



Fast stixel computation for fast pedestrian detection

R. Benenson, M. Mathias, R. Timofte and L. Van Gool



We want *very* fast pedestrian detections



In this paper,
we reach *high quality*
pedestrian detection *at 165 Hz*

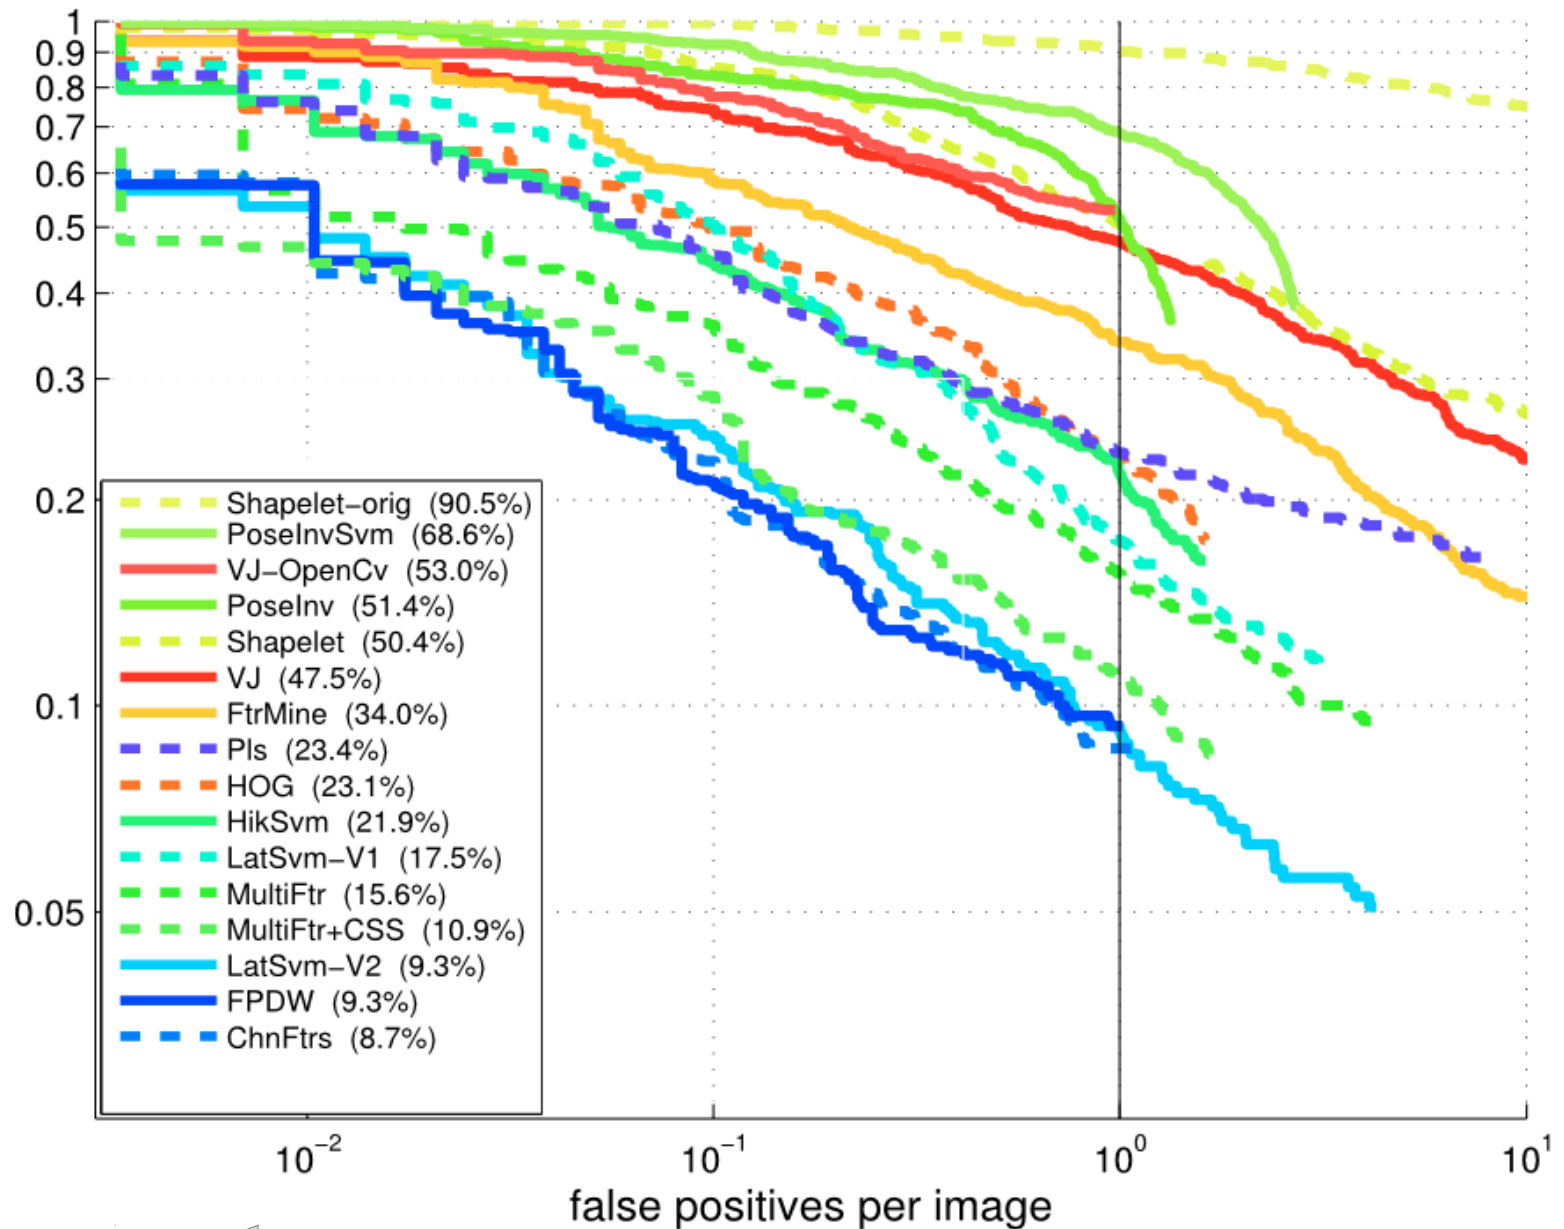


Why 165 Hz ?

- Detection is one module amongst many
- Less computational power
- *Latency matters*



INRIA dataset

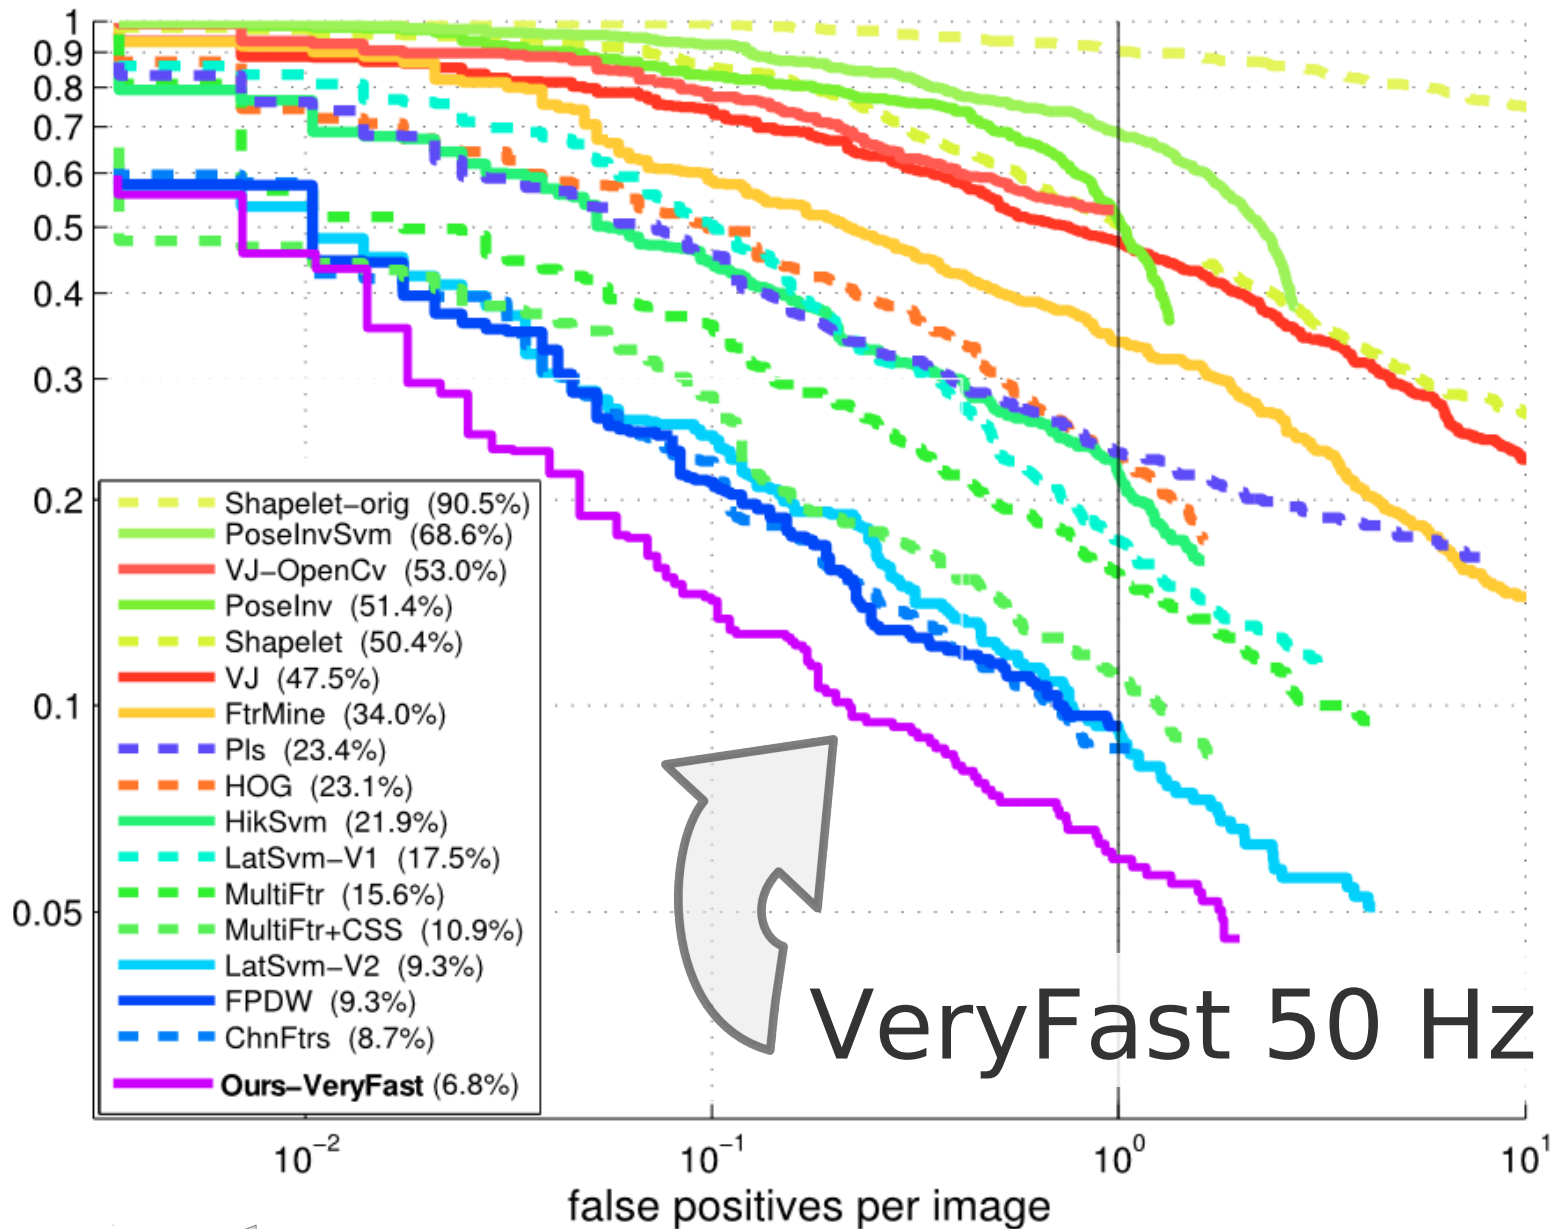


Better

Better

[Dollar et al. 2011]

INRIA dataset



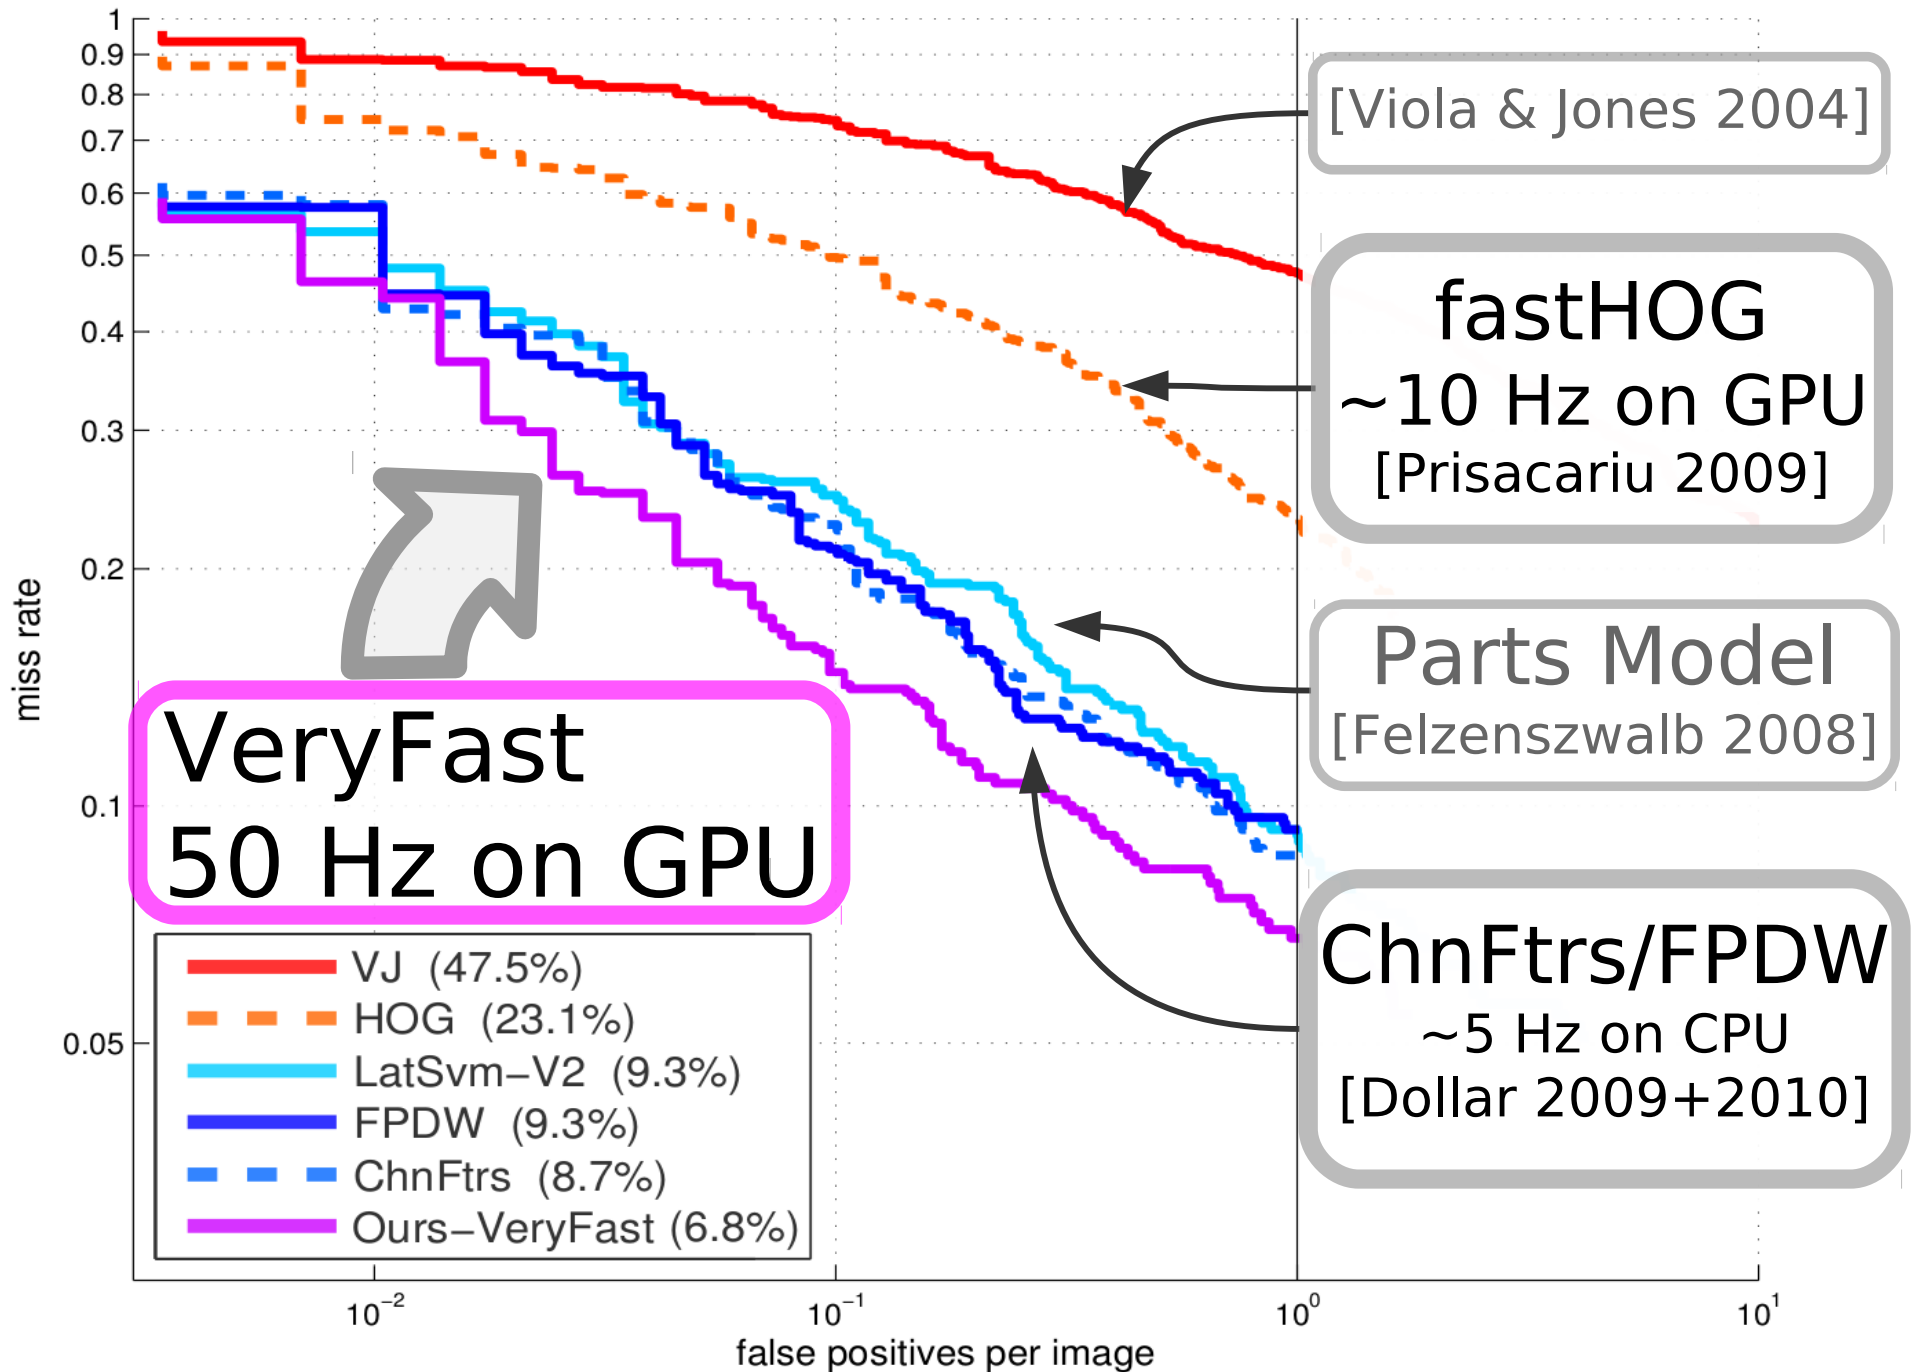
Better

Better


VeryFast 50 Hz

[Benenson et al. 2012]

INRIA dataset



Detecting objects without image resizing



1 model,
50 image scales

50 models,
1 image scale


We want to use scene geometry to guide the detections



Monocular


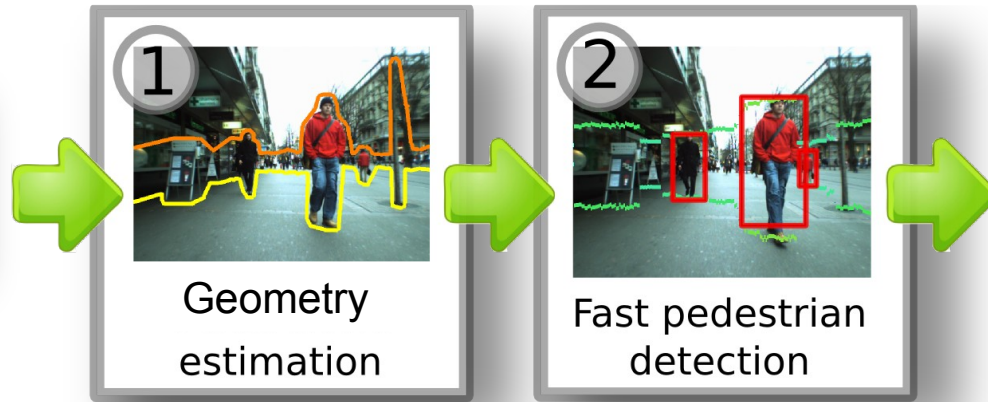


Stereo





Input stereo image




Detected pedestrians

>50 Hz on GPU

Depth maps are slow to compute



< 50 Hz on CPU



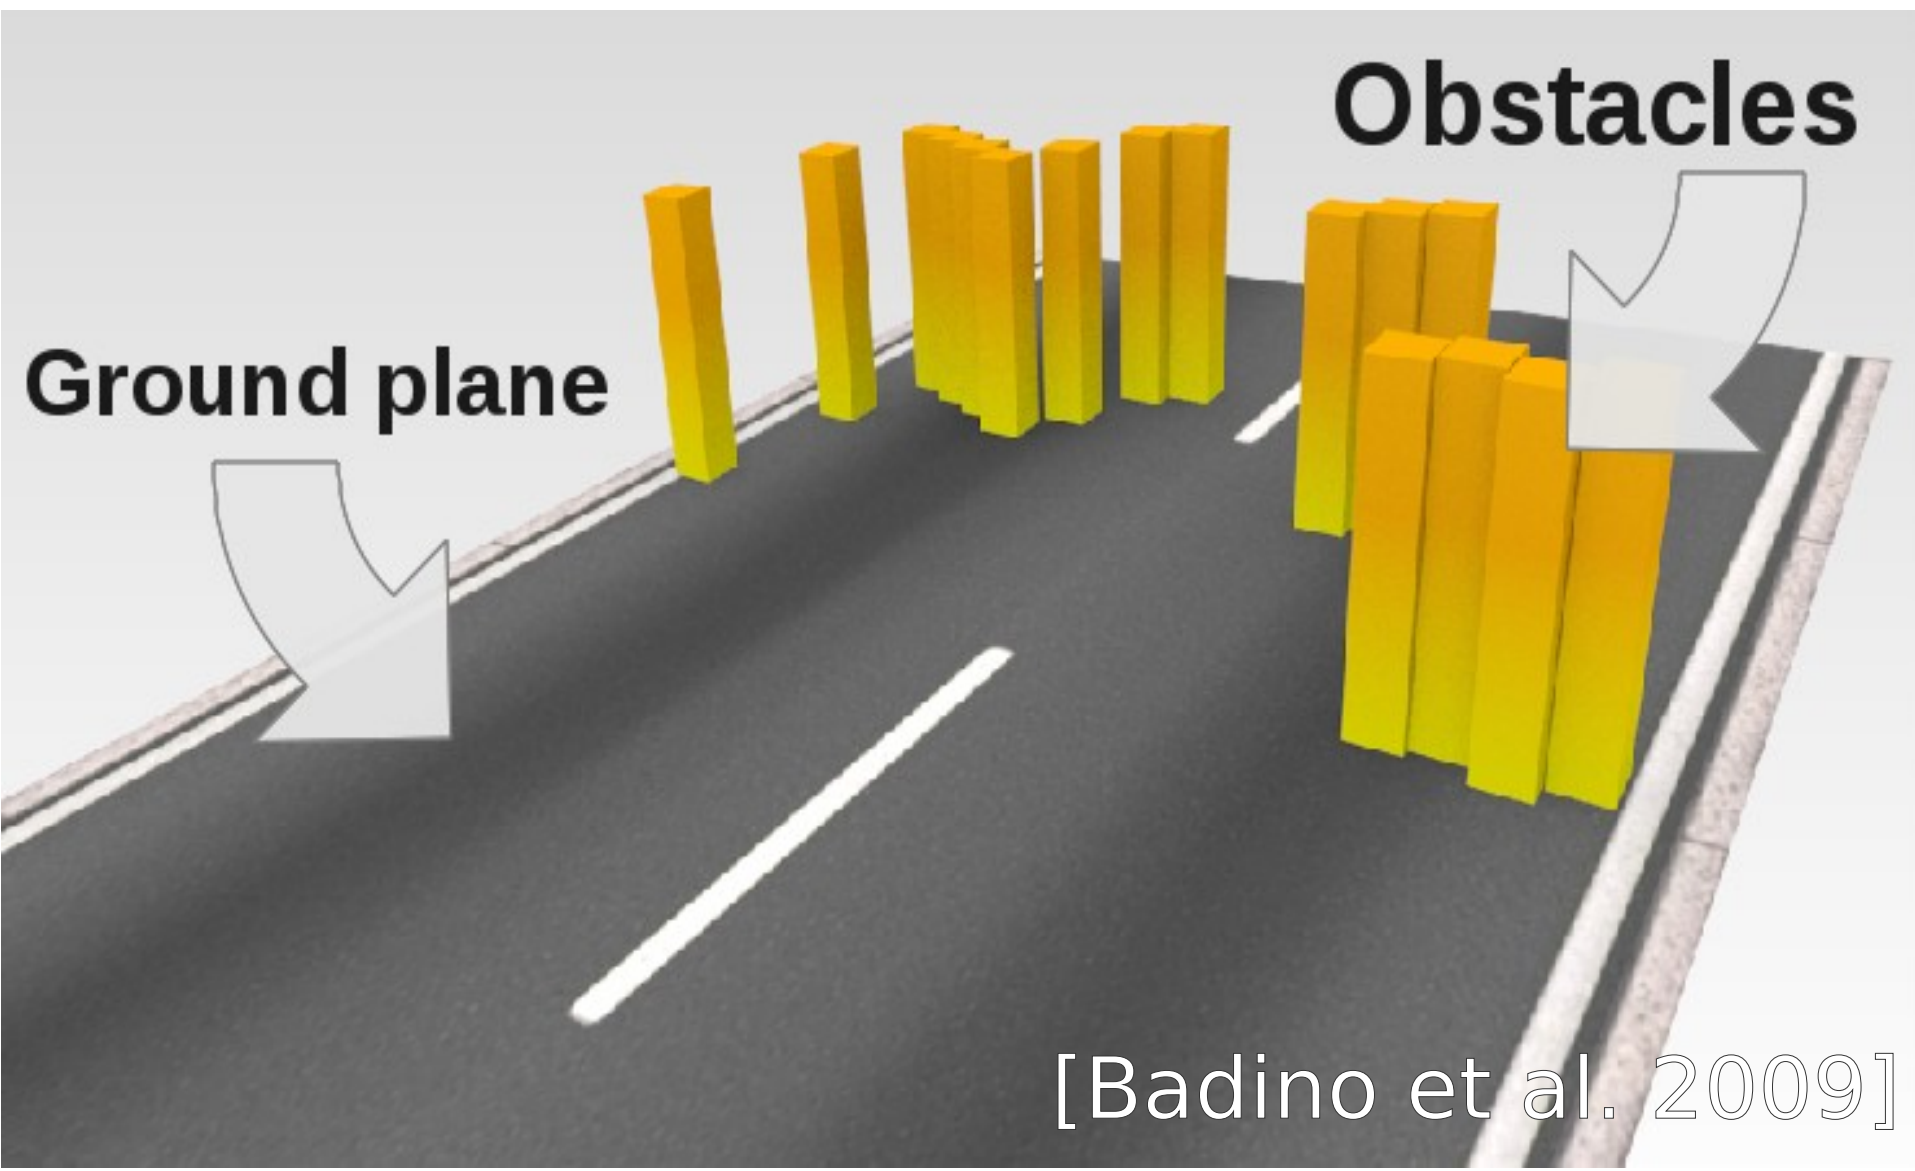
What to do when
depth map computation
is much slower than detection?

Depth maps compute more information that we need



Key insight

Stixel world




[Badino et al. 2009]

Stixel world without depth map



90 Hz on CPU




[Benenson et al. 2011]

Stixel world without depth map



135 Hz

~~90 Hz~~ on CPU




[Benenson et al. 2011 & 2012]

Stixel world without depth map



300 Hz on CPU



CVVT@ICCV 2011

“we can extract scene geometry from stereo at 90 Hz”

CVPR 2012

“we can detect pedestrian at 50 Hz,
135 Hz when using scene geometry”

CVVT@ECCV 2012


“we can extract **scene geometry from stereo at ~300 Hz**,
and thus detect pedestrian at 165 Hz”

Is (direct) stixel world a good idea?


It opens new possibilities



Monocular



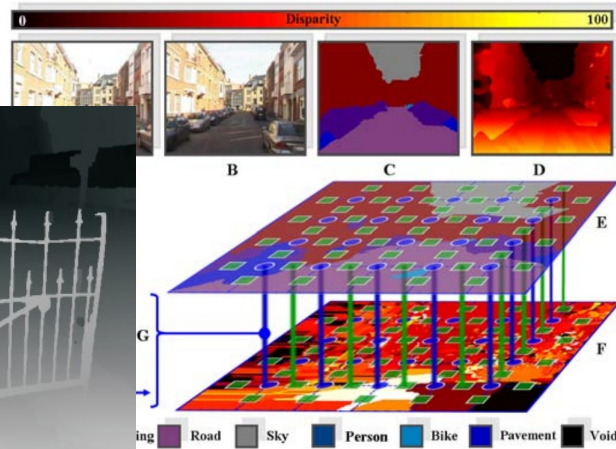




Pixel-wise
depth map



Stixel world

Stereo and object detection is a chicken and egg problem!



[Ladicky et al. 2010], [Bleyer et al. 2011]

Stixel world without depth map



Stixel world without depth map



1) Ground plane estimation

Stixel world without depth map




1) Ground plane estimation

Stixel world without depth map



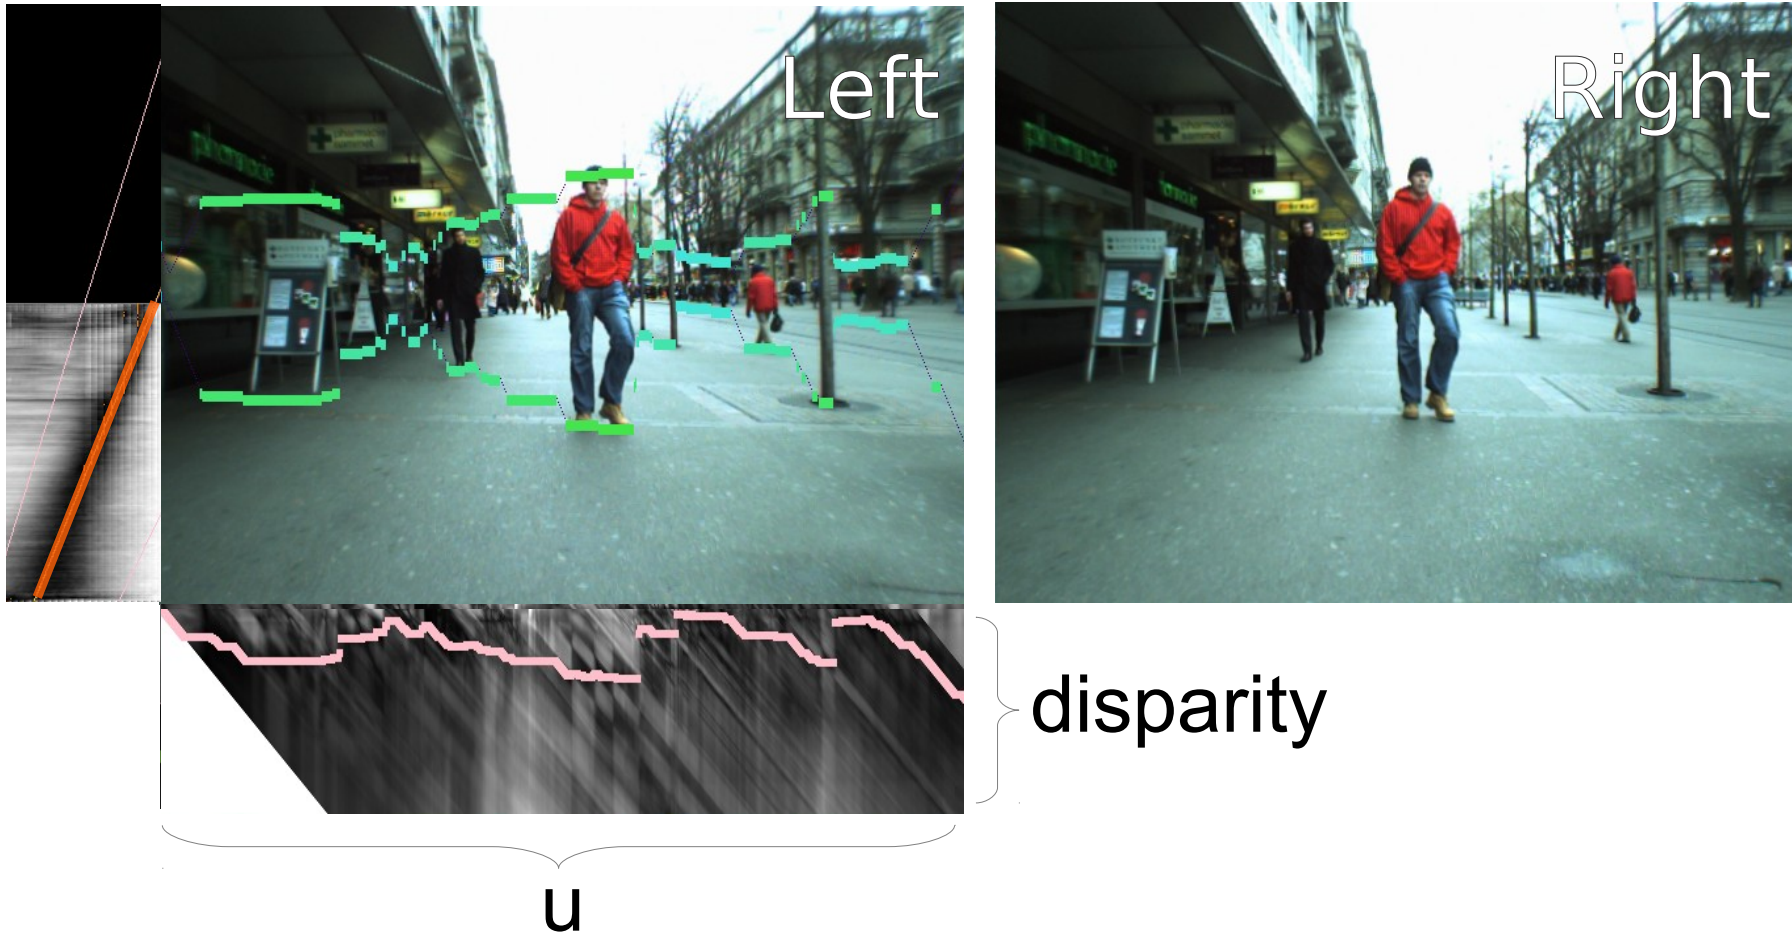
2) Stixel distance estimation

Stixel world without depth map



2) Stixel distance estimation

Stixel world without depth map




2) Stixel distance estimation @ **135 Hz CPU**


100 Hz is too slow!

Our initial proposal has
(at least) two issues

Too focused in u-disparity, instead of the image space




Too focused in u-disparity, instead of the image space



We will **quantize** both rows and columns

Too focused in u-disparity, instead of the u-v space




We will **quantize** both rows and columns

Only half of the image information is used




Only half of the image information is used



Only gradients along rows are used

Only half of the image information is used



Gradients along columns are **ignored**


Now we fully exploit the left image



Now we fully exploit the left image

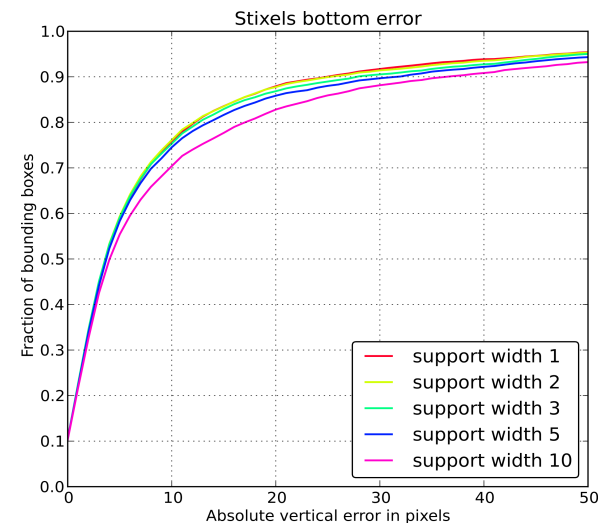
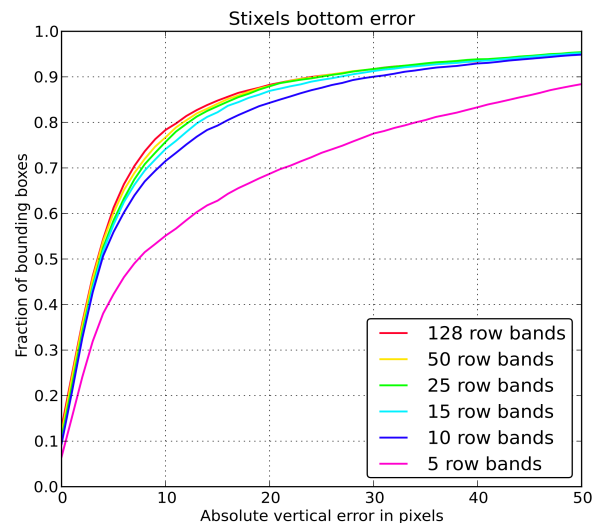
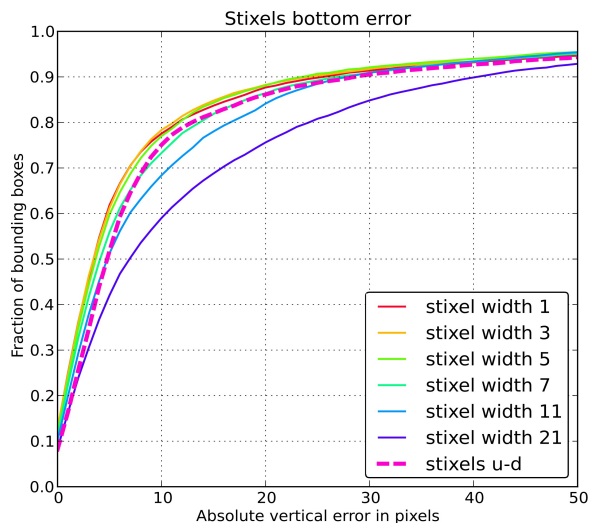


Now we fully exploit the left image



What happens when ?

- Stixels vertical sampling changes?
- Stixels horizontal sampling changes?
- Stixel support width changes?



What happens when ?


- Stixels vertical sampling changes?
→ each **3 columns** is good enough
- Stixels horizontal sampling changes?
→ **25 row bands** is good enough
- Stixel support width changes?
→ **width 1 column** is best

~100 Hz \Rightarrow ~**300** Hz on CPU

(on 640x480 pixels images)




Input stereo image



Stixel world estimation



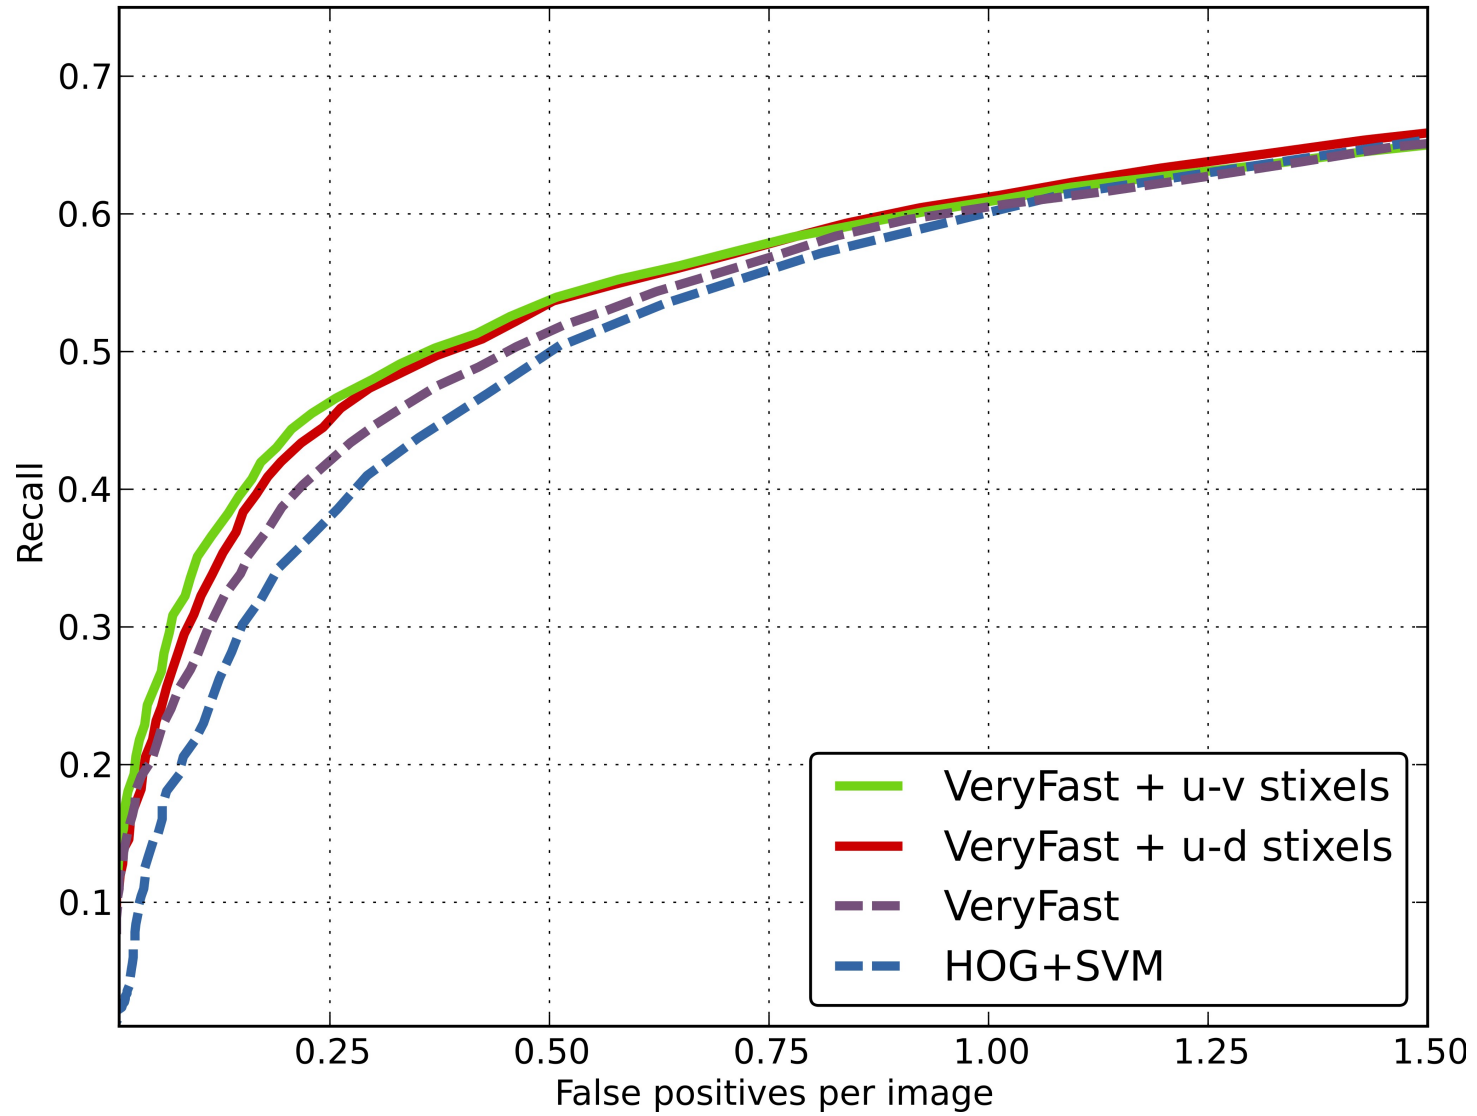
Fast pedestrian detection



Detected pedestrians


More speed, same quality

Recall versus FPPI over Bahnhof dataset,
considering all windows with height > 40 [pixels]




More speed, same quality

Recall versus FPPI over Bahnhof dataset,
considering all windows with height > 40 [pixels]

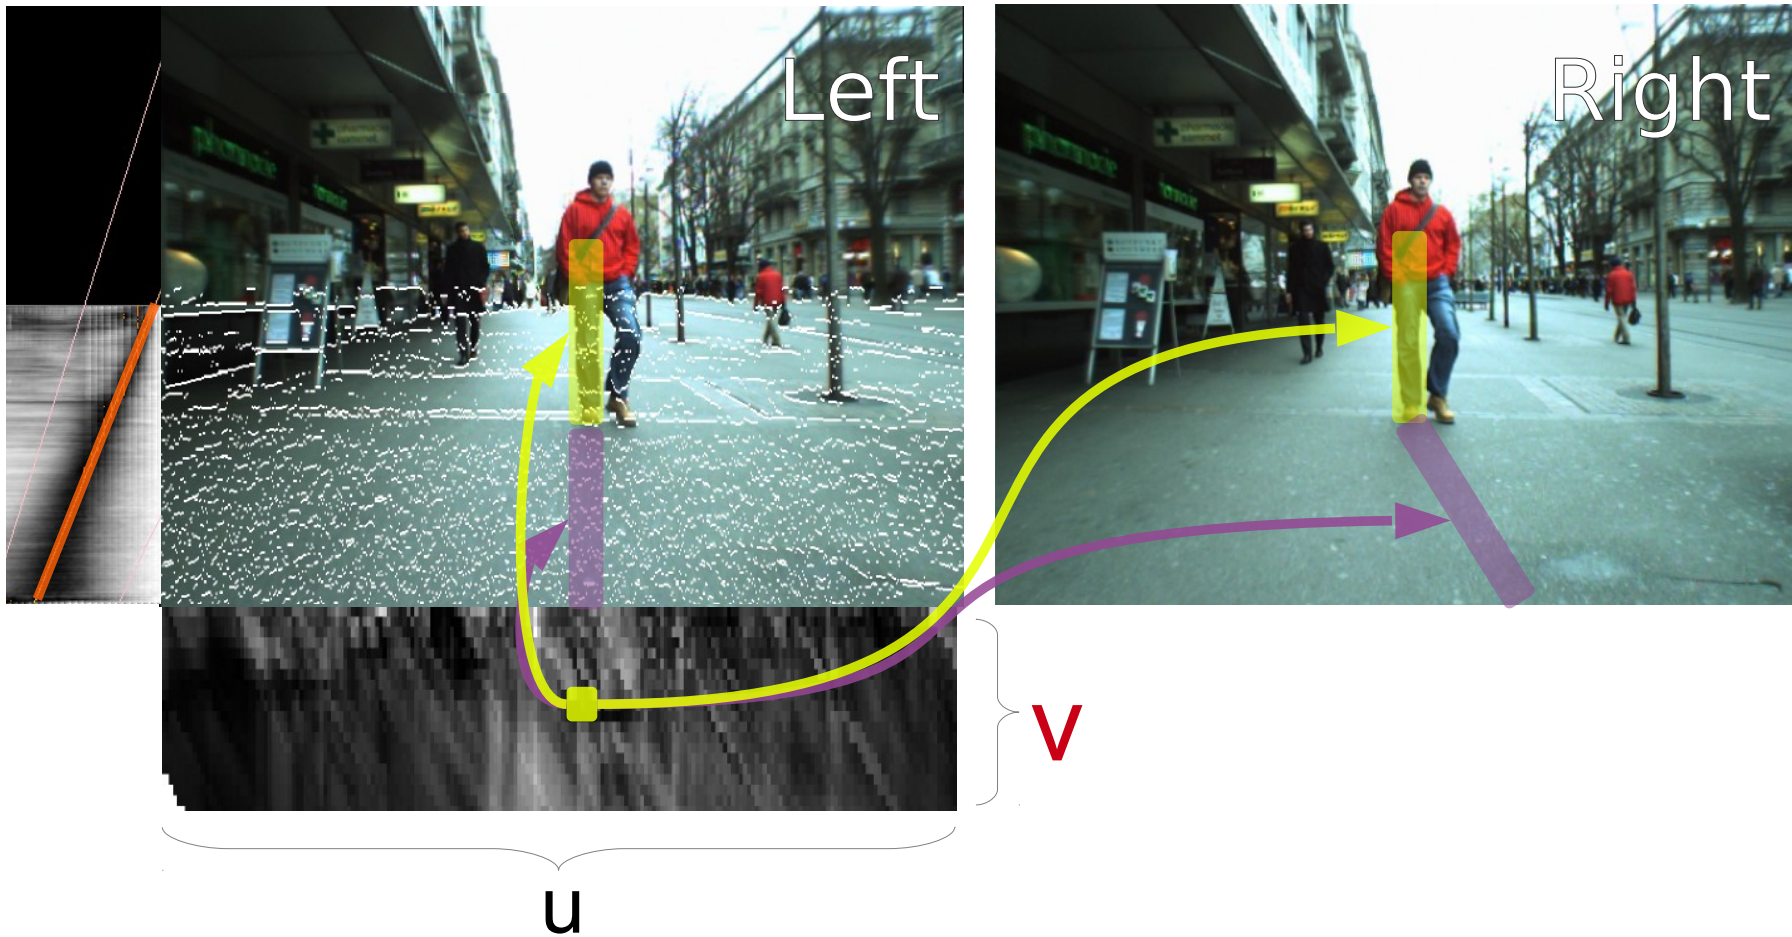


Stixel estimation in u-disparity space is good



Stixel distance estimation @ 135 Hz CPU

Stixel estimation in u-v space is better



2) Stixel distance estimation @ 300 Hz CPU




Take-away message:

Stixels are useful

and



they can be computed very fast
(300 Hz on CPU, less than 4 ms)



Source code available at
<http://rodrigob.github.com>

Rodrigo Benenson

Direct stixels motion estimation



motion

Direct stixels motion estimation

