional Layer	2	1	1024	5	-	-
Leaky ReLU ($\alpha = 0.33$)	-	-	160	5	97.54	100.00
Sparse MaxPooling Layer	3	2	160	2	100.00	100.00
Sparse Convolutional Layer	2	1	1280	1	-	-
Leaky ReLU ($\alpha = 0.33$)	-	-	192	1	100.00	100.00

¹column "sparsity" is computed for render size = **40**




²column "MRI sparsity" is computed for render size = **110**








Forward Pass Activations








Input



Layer 2



Layer 3



Layer 11




Layer 15

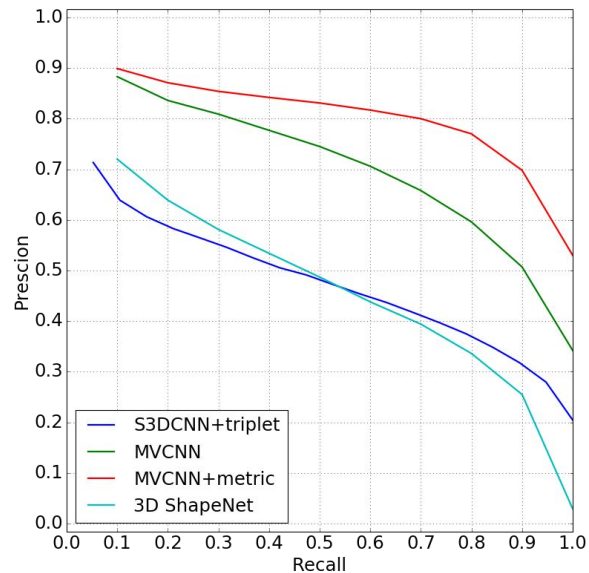


Layer 17

Experimental results



Optimisation algorithm:
Nesterov Accelerated Gradient:
momentum = 0.99
Constant Learning Rate = 0.002



method	Classification	Retrieval mAP
3DShapeNet	77.32%	49.23%
MVCNN	90.10%	80.20%
3DSCNN	90.3%	45.16%
S3DCNN + triplet	---	46.71%

Conclusions

- For Shape Datasets in voxel form - resolution beyond 30^3 doesn't improve performance very much
- More voxels - change scale of features, probably needs more layers
- 3D CNNs are more efficient on volumetric data

Organizing hyperparameter search

Retrieval_Experiments																				avnotcher
File Edit View Insert Format Data Tools Add-ons Help Last edit was on 9 December 2016																				Comments
dataset																				
	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
1	dataset	net_arch	batch_size	unique_c	test_even	pair_taking_method	lr_decay_rate	render_in	batch_sam	margin	norm	exp_hash	server	status	current_ep	mAP	last_update	links	total number of	
2	ShapeNet55	deepC2	252	6	100	lazy_almost_shuffled_pe	0.005	32	FALSE	0.1	cosine	aaaa111vvvv	burn	stopped	828	0.3963241766	Mon Sep 26 17:57:04 2016	png_plot	2.09E+07	
3	ModelNet40	deepC2	252	6	100	lazy_almost_shuffled_pe	0.005	32	FALSE	0.1	cosine	adfcndd33nd	labten-01	done				png_plot	0.00E+00	
4	ModelNet40	deepC2	60	4	100	random_subsampling	0.005	40	FALSE	0.1	cosine	bbbb322adfd	iitp-01	done				png_plot	0.00E+00	
5	ModelNet40	deepC2	270	6	100	lazy_almost_shuffled_pe	0.005	60	FALSE	0.1	cosine	cccc12515af	titan	done				png_plot	0.00E+00	
6	ModelNet40	deepC2	180	5	200	random_subsampling	0.005	50	FALSE	0.1	cosine	dddd34252gd	labten-02	stopped	286	0.3312430251	Wed Sep 28 02:16:13 2016	png_plot	1.03E+07	
7	ModelNet40	deepC2	90	5	200	lazy_almost_shuffled_pe	0.005	70	FALSE	1	L2	fsfwr3311v	titan	stopped	244	0.2998781443	Wed Sep 28 02:10:57 2016	png_plot	4.39E+06	
8	ModelNet40	deepC2	90	3	200	lazy_almost_shuffled_pe	0.005	32	FALSE	1	cosine	mZLKABODnxl	burn	stopped	510	0.2110905835	Tue Sep 27 16:43:21 2016	png_plot	9.18E+06	
9	ModelNet40	deepC2	90	6	200	random_subsampling	0.005	45	FALSE	1	L2	mdqY14gfwXu	deepburn	done	639	0.3035397731		png_plot	1.15E+07	
10	ModelNet40	deepC2	90	6	200	random_subsampling	0.005	60	FALSE	1	L2	aEKclGUNWX4	labten-01	stopped	335	0.3137650743		png_plot	6.03E+06	
11	ModelNet40	deepC2	90	5	200	random_subsampling	0.005	55	FALSE	0.5	cosine	5ciPqKGBUHE	deepburn	stopped	368	0.2896838771		png_plot	6.62E+06	
12	ModelNet40	deepC2	120	10	800	random_subsampling	0.002	32	FALSE	1	cosine	HyBxR0QAZkh	burn	stopped	104	0.1979471797	Wed Oct 5 16:36:27 2016	png_plot	9.98E+06	
13	ModelNet40	deepC2	90	10	800	random_subsampling	0.002	40	FALSE	1	cosine	PaknYB9i3R0	iitp-01	stopped	124	0.2085503059	Sat Oct 1 04:45:20 2016	png_plot	8.93E+06	
14	ModelNet40	deepC2	72	8	800	random_subsampling	0.002	45	FALSE	1	cosine	0fWOUruVEaz	iitp-03	done	199	0.2303943628	Sun Oct 2 02:51:04 2016	png_plot	1.15E+07	
15	ModelNet40	deepC2	120	5	400	random_subsampling	0.004	50	FALSE	1	cosine	AEWxBRuSJ4n	burn	stopped	226	0.2542267267	Tue Oct 11 19:54:57 2016	png_plot	1.08E+07	
16	ModelNet40	deepC2	90	10	800	random_subsampling	0.002	55	FALSE	1	cosine	87ALzZ1QwBK	deepburn	done	159	0.2326888225	Mon Oct 3 13:47:29 2016	png_plot	1.14E+07	
17	ModelNet40	deepC2	120	10	800	random_subsampling	0.002	60	FALSE	1	cosine	tew16fZSoyl	labten-02	stopped	85	0.2235581123	Sat Oct 1 21:22:33 2016	png_plot	8.16E+06	
18	ModelNet40	deepC2	120	10	800	random_subsampling	0.002	65	FALSE	1	cosine	EvO6ftwRjSl	labten-01	stopped	58	0.2526285167	Fri Sep 30 13:11:04 2016	png_plot	5.57E+06	
19	ModelNet40	deepC2	90	10	800	random_subsampling	0.002	70	FALSE	1	cosine	9KCNSE4huYB	deepburn	done	159	0.233684883	Tue Oct 4 02:47:30 2016	png_plot	1.14E+07	
20	ModelNet40	deepC2	90	10	800	random_subsampling	0.002	75	FALSE	1	cosine	9nhgJRCoO0Q	iitp-04	stopped	78	0.2224471503	Sat Oct 1 01:47:54 2016	png_plot	5.62E+06	
21	ModelNet40	deepC2	120	10	800	random_subsampling	0.002	80	FALSE	1	cosine	K7IG4U3cj5V	titan	done	119	0.2205873523	Tue Oct 4 12:01:47 2016	png_plot	1.14E+07	
22	ModelNet40	deepC2	120	10	800	random_subsampling	0.002	25	FALSE	1	cosine	OclCWJjVLmb	iitp-03	stopped	86	0.19171750831	Tue Oct 4 10:53:00 2016	png_plot	8.26E+06	
23	ShapeNet55	deepC2	150	10	600	random_subsampling	0.002	32	FALSE	1	cosine	AksKiQrRh8	burn	stopped	20	0.2578580269	Thu Sep 29 15:55:45 2016	png_plot	1.80E+06	
24	ModelNet40	deepC2_wide	90	5	500	shuffled_permutations	["lr_step": 5, "lr_base": 0.01, "lr_decay": 0.66, "m_step": 12, "m_base": 0.5, "m_inc": 0.1, "m_max": 0.99]	50	FALSE	2	L2	llbtmHoQ9UF	titan	stopped	142	0.2389334476	Fri Oct 14 12:28:42 2016	png_plot	6.39E+06	
25	ModelNet40	deepC2_wide	90	5	500	shuffled_permutations	["lr_step": 5, "lr_base": 0.01, "lr_decay": 0.66, "m_step": 20, "m_base": 0.7, "m_inc": 0.1, "m_max": 0.99]	50	FALSE	2	L2	pHEKMNesSJ38	titan	stopped	139	0.237851468	Fri Oct 14 12:24:56 2016	png_plot	6.26E+06	
26	ModelNet40	deepC2_wide	90	5	500	shuffled_permutations	["lr_step": 5, "lr_base": 0.01, "lr_decay": 0.5, "m_step": 20, "m_base": 0.5, "m_inc": 0.1, "m_max": 0.99]	50	FALSE	2	L2	wFnYuDixHdZ	burn	stopped	117	0.2495343194	Fri Oct 14 10:52:24 2016	png_plot	5.27E+06	
27	ModelNet40	deepC2_wide	90	5	500	shuffled_permutations	["lr_step": 10, "lr_base": 0.003, "lr_decay": 0.5, "m_step": 20, "m_base": 0.5, "m_inc": 0.1, "m_max": 0.99]	50	FALSE	2	L2	cMR2RzgeVqW	labten-01	stopped	107	0.2529385175	Fri Oct 14 12:09:14 2016	png_plot	4.82E+06	


Please contribute:

<https://github.com/gangiman/PySparseConvNet>

GPLv3

ADISE@Skoltech

Special thanks to:



Alexandr Notchenko, Ermek Kapushev, Evgeny Burnaev

Dmitry Yarotsky

Rasim Akhunzyanov

From
OpenDataScience
with Love

