AVANTIX

Avancement Simulateur IQ

AVANTIX

Tensor-RT

10/01/2024

Tensor-RT

Pytorch → ONNX → TensorRT

Résultats (CBE/segmentation/_TRT/*.ipynb):

Conv(32)+Amax
RTR = 4,39 → RTR = 3,87

Gain de ~12%

Unet(8)
RTR = 17 → RTR = 16

Gain de ~7%

Gain variable selon le type d'architecture

AVANTIX

Segmentation : Modèle bloc 2x2

10/01/2024

Conv + Amax: k=3, s=2, p=1

Tous les npersegs, sans overlap
Sxx_tiles (8192, 1, height=128, width=512)
Correspond à un dt_iq de 16,78 ms

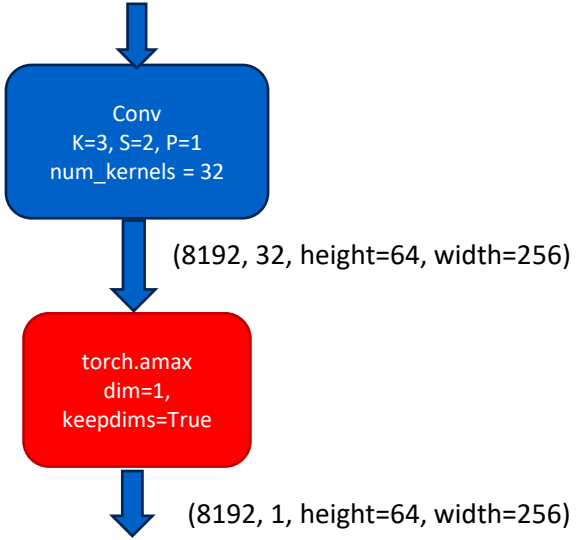

RTR = 4,4

Résultats à 32 channels

```
Result for SNR=-3
        recall = 0.00% | precision = 6.56% | fdr = 93.44% | pfa = 0.00%
Result for SNR=0
        recall = 0.00% | precision = 33.71% | fdr = 66.29% | pfa = 0.00%
Result for SNR=3
        recall = 0.21% | precision = 97.17% | fdr = 2.83% | pfa = 0.00%
Result for SNR=6
        recall = 39.00% | precision = 99.98% | fdr = 0.02% | pfa = 0.00%
Result for SNR=9
        recall = 81.99% | precision = 99.98% | fdr = 0.02% | pfa = 0.00%
Result for SNR=12
        recall = 96.99% | precision = 97.41% | fdr = 2.59% | pfa = 0.04%
Result for SNR=15
        recall = 100.00% | precision = 87.57% | fdr = 12.43% | pfa = 0.20%
Result for SNR=18
        recall = 100.00% | precision = 80.29% | fdr = 19.71% | pfa = 0.35%
Result for SNR=21
        recall = 100.00% | precision = 80.22% | fdr = 19.78% | pfa = 0.34%
```

Résultats à 24 channels

```
Result for SNR=-3
        recall = 0.00% | precision = 7.50% | fdr = 92.50%
Result for SNR=0
        recall = 0.00% | precision = 38.18% | fdr = 61.82%
Result for SNR=3
        recall = 0.24% | precision = 96.38% | fdr = 3.62%
Result for SNR=6
        recall = 39.66% | precision = 99.98% | fdr = 0.02%
Result for SNR=9
        recall = 82.18% | precision = 99.98% | fdr = 0.02%
Result for SNR=12
        recall = 97.23% | precision = 97.27% | fdr = 2.73%
Result for SNR=15
        recall = 100.00% | precision = 87.59% | fdr = 12.41%
Result for SNR=18
        recall = 100.00% | precision = 80.06% | fdr = 19.94%
Result for SNR=21
        recall = 100.00% | precision = 80.40% | fdr = 19.60%
```

Perf stable

Résultats à 16 channels

```
Result for SNR=-3
        recall = 0.00% | precision = 6.67% | fdr = 93.33%
Result for SNR=0
        recall = 0.00% | precision = 36.73% | fdr = 63.27%
Result for SNR=3
        recall = 0.22% | precision = 96.27% | fdr = 3.73%
Result for SNR=6
        recall = 38.87% | precision = 99.98% | fdr = 0.02%
Result for SNR=9
        recall = 81.85% | precision = 99.98% | fdr = 0.02%
Result for SNR=12
        recall = 97.12% | precision = 97.35% | fdr = 2.65%
Result for SNR=15
        recall = 100.00% | precision = 87.64% | fdr = 12.36%
Result for SNR=18
        recall = 100.00% | precision = 80.06% | fdr = 19.94%
Result for SNR=21
        recall = 100.00% | precision = 80.40% | fdr = 19.60%
```

Perf stable

Résultats à 8 channels

```
Result for SNR=-3
        recall = 0.00% | precision = 9.52% | fdr = 90.48%
Result for SNR=0
        recall = 0.00% | precision = 30.00% | fdr = 70.00%
Result for SNR=3
        recall = 0.11% | precision = 95.83% | fdr = 4.17%
Result for SNR=6
        recall = 33.86% | precision = 99.99% | fdr = 0.01%
Result for SNR=9
        recall = 79.71% | precision = 99.99% | fdr = 0.01%
Result for SNR=12
        recall = 96.58% | precision = 97.78% | fdr = 2.22%
Result for SNR=15
        recall = 100.00% | precision = 87.83% | fdr = 12.17%
Result for SNR=18
        recall = 100.00% | precision = 80.06% | fdr = 19.94%
Result for SNR=21
        recall = 100.00% | precision = 80.40% | fdr = 19.60%
```

Perf stable

Résultats à 4 channels

```
Result for SNR=-3
        recall = 0.00% | precision = 7.59% | fdr = 92.41%
Result for SNR=0
        recall = 0.00% | precision = 38.18% | fdr = 61.82%
Result for SNR=3
        recall = 0.23% | precision = 96.43% | fdr = 3.57%
Result for SNR=6
        recall = 39.51% | precision = 99.98% | fdr = 0.02%
Result for SNR=9
        recall = 82.11% | precision = 99.98% | fdr = 0.02%
Result for SNR=12
        recall = 97.21% | precision = 97.29% | fdr = 2.71%
Result for SNR=15
        recall = 100.00% | precision = 87.60% | fdr = 12.40%
Result for SNR=18
        recall = 100.00% | precision = 80.06% | fdr = 19.94%
Result for SNR=21
        recall = 100.00% | precision = 80.40% | fdr = 19.60%
```

Perf stable

Résultats à 2 channels

```
Result for SNR=-3
        recall = 0.00% | precision = 7.59% | fdr = 92.41%
Result for SNR=0
        recall = 0.00% | precision = 37.04% | fdr = 62.96%
Result for SNR=3
        recall = 0.23% | precision = 96.43% | fdr = 3.57%
Result for SNR=6
        recall = 39.42% | precision = 99.98% | fdr = 0.02%
Result for SNR=9
        recall = 82.07% | precision = 99.98% | fdr = 0.02%
Result for SNR=12
        recall = 97.20% | precision = 97.31% | fdr = 2.69%
Result for SNR=15
        recall = 100.00% | precision = 87.61% | fdr = 12.39%
Result for SNR=18
        recall = 100.00% | precision = 80.06% | fdr = 19.94%
Result for SNR=21
        recall = 100.00% | precision = 80.40% | fdr = 19.60%
```

Perf stable

Résultats à 1 channel

```
Result for SNR=-3
        recall = 0.00% | precision = 7.41% | fdr = 92.59%
Result for SNR=0
        recall = 0.00% | precision = 38.39% | fdr = 61.61%
Result for SNR=3
        recall = 0.24% | precision = 96.43% | fdr = 3.57%
Result for SNR=6
        recall = 39.77% | precision = 99.98% | fdr = 0.02%
Result for SNR=9
        recall = 82.22% | precision = 99.98% | fdr = 0.02%
Result for SNR=12
        recall = 97.25% | precision = 97.26% | fdr = 2.74%
Result for SNR=15
        recall = 100.00% | precision = 87.58% | fdr = 12.42%
Result for SNR=18
        recall = 100.00% | precision = 80.06% | fdr = 19.94%
Result for SNR=21
        recall = 100.00% | precision = 80.40% | fdr = 19.60%
```

Perf stable

MaxPool2D

Tous les npersegs, sans overlap
Sxx_tiles (8192, 1, height=128, width=512)
Correspond à un dt_iq de 16,78 ms

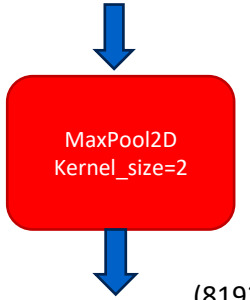

(8192, 1, height=64, width=256)

RTR = 4,4

Résultats avec MaxPool2D

```
Result for SNR=-3
        recall = 0.01% | precision = 1.77% | fdr = 98.23%
Result for SNR=0
        recall = 0.02% | precision = 3.37% | fdr = 96.63%
Result for SNR=3
        recall = 0.05% | precision = 7.11% | fdr = 92.89%
Result for SNR=6
        recall = 0.32% | precision = 34.32% | fdr = 65.68%
Result for SNR=9
        recall = 15.61% | precision = 96.20% | fdr = 3.80%
Result for SNR=12
        recall = 100.00% | precision = 99.41% | fdr = 0.59%
Result for SNR=15
        recall = 100.00% | precision = 99.41% | fdr = 0.59%
Result for SNR=18
        recall = 100.00% | precision = 99.36% | fdr = 0.64%
Result for SNR=21
        recall = 100.00% | precision = 99.38% | fdr = 0.62%
```

Forte baisse de perf à SNR=9dB sinon tjrs ok à SNR>=12dB

Conv + Amax : k=8

Tous les npersegs, sans overlap
Sxx_tiles (8192, 1, height=128, width=512)
Correspond à un dt_iq de 16,78 ms

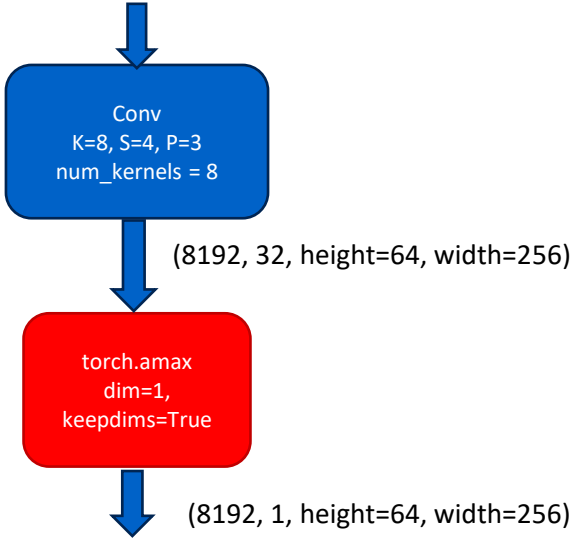

Résultats avec Conv + Amax ks=8

```
Result for SNR=-3
        recall = 0.04% | precision = 4.88% | fdr = 95.12%
Result for SNR=0
        recall = 0.14% | precision = 17.15% | fdr = 82.85%
Result for SNR=3
        recall = 2.55% | precision = 78.54% | fdr = 21.46%
Result for SNR=6
        recall = 52.24% | precision = 98.74% | fdr = 1.26%
Result for SNR=9
        recall = 83.59% | precision = 99.17% | fdr = 0.83%
Result for SNR=12
        recall = 100.00% | precision = 99.34% | fdr = 0.66%
Result for SNR=15
        recall = 100.00% | precision = 99.32% | fdr = 0.68%
Result for SNR=18
        recall = 100.00% | precision = 99.27% | fdr = 0.73%
Result for SNR=21
        recall = 100.00% | precision = 98.57% | fdr = 1.43%
```

Amélioration sur snr=6db

On passe de 39% à 52% en recall

UNet-Like

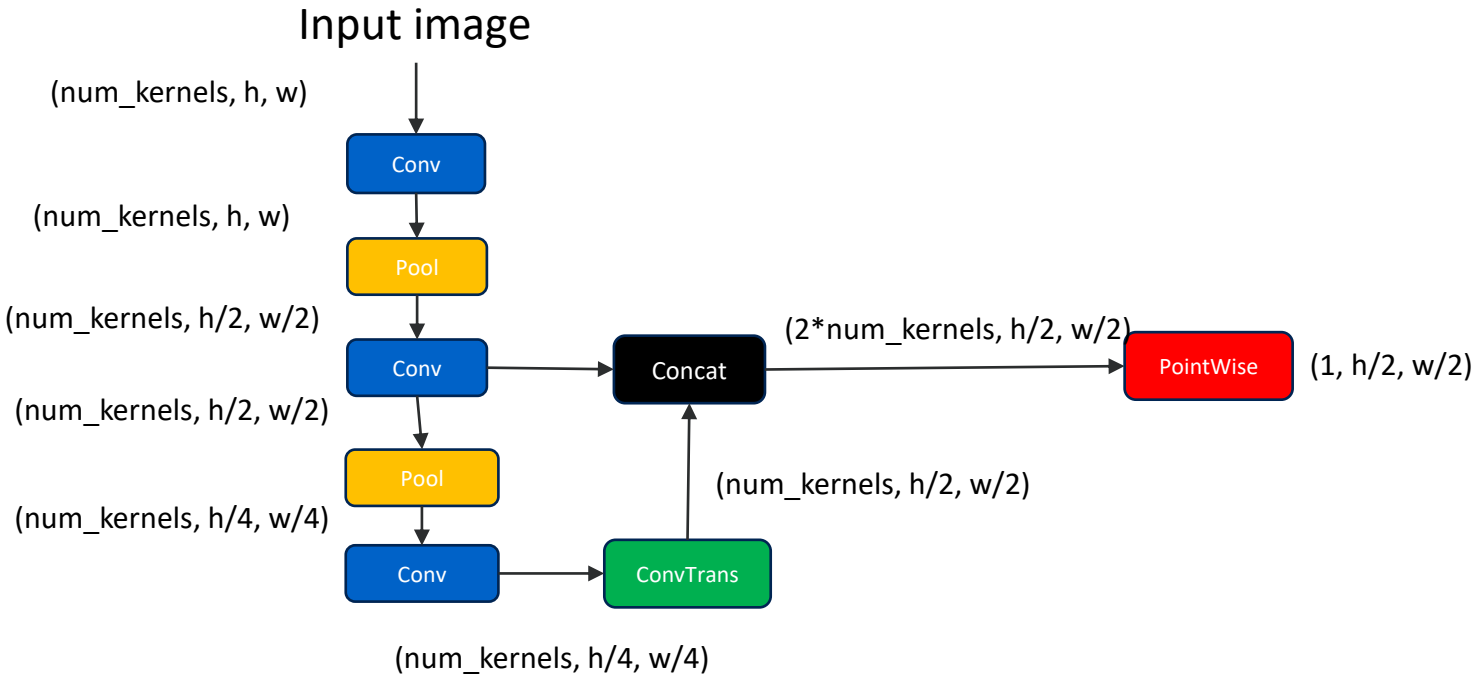
Input image
(num_kernels, h, w)
Conv
(num_kernels, h, w)
Pool
(num_kernels, h/2, w/2)
Conv
(num_kernels, h/2, w/2)
Pool
(num_kernels, h/4, w/4)
Conv
ConvTrans
(num_kernels, h/4, w/4)
Concat
(num_kernels, h/2, w/2)
(2*num_kernels, h/2, w/2)
PointWise
(1, h/2, w/2)

Résultats num_kernels= 8

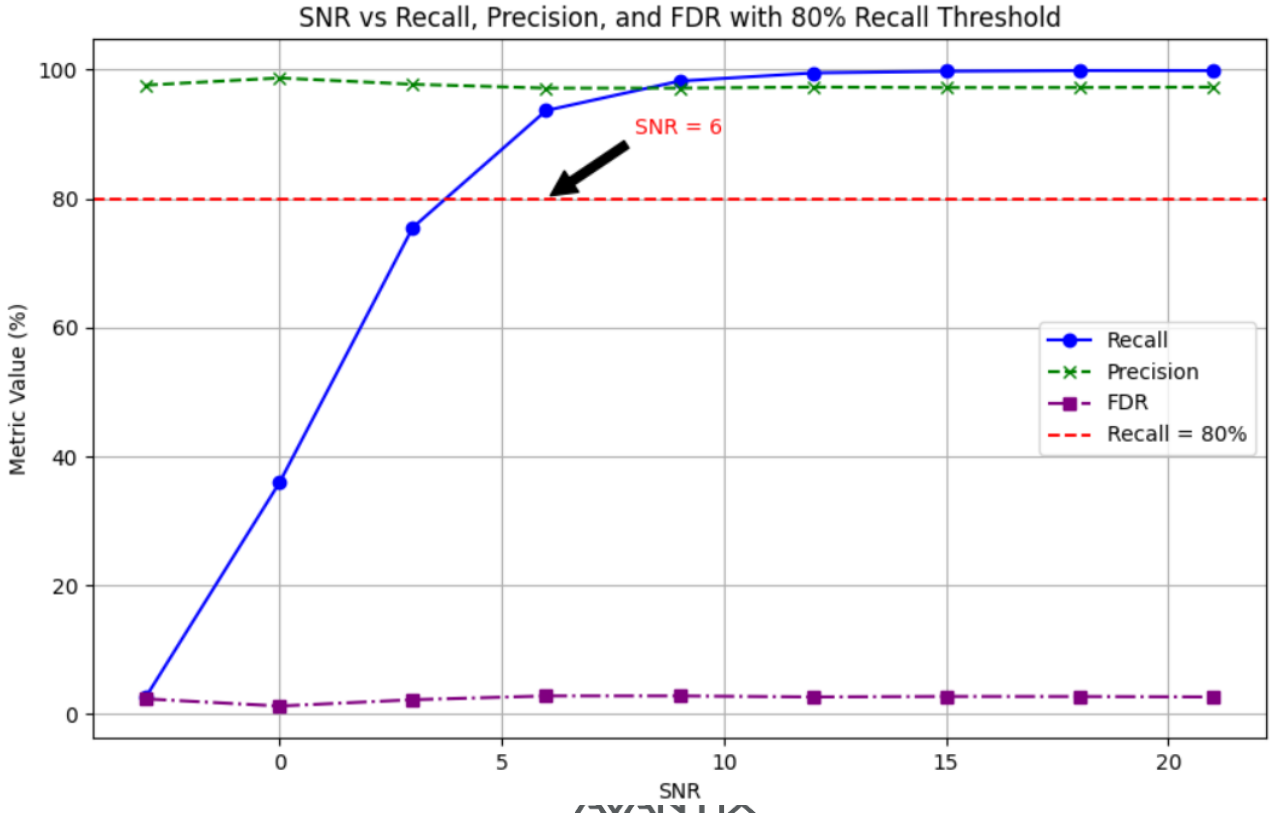

Résultats num_kernels= 32

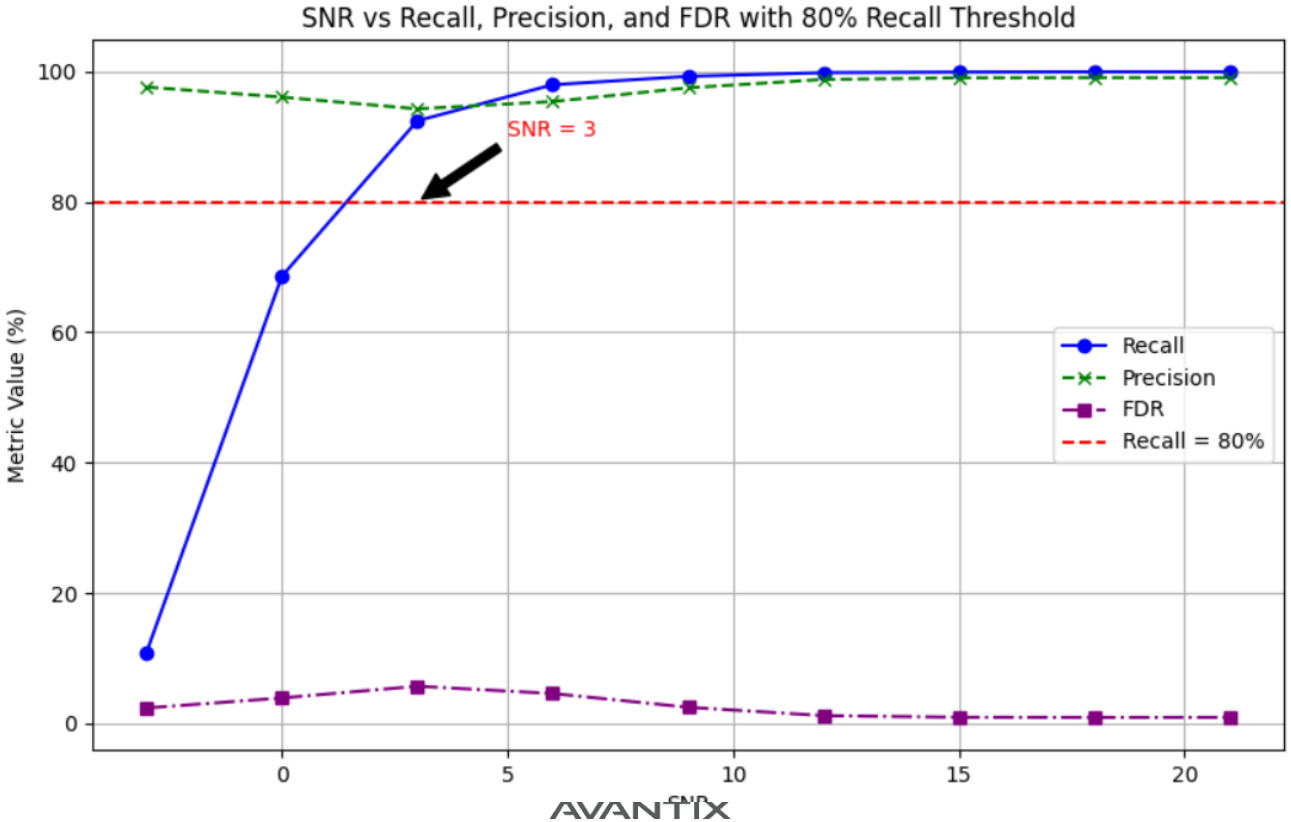

Résultats num_kernels= 64

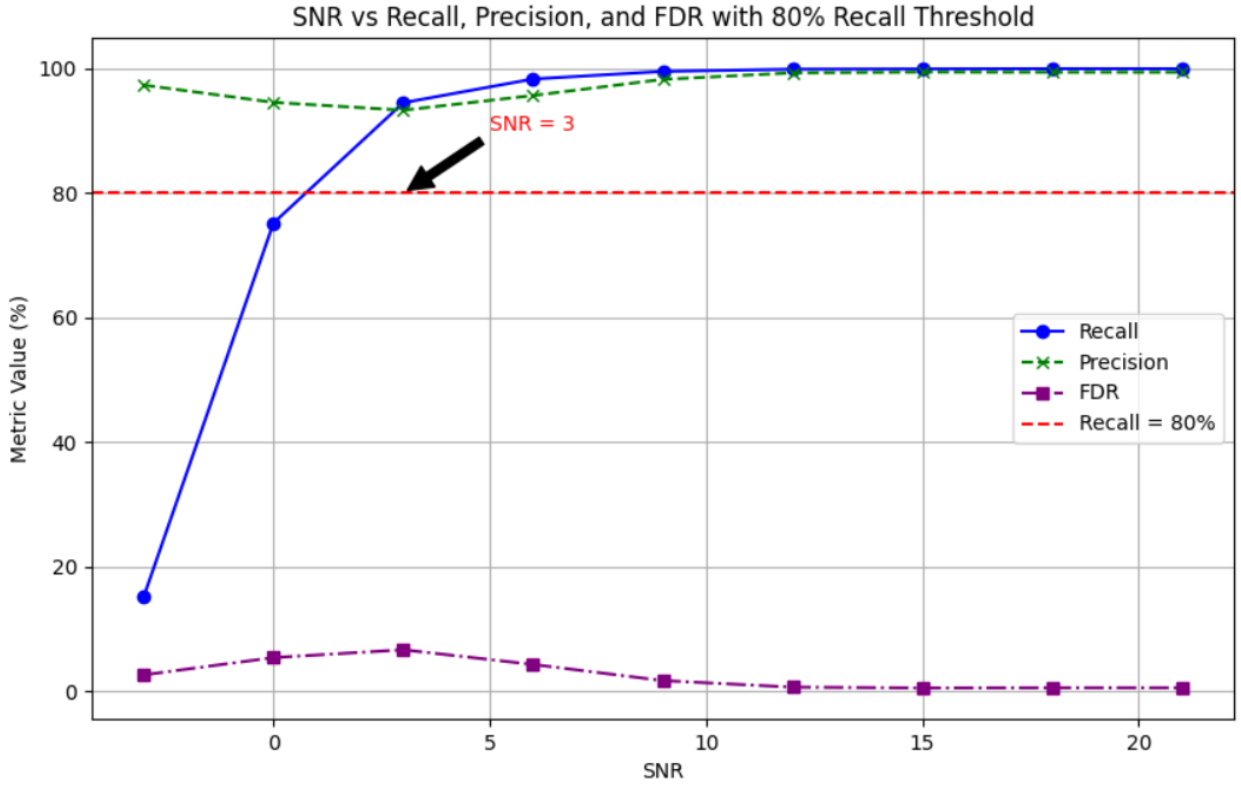

AVANTIX

Résultats num_kernels= 128

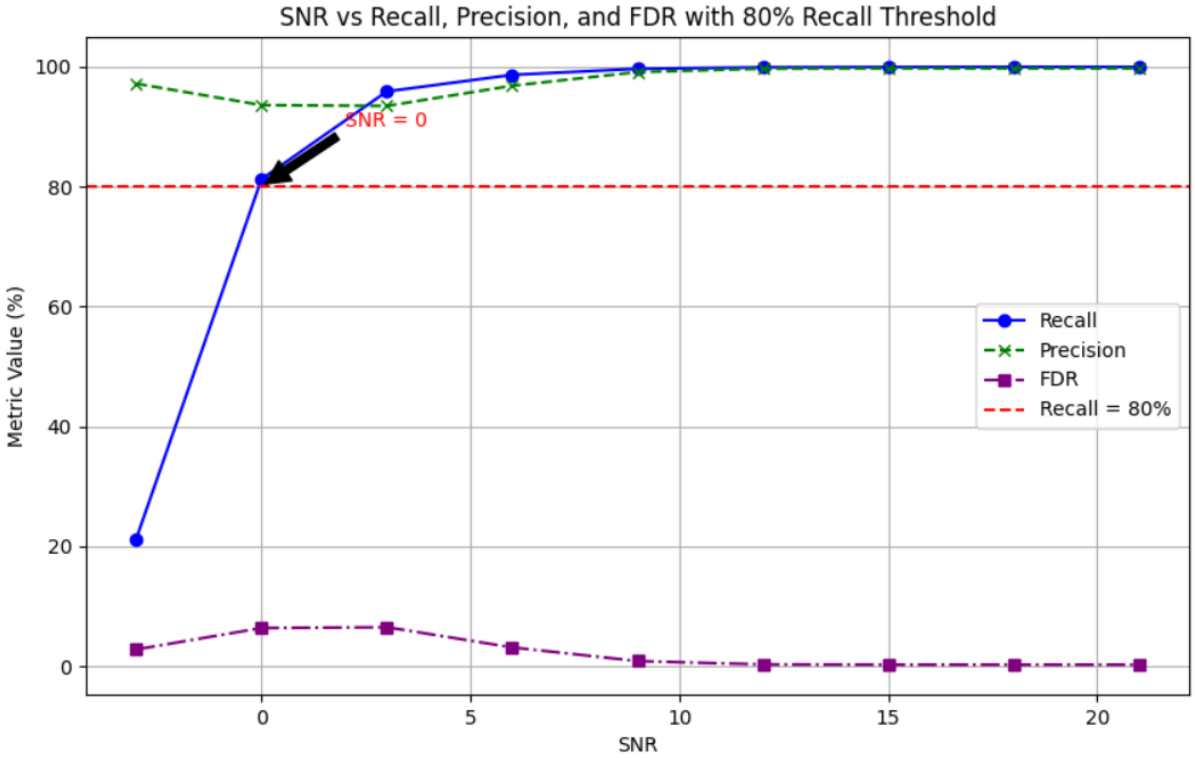

Résultats num_kernels= 256

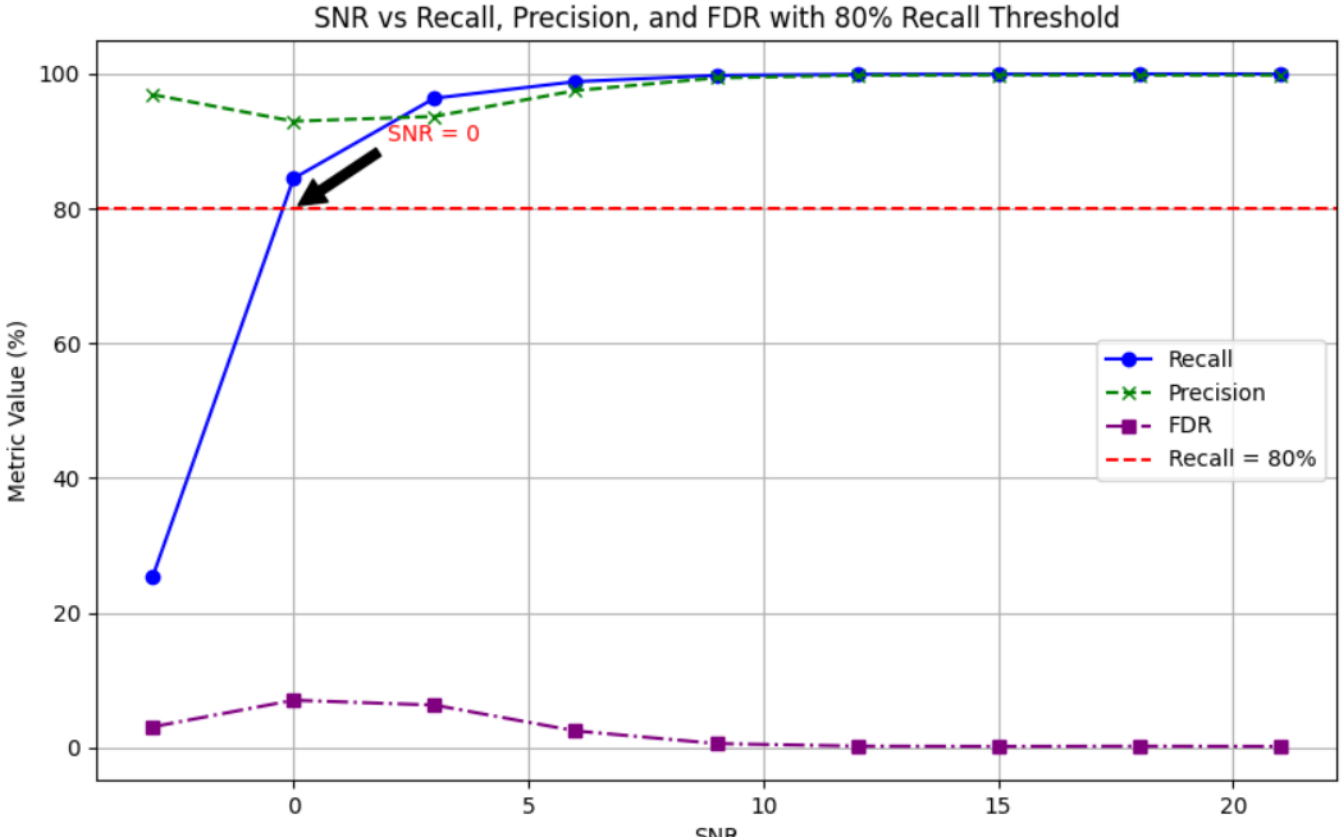

Conclusion

- A 12dB, un MaxPool2D suffit avec un **RTR = 0,34 @ 2% gpu util**
- A 9dB, Conv+Amax avec un 1 seul kernel suffirait avec un **RTR = 0,39@39% gpu util**
- A SNR < 9dB, il faut augmenter le nombre de couches avec une architecture type « Unet » :
  - A 6dB num_kernels = 8 | **RTR = 17@100% gpu util**
  - A 3dB (limite YOLOv8) num_kernels = 32 | **RTR = 49@68% gpu util**
  - A 0dB num_kernels = 128 | **RTR = 279 @ 64% gpu util**
- TODO Conv+Amax : Voir le problème de fdr élevé sur SNR > 15dB

AVANTIX

Régression angle : Modèle bloc 8x8

10/01/2024

Segmentation + Angle

Principe :

1. Segmenter l'image et inférer l'angle sur chaque bloc 8x8
2. Calculer le barycentre sur chaque bloc 8x8
3. Sur chaque bloc 8x8 ayant du signal : générer une obb centrée sur le barycentre et orienté de l'angle inféré
4. Enchainer avec les algos de merge intra et inter

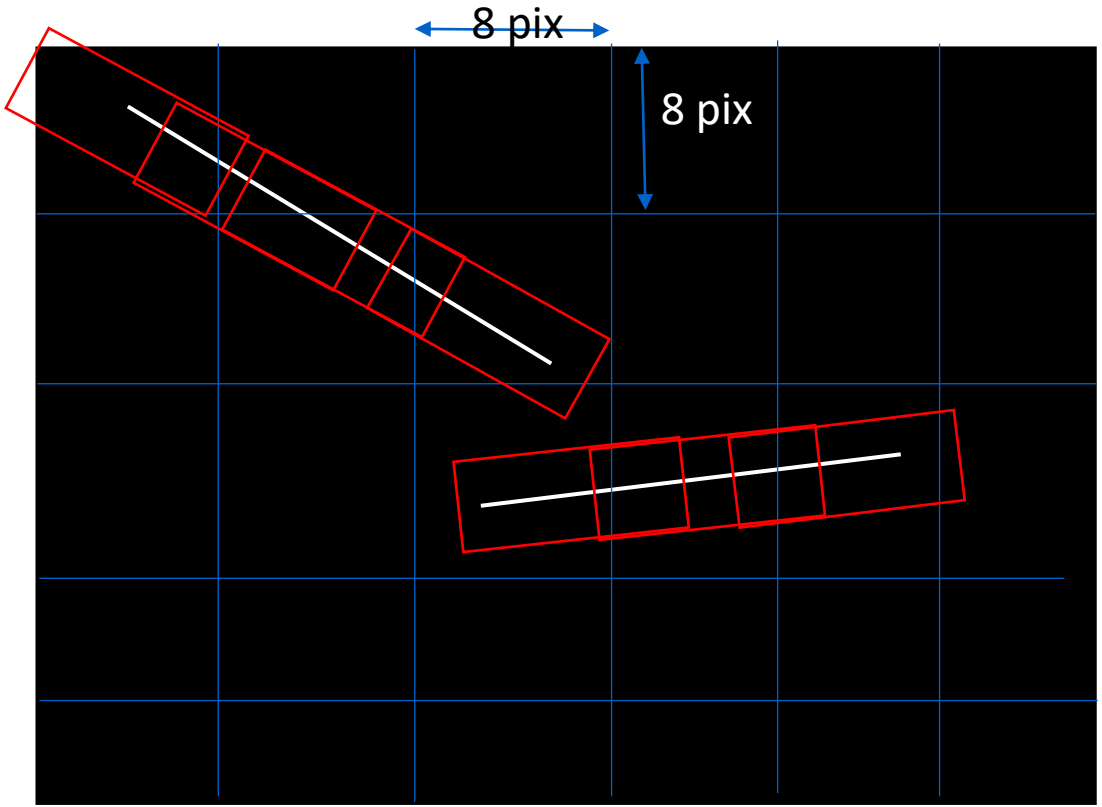

Résultats apprentissage bloc 8x8

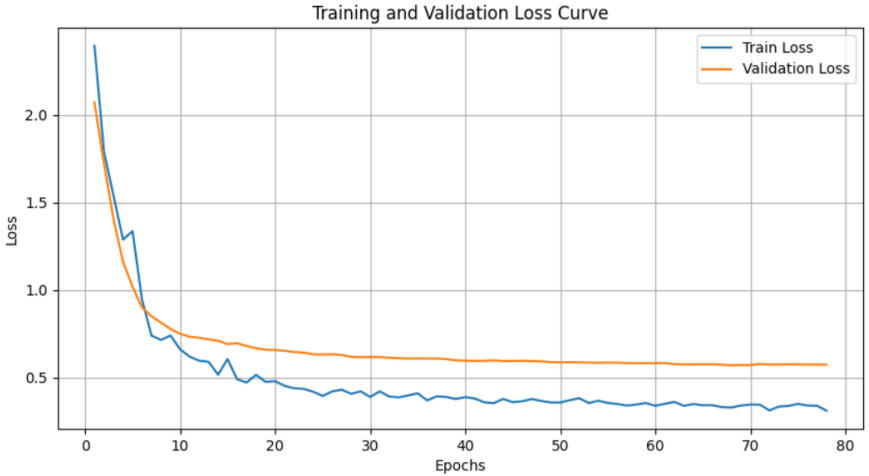

Conv + Amax

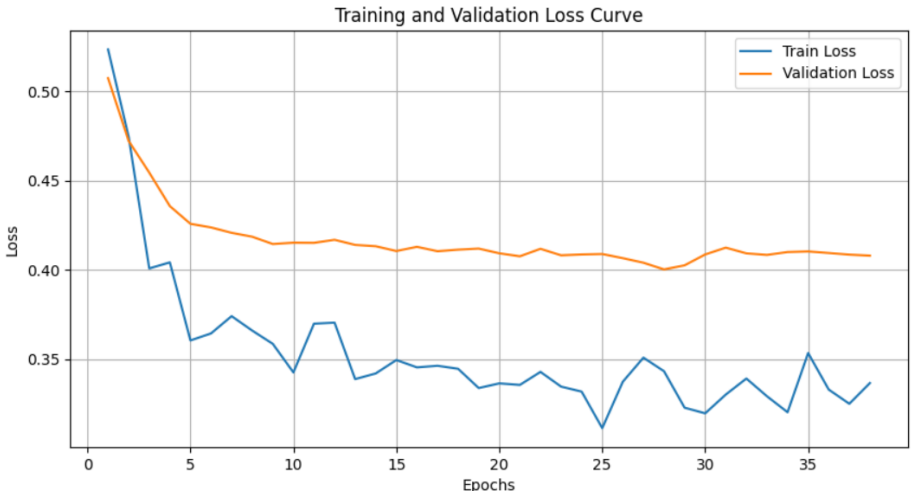

Conv + PW

Résultats préliminaires mauvais. Essayer de voir s'il est possible de réduire l'erreur sans grossir le modèle

AVANTIX

Segmentation + Algo custom

10/01/2024

Principe

1. Segmenter par bloc NxN
2. Utiliser un algo custom pour regrouper les pixels appartenant à la même impulsion

Algo custom en cours d'implémentation :

1. Calculer les pentes et ordonnée à l'origine 2 à 2 de chaque pixel segmenté
2. Ordonner les pentes et faire un diff ce qui va créer des tronçons de zéro
3. Moyenner les pentes et OàO sur les indices appartenant au même tronçon de zéro
4. Rattacher les points ayant même pentes et OàO

AVANTIX

MERCI