

Objective

Speed-up the convolution of Haar-like features (HLFs) for computer vision applications

Approach

Classical convolution: read the values at corners
⊗ Redundant I/Os

Bottleneck: data transfer
☞ Reduce I/O !

Evaluation

Example configurations:

- Speeded-Up-Robust-Features (SURF)
- Initial HLF for OpenCV face detection (FACE)

Theoretical acceleration:

$$\frac{I_{class} + 1}{I_{prop} + O_{prop} + 1} \quad \text{where}$$

- I_{class} = number of inputs per pixel of the classical method,
- I_{prop} , O_{prop} = number of I/O for our method.


Measured acceleration:

Exam.	Meth.	I/O	t (ms)	Acceleration	
				Theor.	Meas.
SURF	Class.	32/0	664	1.65	1.63
	Prop.	14/5	407		
FACE	Class.	22/0	409	1.35	1.36
	Prop.	13/3	300		


Haar-like features in scanning window

Input :
2nd order derivatives of HLFs

Example configuration




Derivates of HLF



Decomposition into simpler kernels

Decompose HLFs into horizontal k_x and vertical k_y kernels forming convolution passes.


Kernels: $k_x * k_x * k_y * k_y$



Assignment of kernels into passes

Permute kernels to convolution passes and order them into trees by testing if simpler kernels are shared across features.


Pass: 1 2 3 4



Alignment of kernels within passes


Align kernels across features so that their inputs always coincide in at least one position.

Pass: p



Best ensemble of kernel trees

Keep the ensemble which yields minimum I/O.



Implementation of best ensemble: B-channel buffer

Store the outputs of the ensemble in an array of contiguous elements.

☞ efficient memory caching

