Microsoft

Microsoft Research Faculty Summit 2013

Microsoft

A Big World of Tiny Motions

William T. Freeman
Professor, Assoc. Dept. Head,
Dept. Elect. Eng. and Computer Science
Massachusetts Institute of Technology

Microsoft Research
Faculty Summit

A big world of tiny motions.

- Joint work with Michael Rubinstein, Hao Yu, Eugine Hsu, Neal Wadhwa, John Guttag, Fredo Durand. SIGGRAPH, 2012, 2013

Michael Rubinstein

summer intern,

Microsoft Research Seattle: 2011,

Microsoft Research New England: 2012

Microsoft Research PhD Fellowship, 2012-2013

Imperceptible Changes in the World

Respiratory motion

Pulse and Blood flow

Buildings swaying
in wind

Magnifying Glass for Temporal Variations

Respiratory motion

Pulse and Blood flow

Magnifying Glass for Temporal Variations

Respiratory motion

Pulse and Blood flow

Amplifying Subtle Color Variations

- 1. Average spatially to overcome sensor and quantization noise

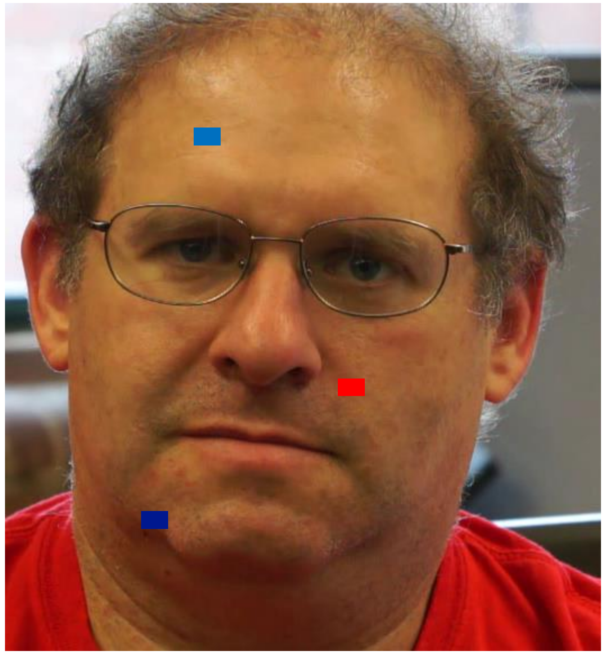

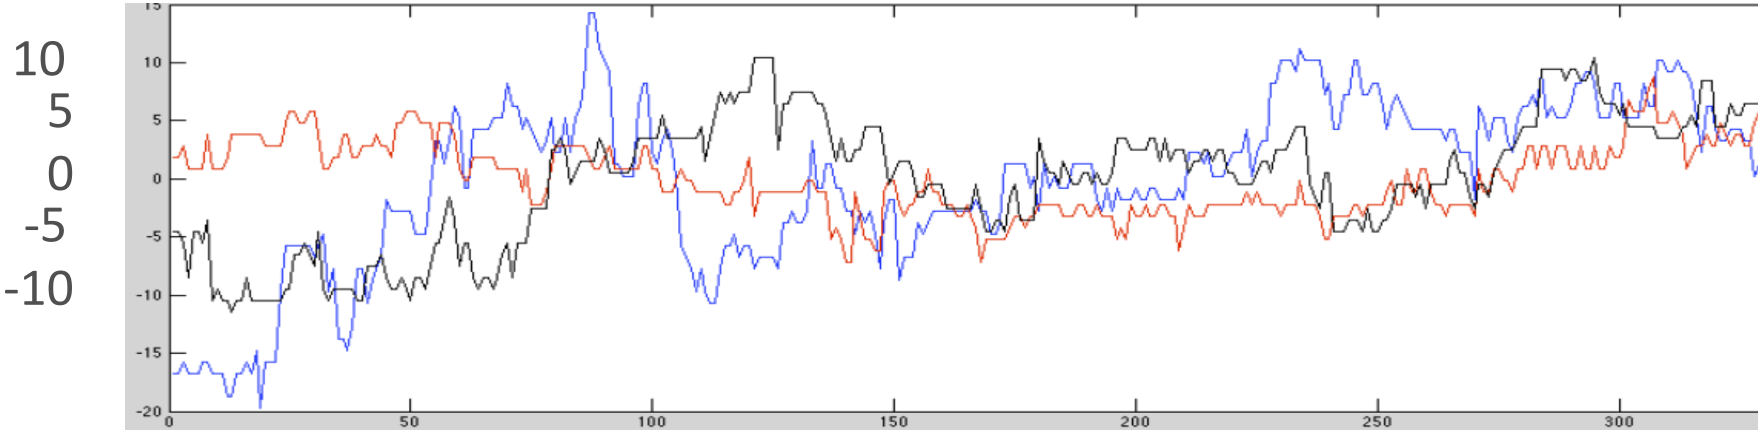

Luminance traces for 3 image locations (zero mean)

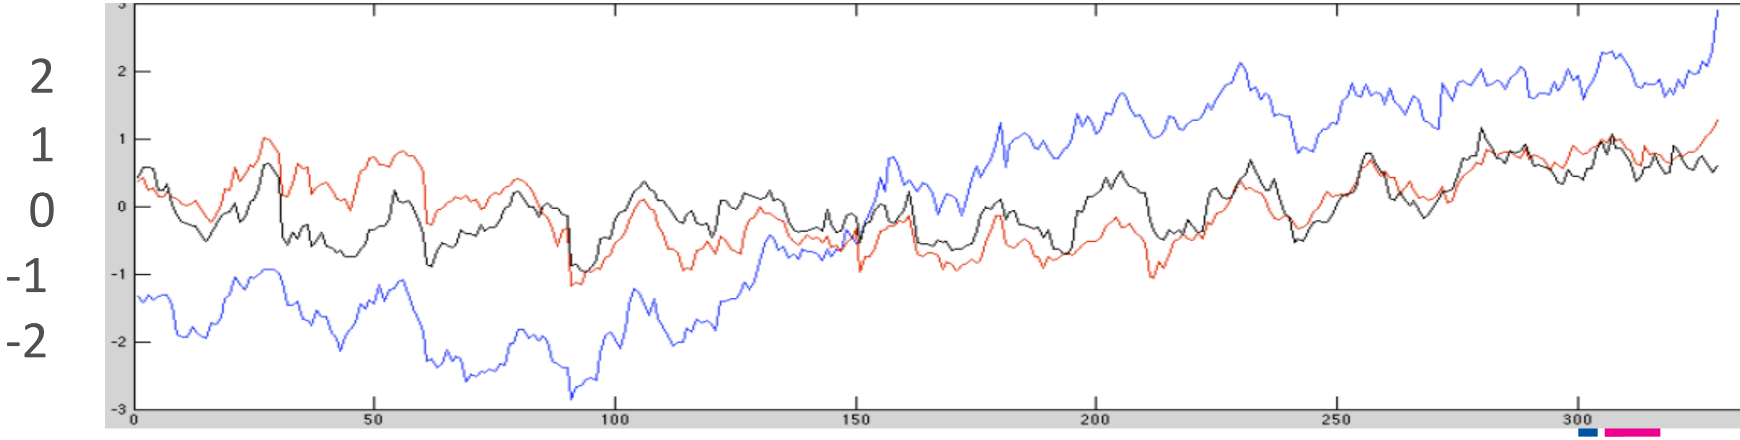

Spatial average in those locations

Amplifying Subtle Color Variations

- 2. Filter Temporally

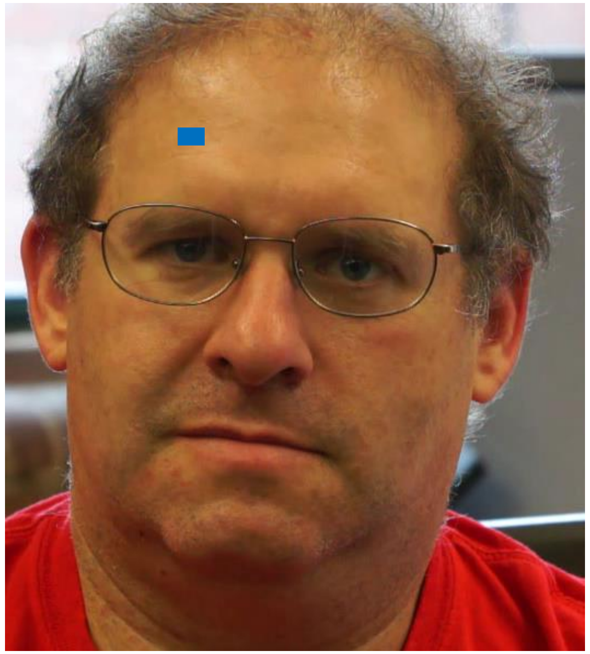

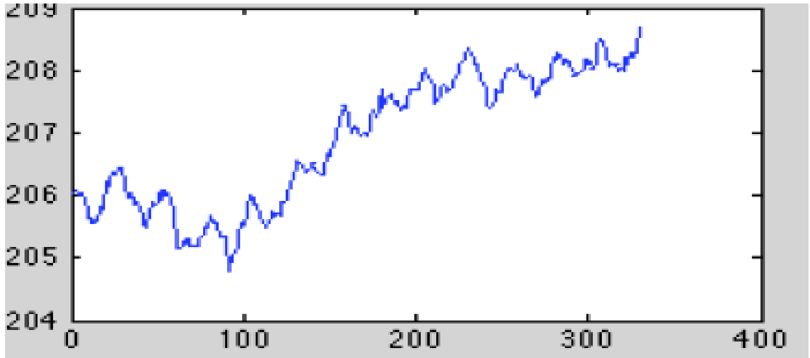

Luminance trace

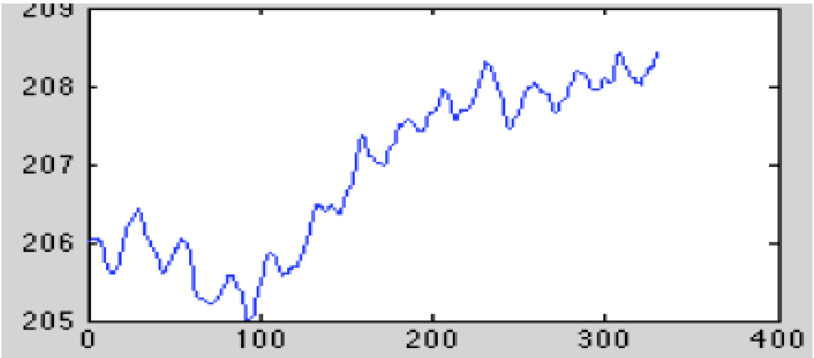

IIR1 [.4,.6]

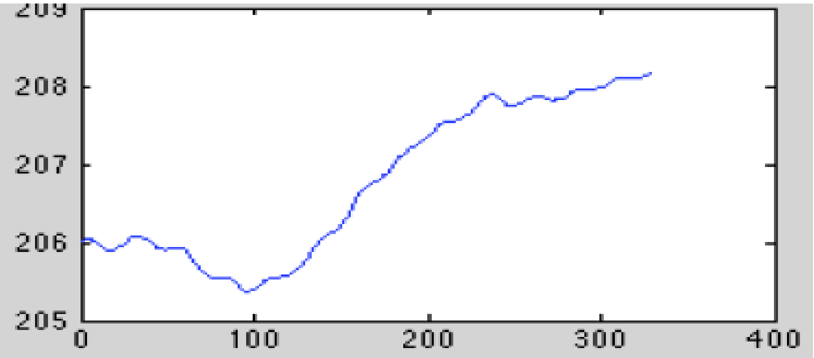

IIR2 [.05,.95]

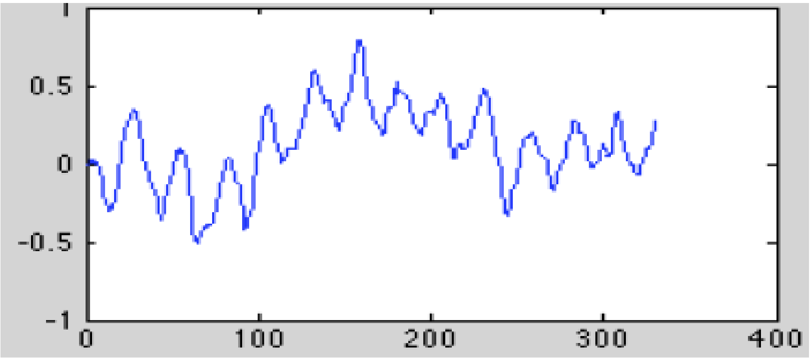

Temporally bandpassed trace

Color Amplification Results

Source

Color-amplified (x100)
0.83-1 Hz (50-60 bpm)
(**ideal filtering**)

Color Amplification Results

Source

Color-amplified (x120)
0.83-1 Hz (50-60 bpm)

Heart Rate Extraction

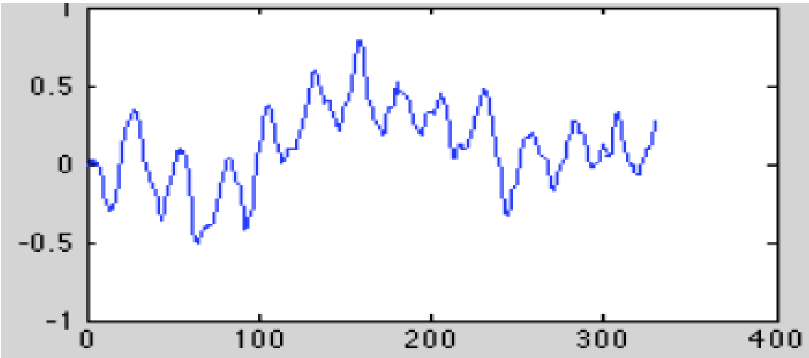

Temporally bandpassed trace
(one patch)

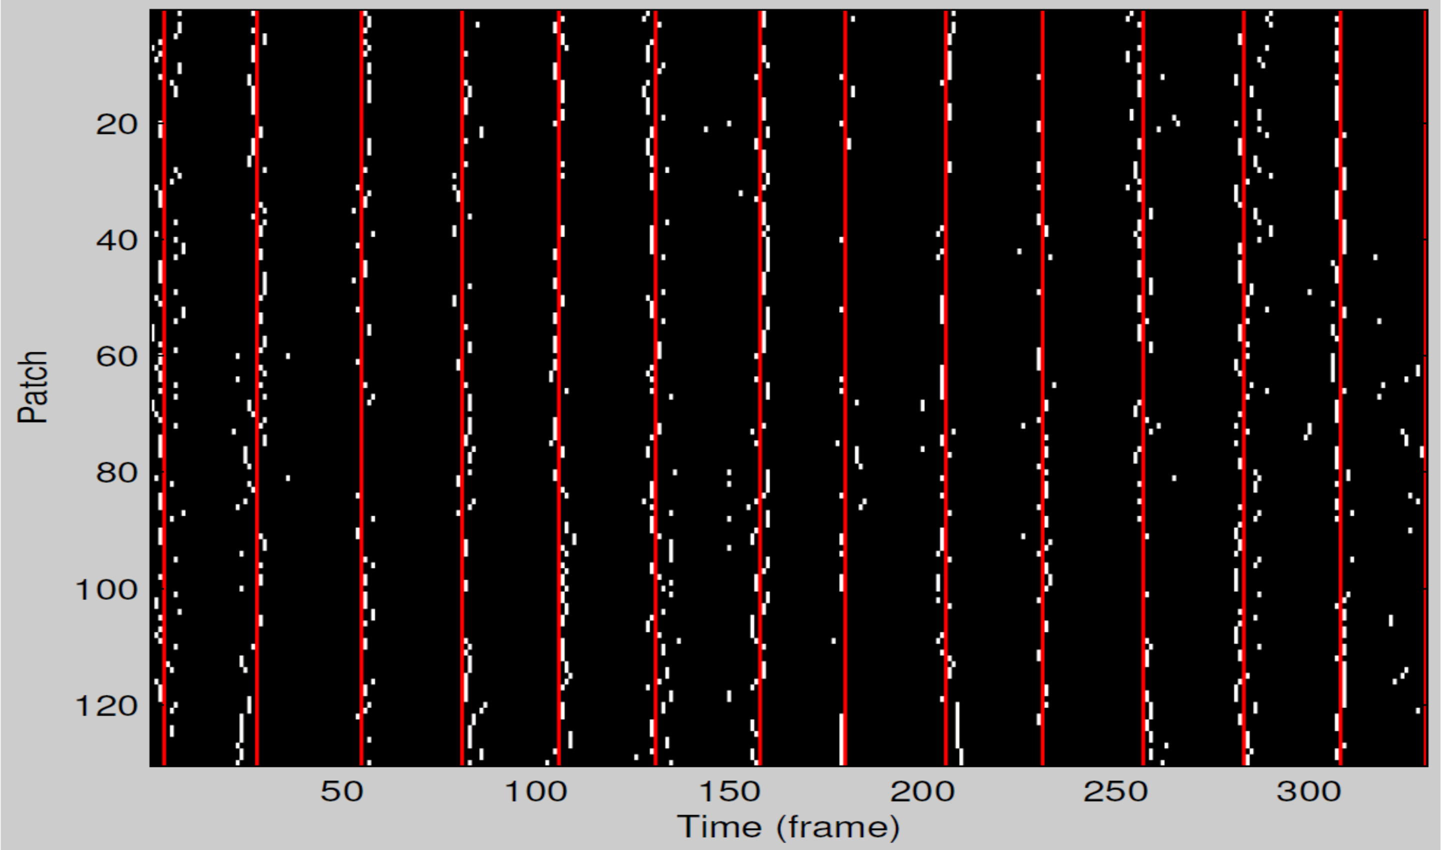

Pulse locations

Heart Rate Extraction

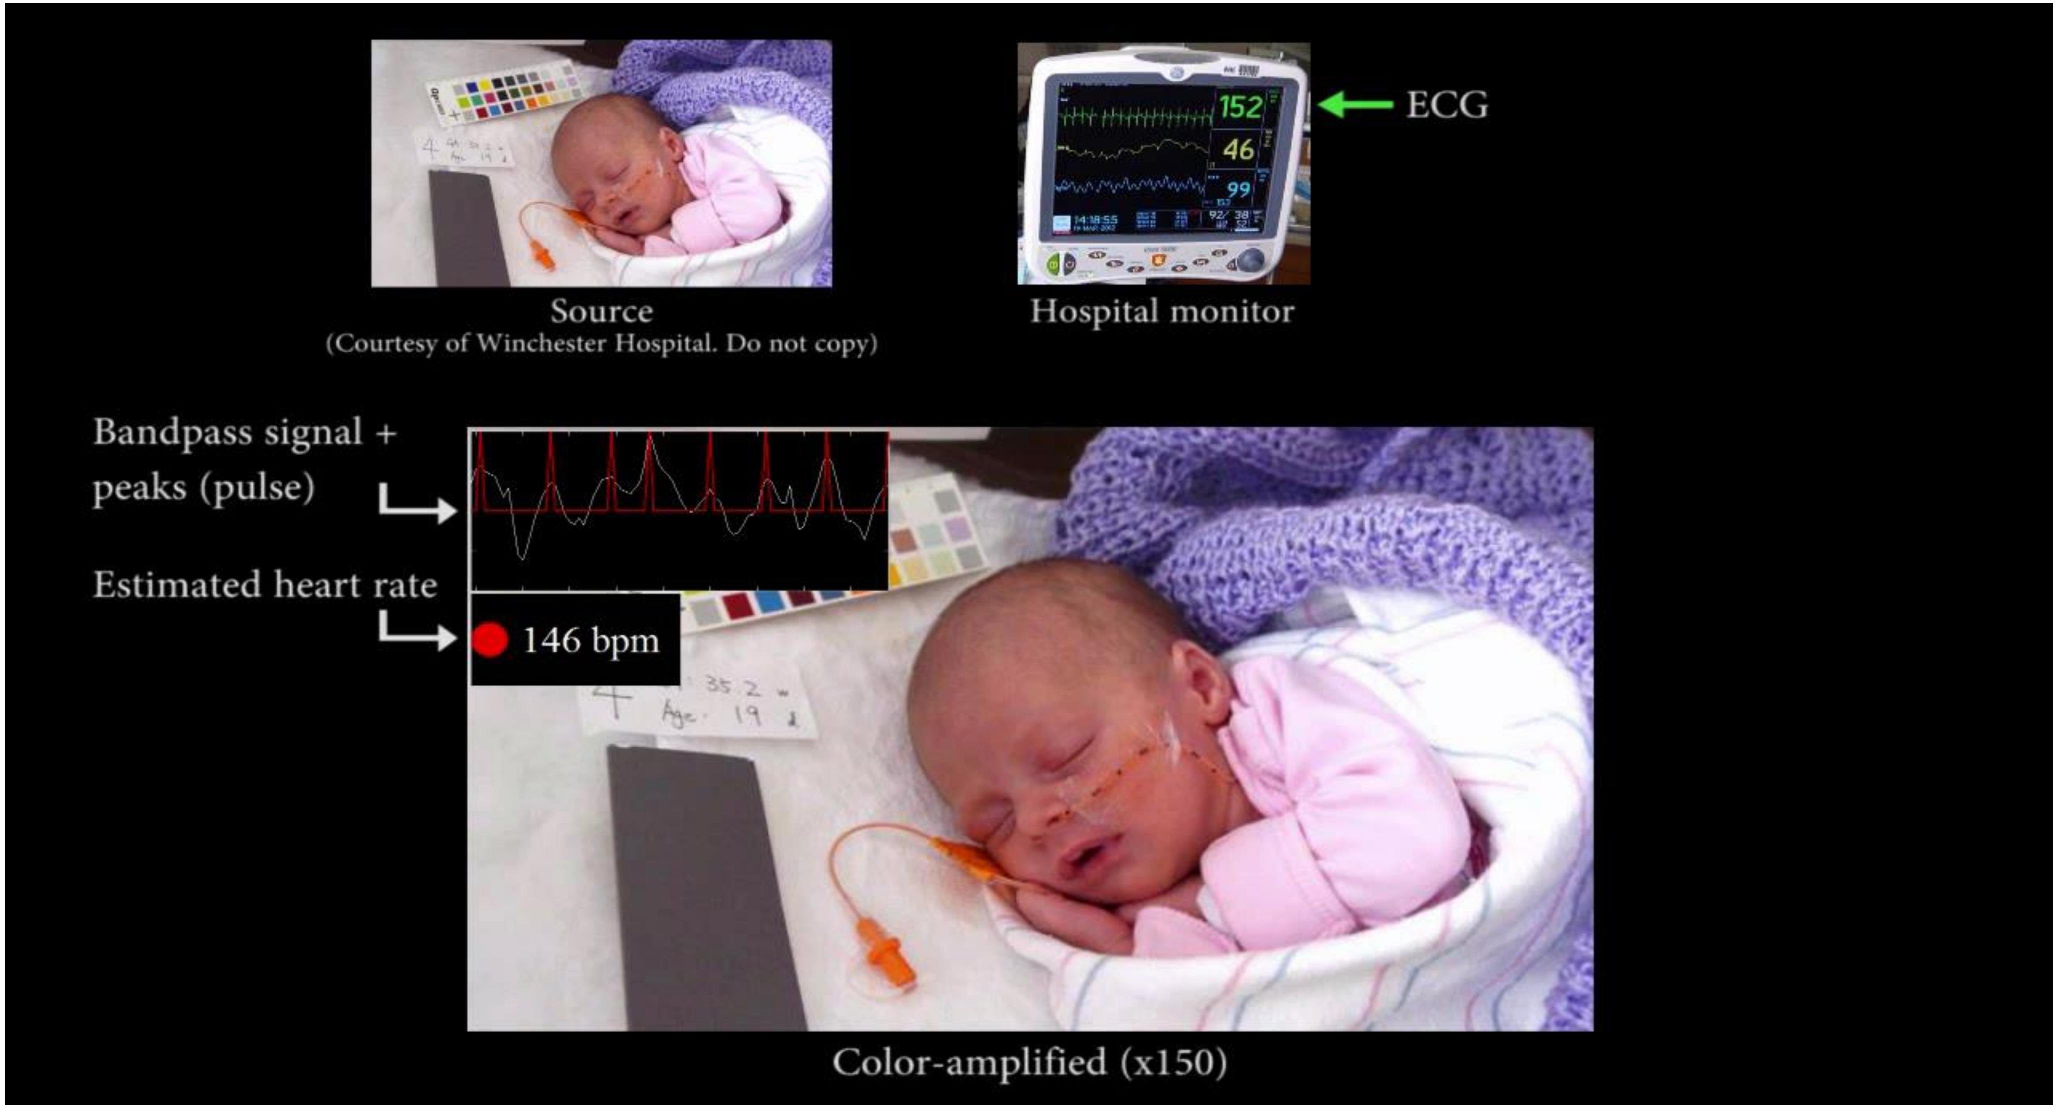

Thanks to Dr. Donna Brezinski and the Winchester Hospital staff

2.33-2.67 Hz (140-160 bpm)

Related Work on Pulse Detection from Videos

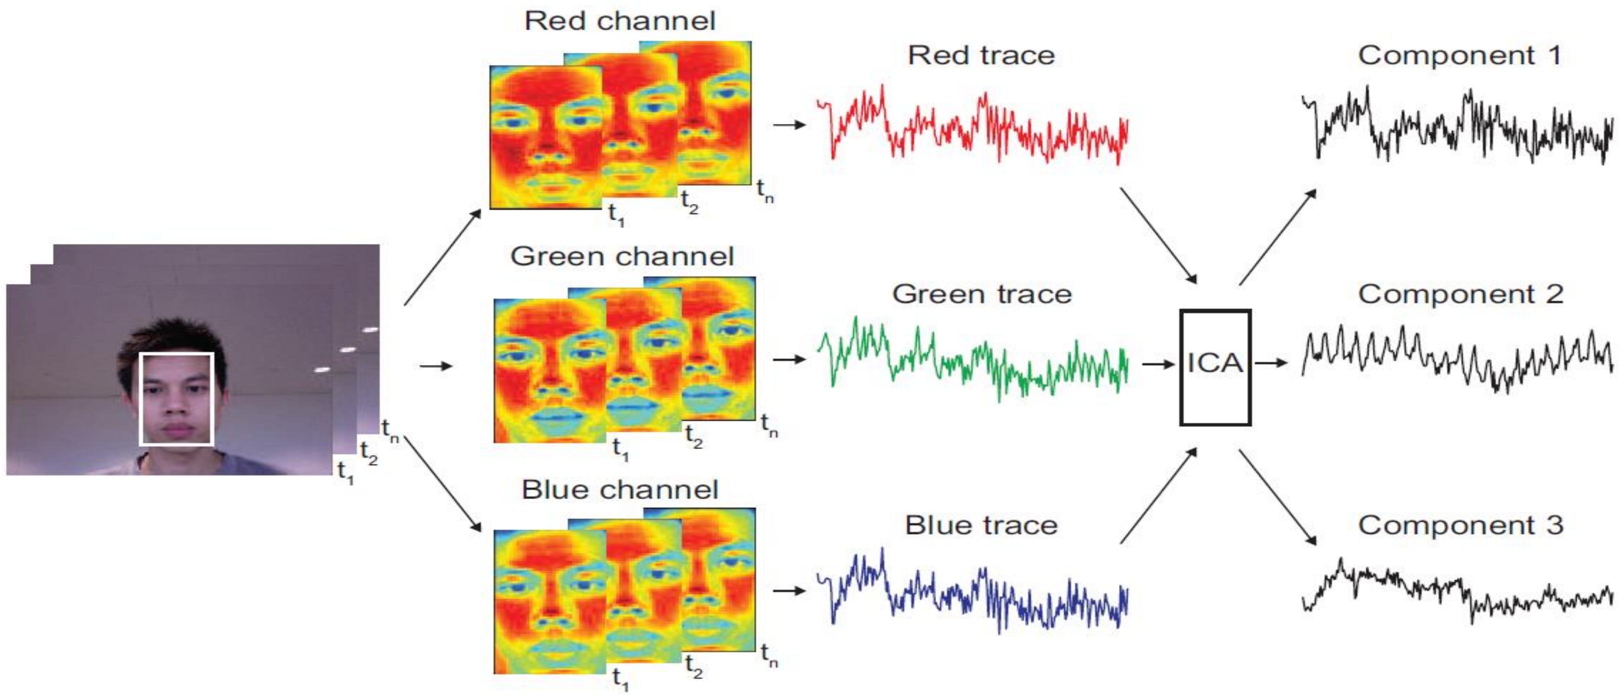

Poh, McDuff and Picard, *Non-contact, automated cardiac pulse measurements using video imaging and blind source separation,* 2010

Face recognition and tracking → RGB averaged over ROI (single number) → Use ICA to find signal sources

Cell Phone Apps…

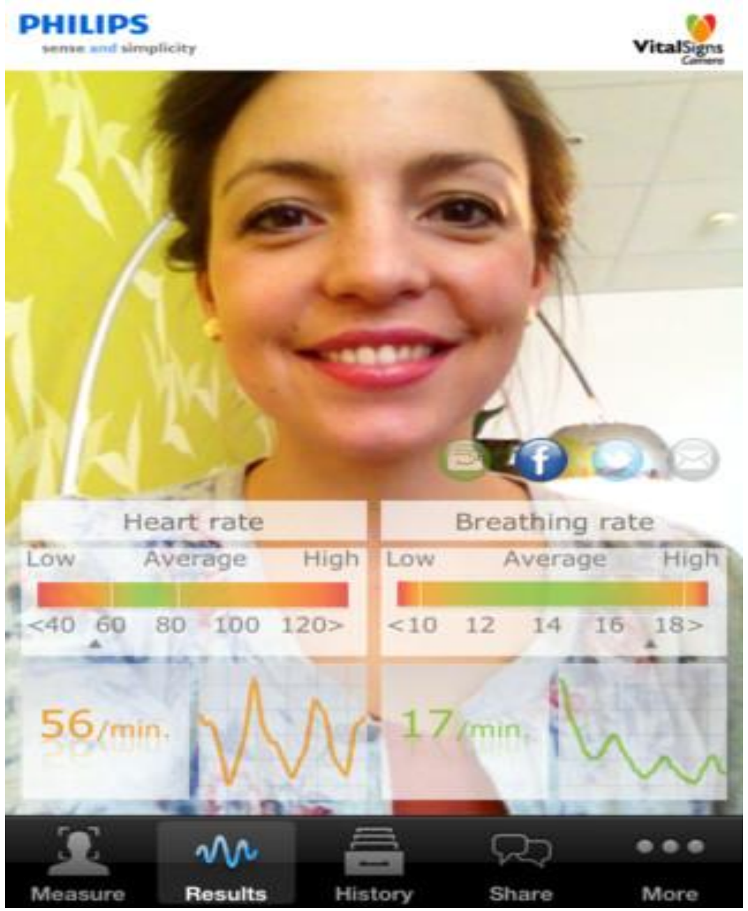

"*Vital Signs Camera*" – Philips (proprietary)

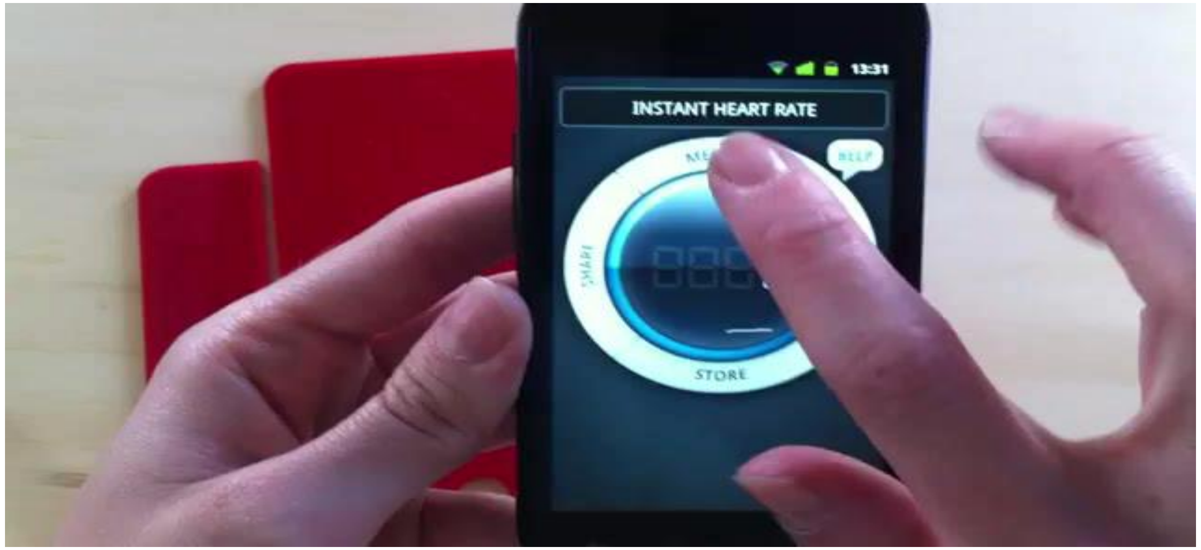

"*Instant Heart Rate*" for Android
Photoplethysmography (PPG)

The extra motion with the color amplification puzzled us

Source

Color-amplified (x100)
0.83-1 Hz (50-60 bpm)
(**ideal filtering**)

Motion Magnification via Temporal Filtering

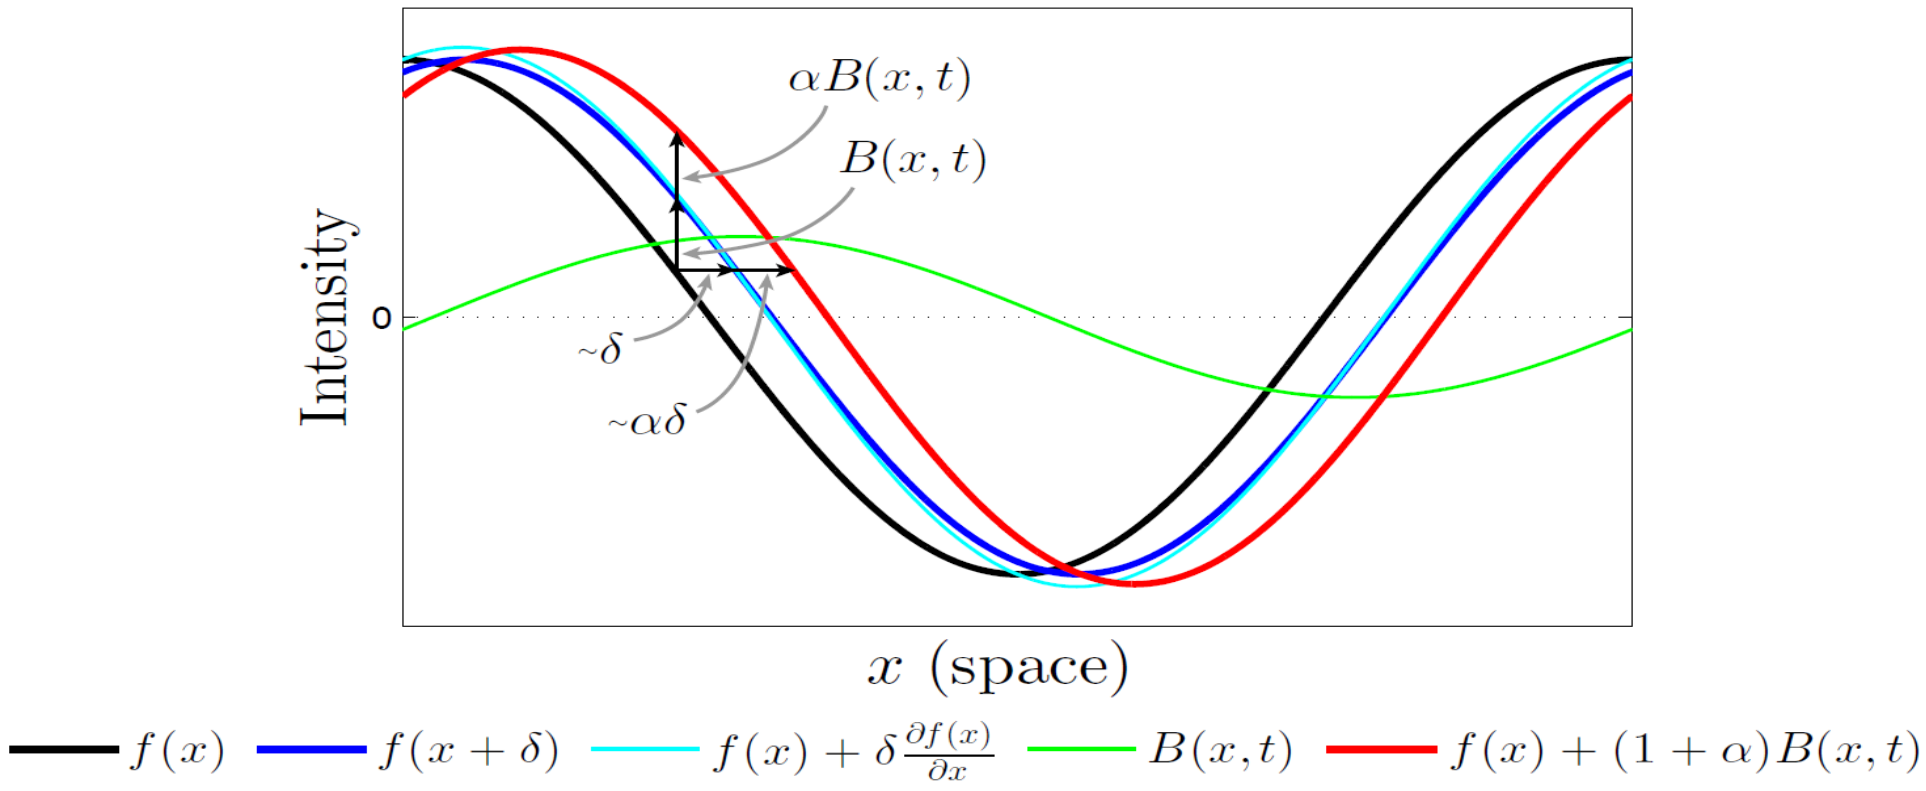

Linearized motion magnification

Rigid translation

$$I(x,t) = f(x+\delta(t))$$

Assume small translation relative to image structures

$$I(x,t) = f(x) + \delta(t)\frac{\partial f(x)}{\partial x}$$

Amplify temporally bandpassed signal

$$\hat{I}(x,t) = I(x,t) + (\alpha - 1)B_t[I(x,t)]$$

Modified signal

$$= f(x) + \alpha\delta(t)\frac{\partial f(x)}{\partial x}$$

Assume the amplified translation is still small relative to image structures

$$= f(x+\alpha\delta(t))$$

Where we expect this to break down

Let's look at it for f(x) being a sinusoid,

Linearized

$$\cos(\omega x + \alpha\omega\delta) = \cos(\omega x) - \alpha\omega\delta \sin(\omega x)$$

$$= \cos(\omega x)\cos(\alpha\omega\delta) - \sin(\omega x)\sin(\alpha\omega\delta)$$

Exact translation by $\alpha\omega\delta$

For the motion magnification approximation to hold:

$$\cos(\alpha\omega\delta) = 1$$

$$\lambda = \frac{2\pi}{\omega}$$

$$\sin(\alpha\omega\delta) = \alpha\omega\delta$$

Condition required for those conditions to be approximately true:

$$\sin(\frac{\pi}{4}) = 0.9\frac{\pi}{4}$$

$$\alpha\omega\delta < \frac{\pi}{4}$$

true displacement

spatial frequency

magnification amount

Synthetic 1D Examples

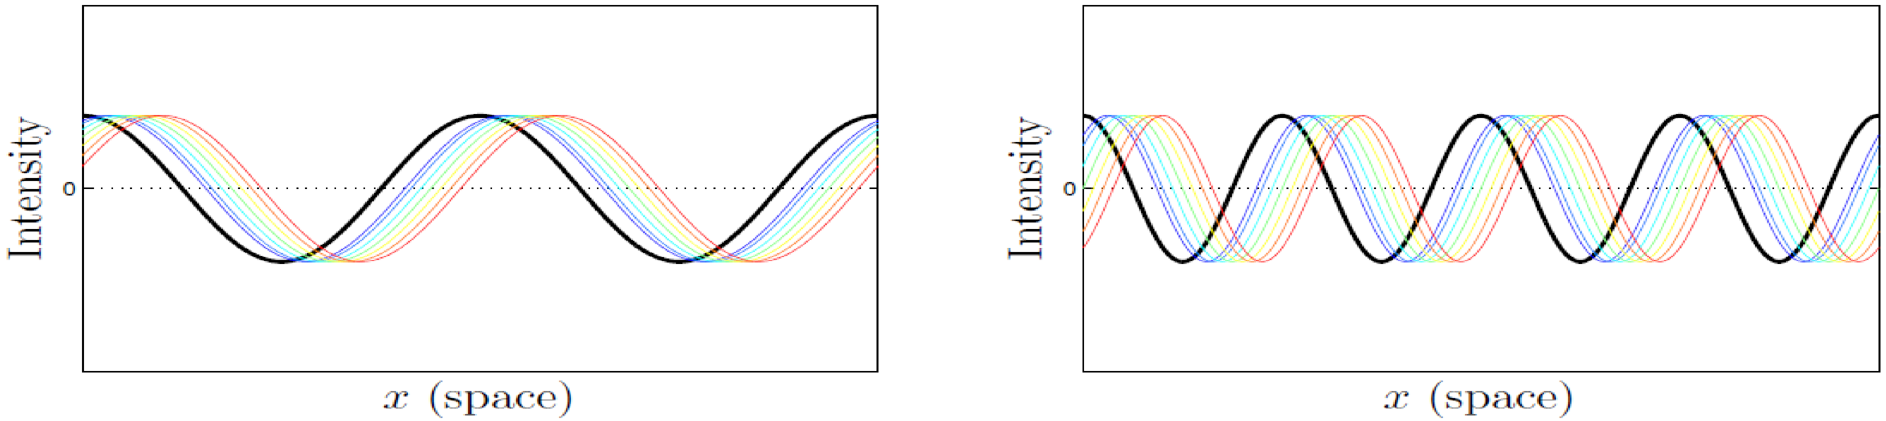

(a) True motion amplification: $\hat{I}(x,t) = f(x + (1+\alpha)\delta(t))$.

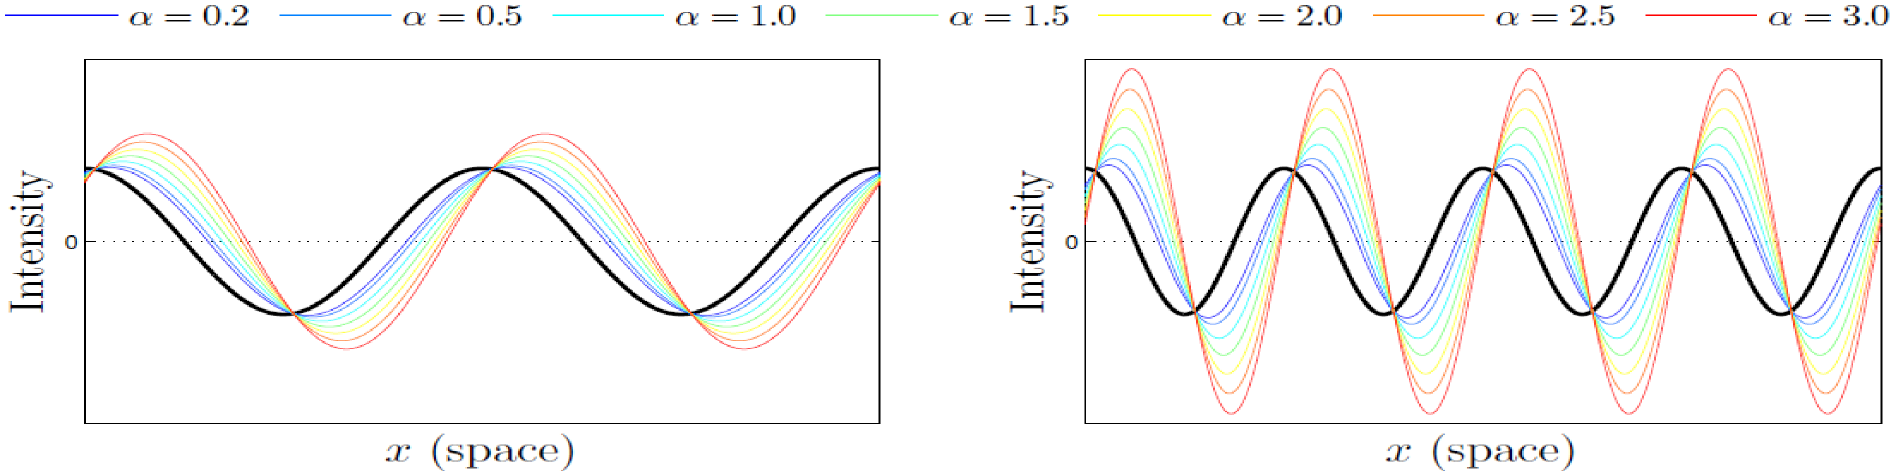

(b) Motion amplification via temporal filtering:
$\tilde{I}(x,t) = I(x,t) + \alpha B(x,t)$.

System Overview

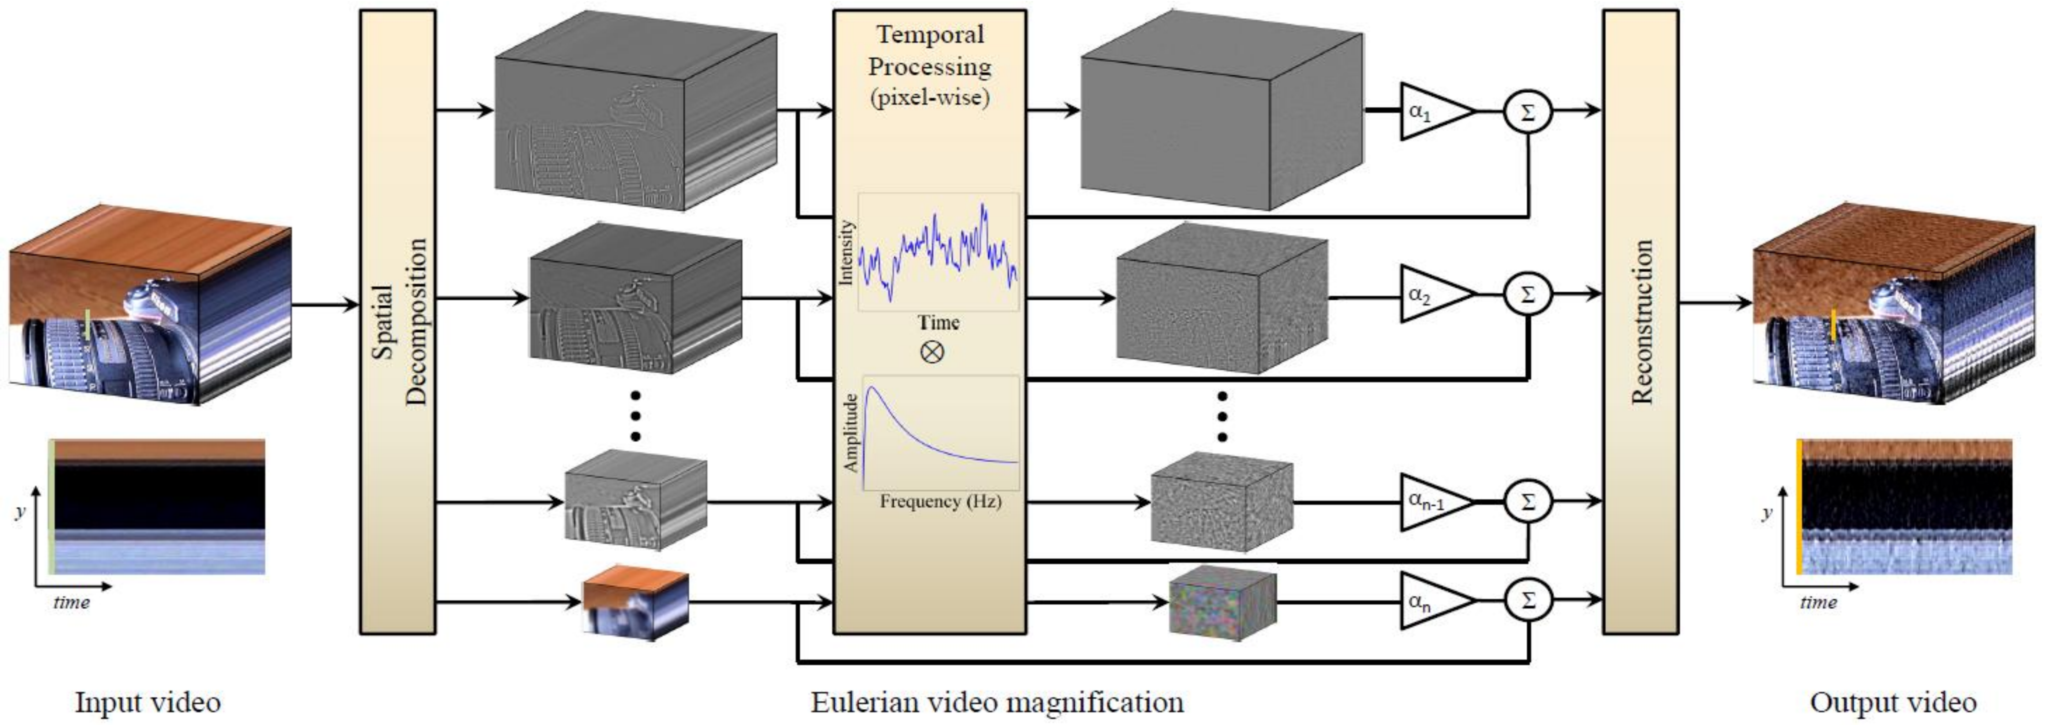

Amplify spatial frequencies where approximation holds, otherwise fail toward zero

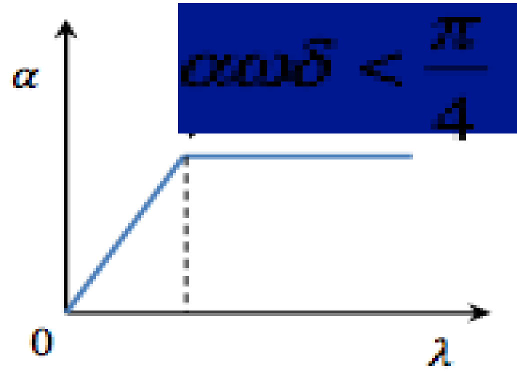

$$\alpha\omega\delta < \frac{\pi}{4}$$

$$\lambda = \frac{2\pi}{\omega}$$

Figure 6: *Amplification factor, α, as function of spatial wavelength λ, for amplifying motion. The amplification factor is fixed to α for spatial bands that are within our derived bound (Eq. 14), and is attenuated linearly for higher spatial frequencies.*

Source

Motion-magnified (3.6-6.2 Hz, x60)

Demo

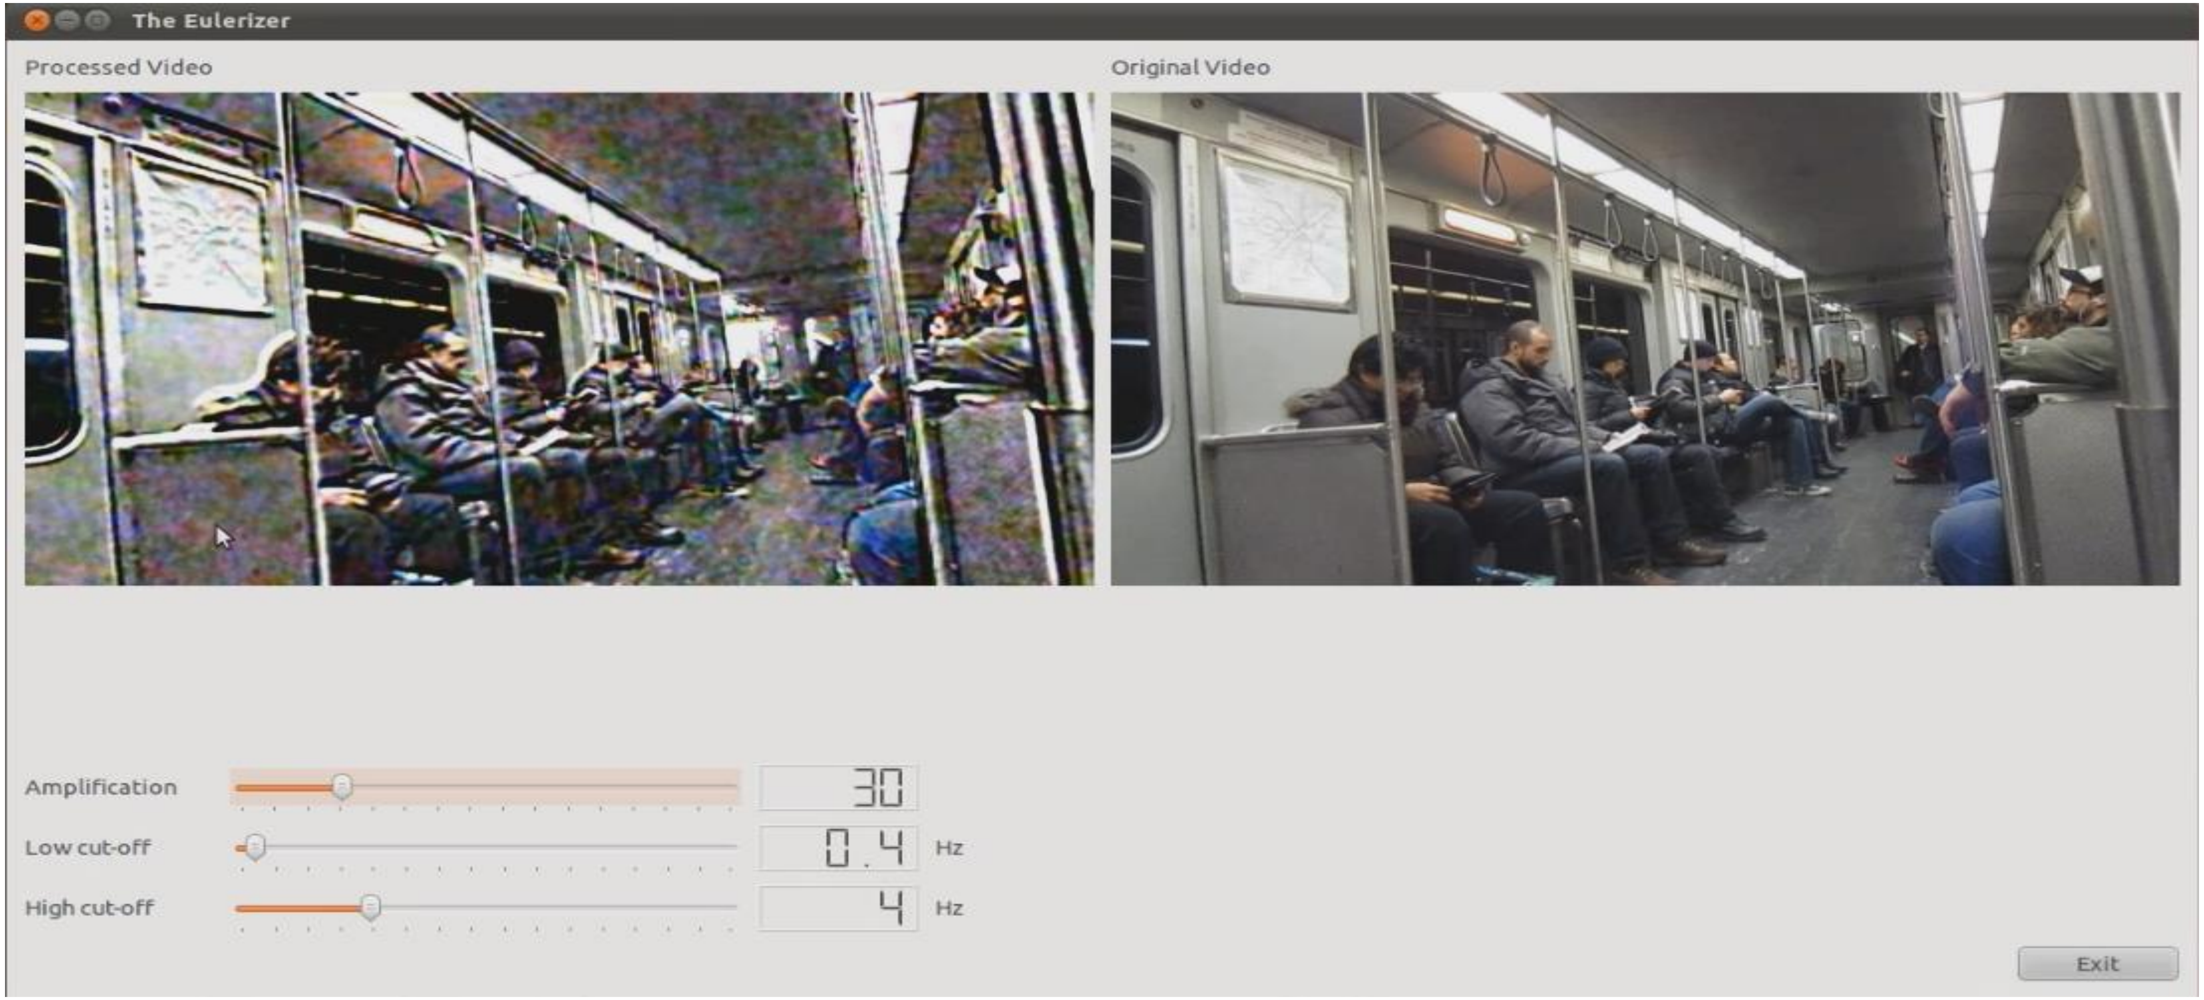
The Eulerizer
Processed Video
Original Video
Amplification
30
Low cut-off
0.4
Hz
High cut-off
4
Hz
Exit

Motion Magnification Results

Source

Motion-magnified (0.4-3 Hz, x10)

Motion Magnification

Source

Motion-magnified (0.4-3 Hz, x10)

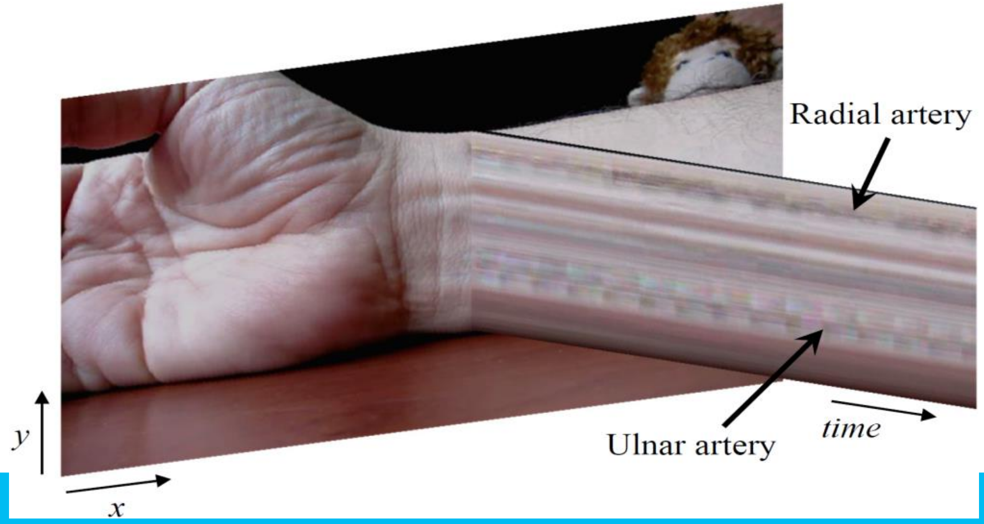

Synthetic 2D Example

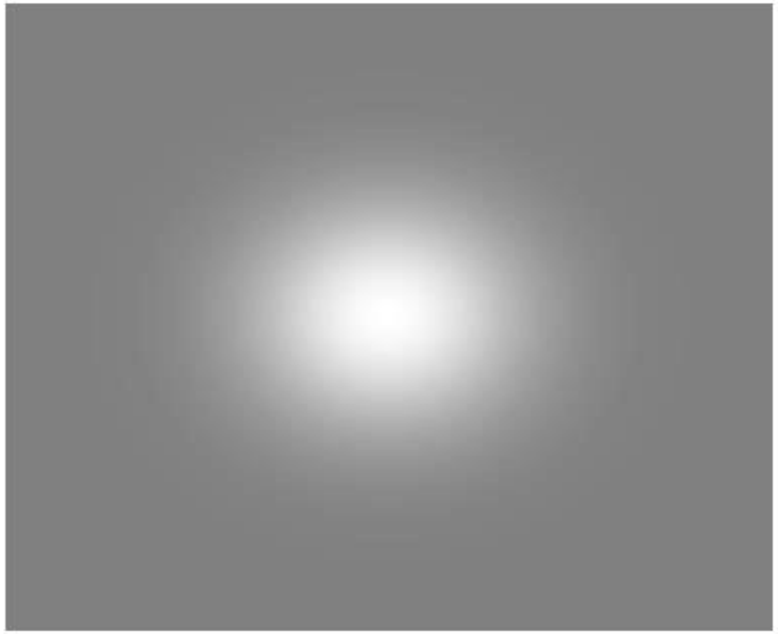

Source

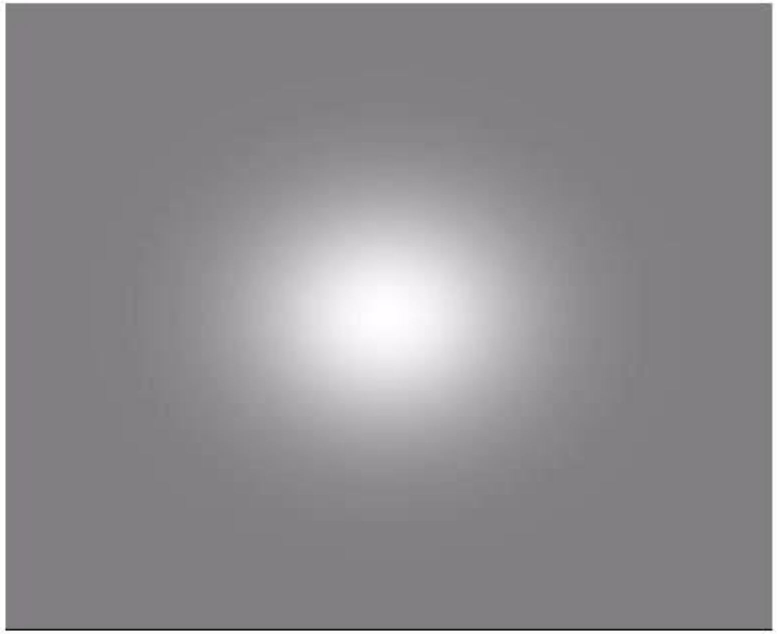

Motion-magnified

Temporal Filters

- Mostly application dependent
  - Configurable by the user

- Some of the filters we used (and their applications):

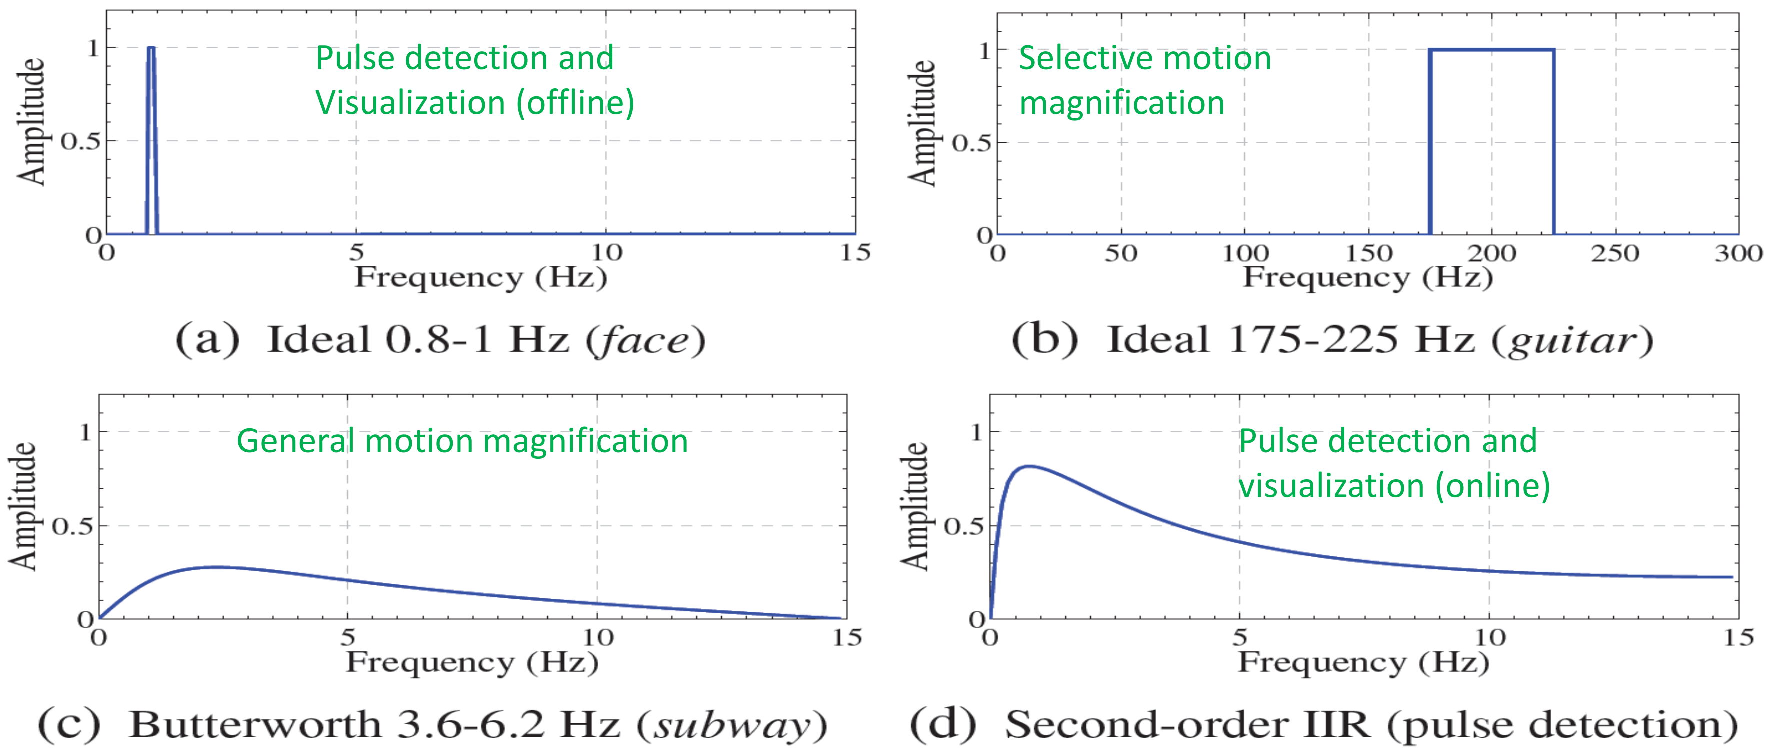

Selective Motion Magnification

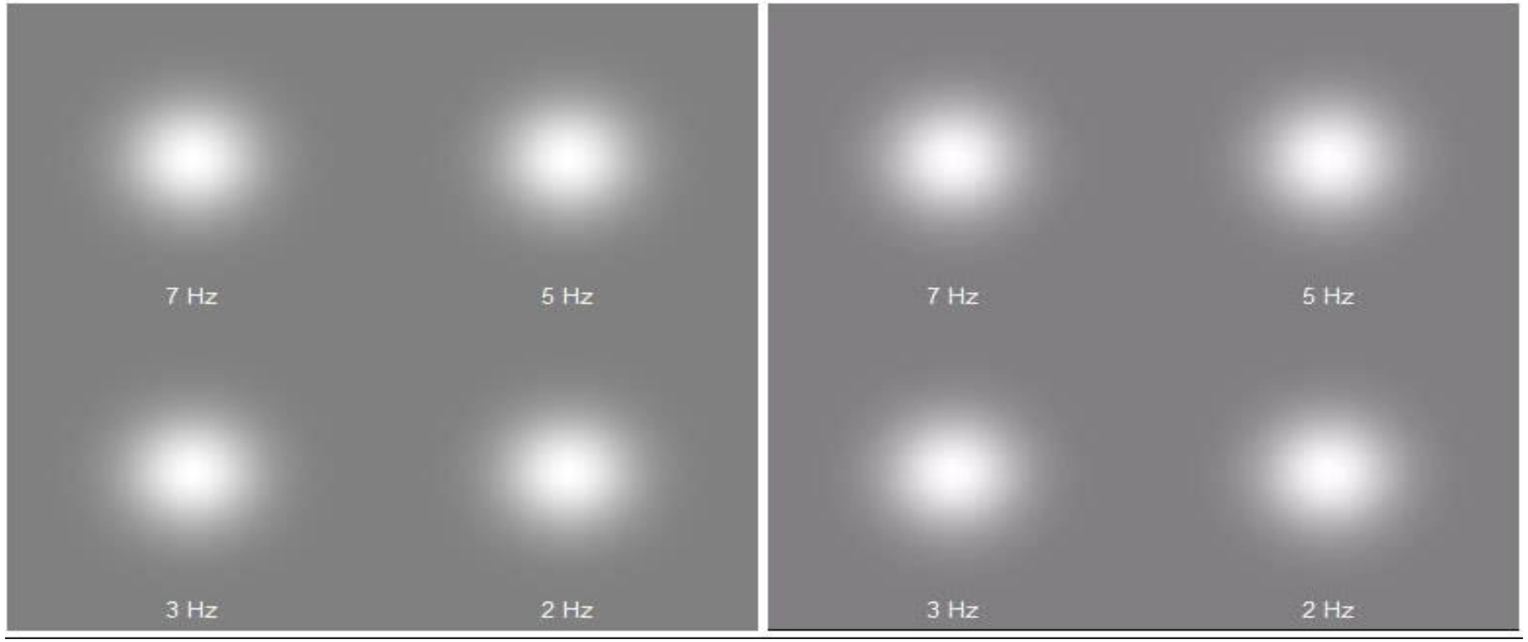

Source

Motion-magnified (2 Hz)

Temporal filter:

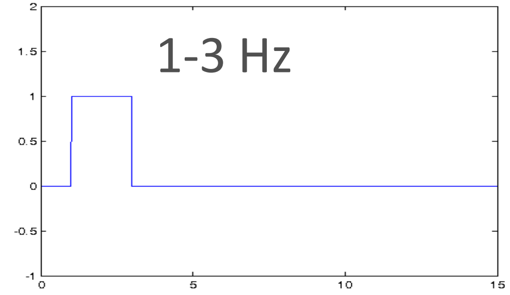

Selective Motion Magnification

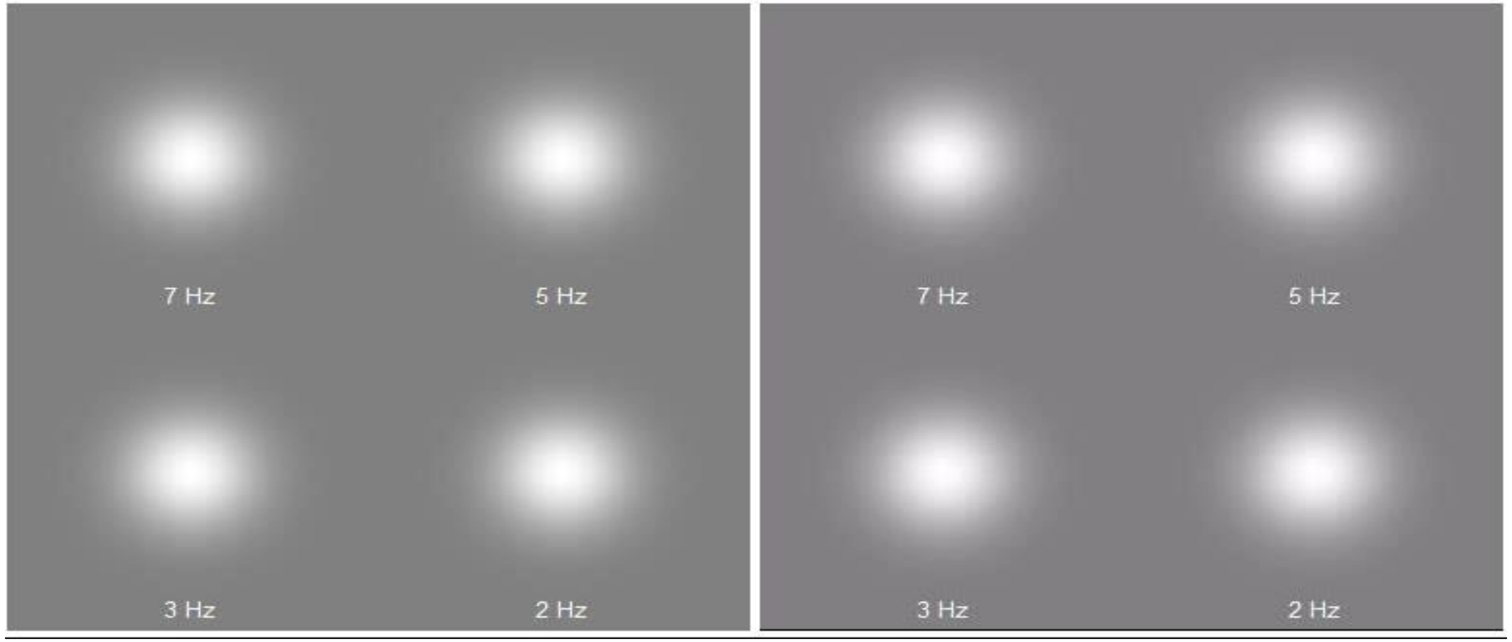

Source

Motion-magnified (3 Hz)

Temporal filter:

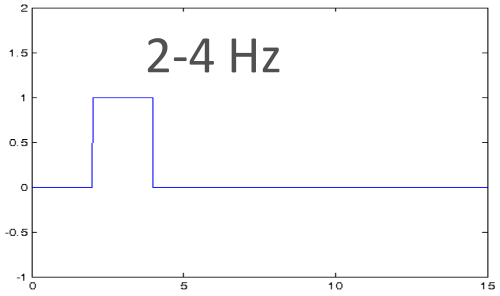

Selective Motion Magnification

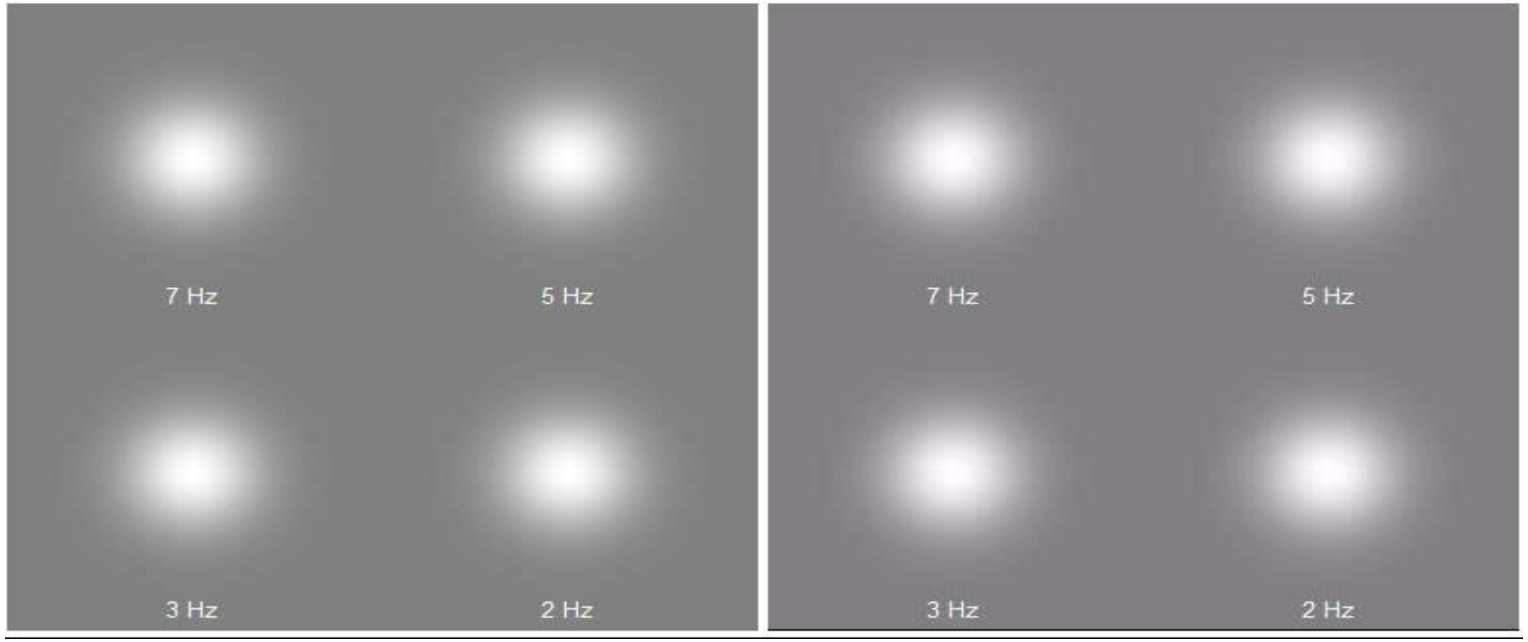

Source

Motion-magnified (5 Hz)

Temporal filter:

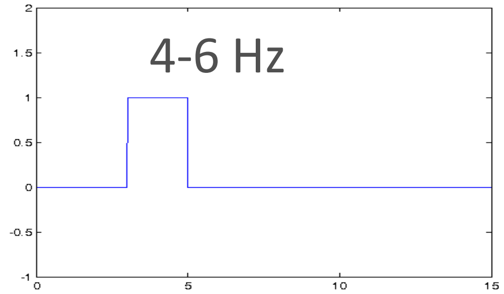

Selective Motion Magnification

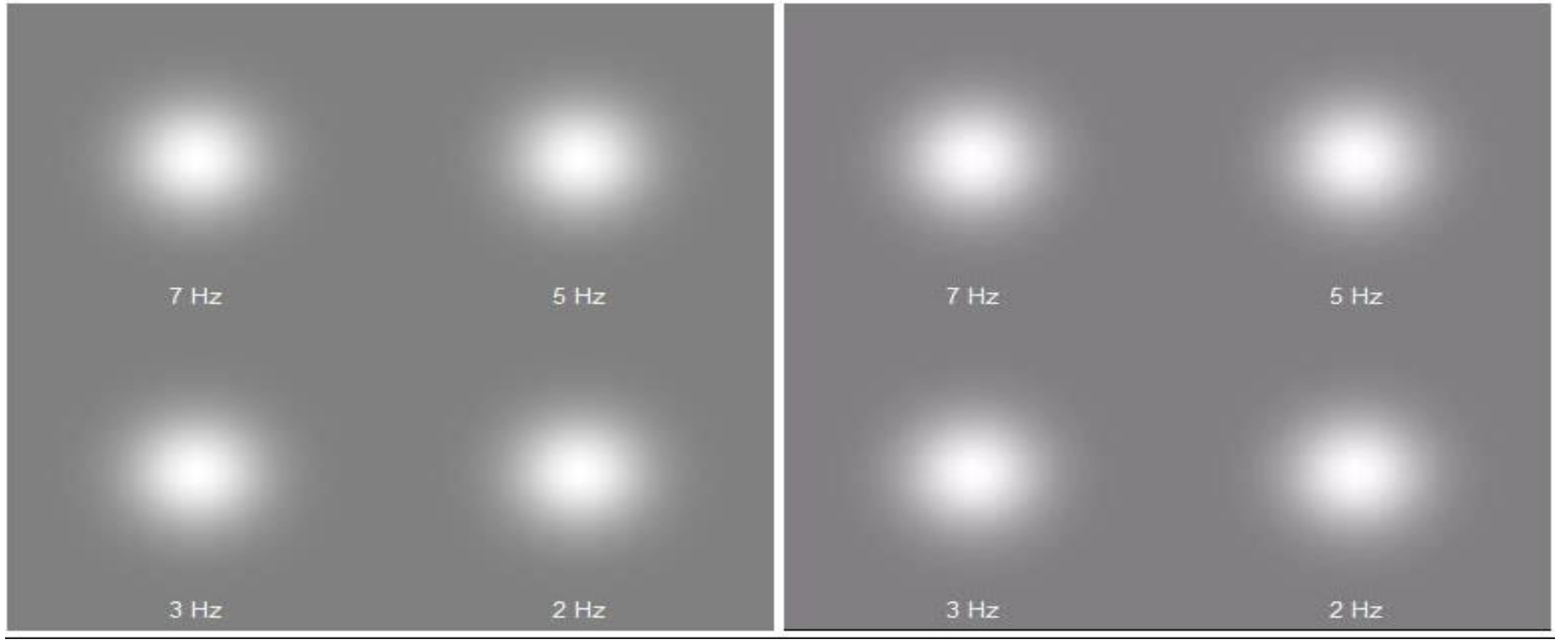

Source

Motion-magnified (7 Hz)

Temporal filter:

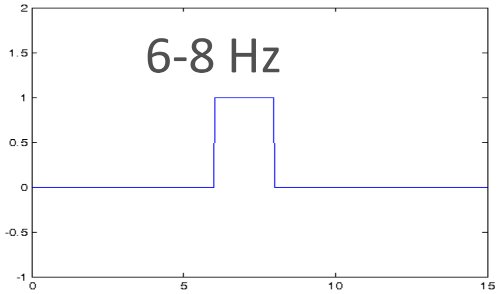

Selective Motion Magnification in Natural Videos

Source
(600 fps)

72-92 Hz
Amplified

100-120 Hz
Amplified

Low E (82.4 Hz)

A (110 Hz)

Motion Magnification Results

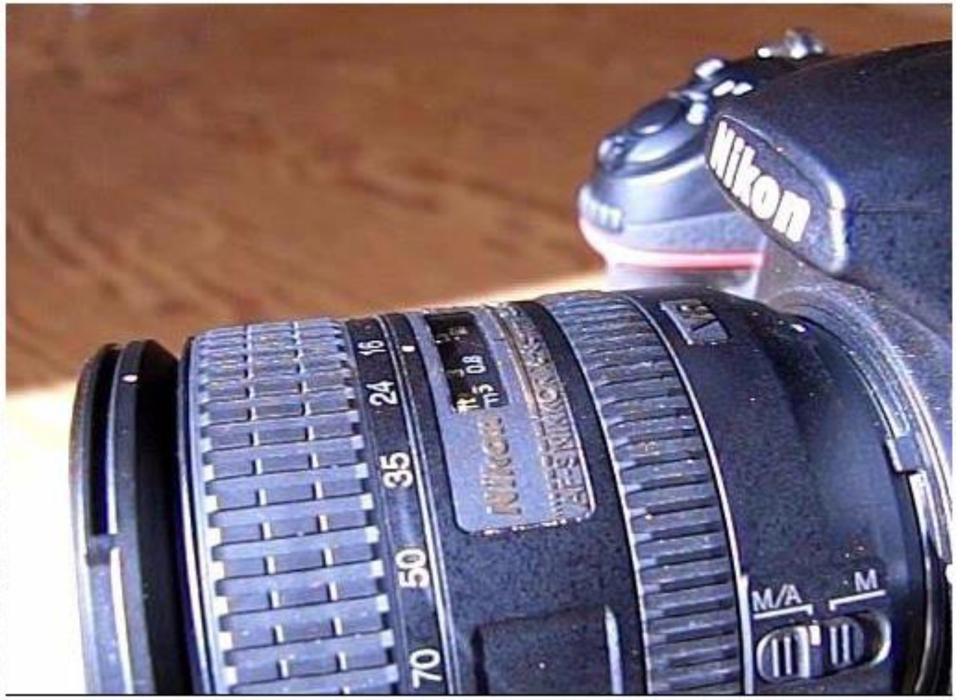

Source (300 fps)

Motion-magnified (45-100 Hz, x100)

DSLR Controlled Setup

Laser pointer

Reflected laser point

Source (300 fps)

Motion-magnified (45-100 Hz, x100)

Same DSLR camera
(Nikon D400)

Standard A4 sheet

Eulerian vs. Lagrangian Motion Magnification

Leonhard Euler

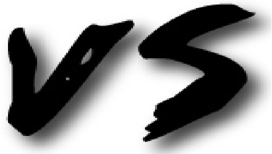

Joseph Louis Lagrange

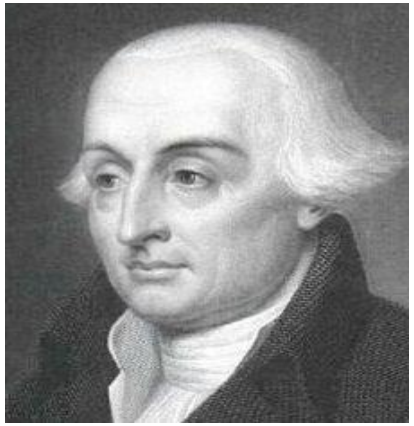

- Works better for smooth structures
- Supports smaller amplification factors
- Real time processing
- Unified framework to amplify spatial motion and purely temporal changes (e.g. heart pulse)
- Supports frequency-based processing and selective amplification

- Works better for point features
- Supports larger amplification factors
- Computationally intensive
- Optical flow may be inaccurate

Bruce Wayne’s Pulse

Batman Begins (2005), courtesy of Warner Bros. Pictures

Eulerian Video Magnification in the wild

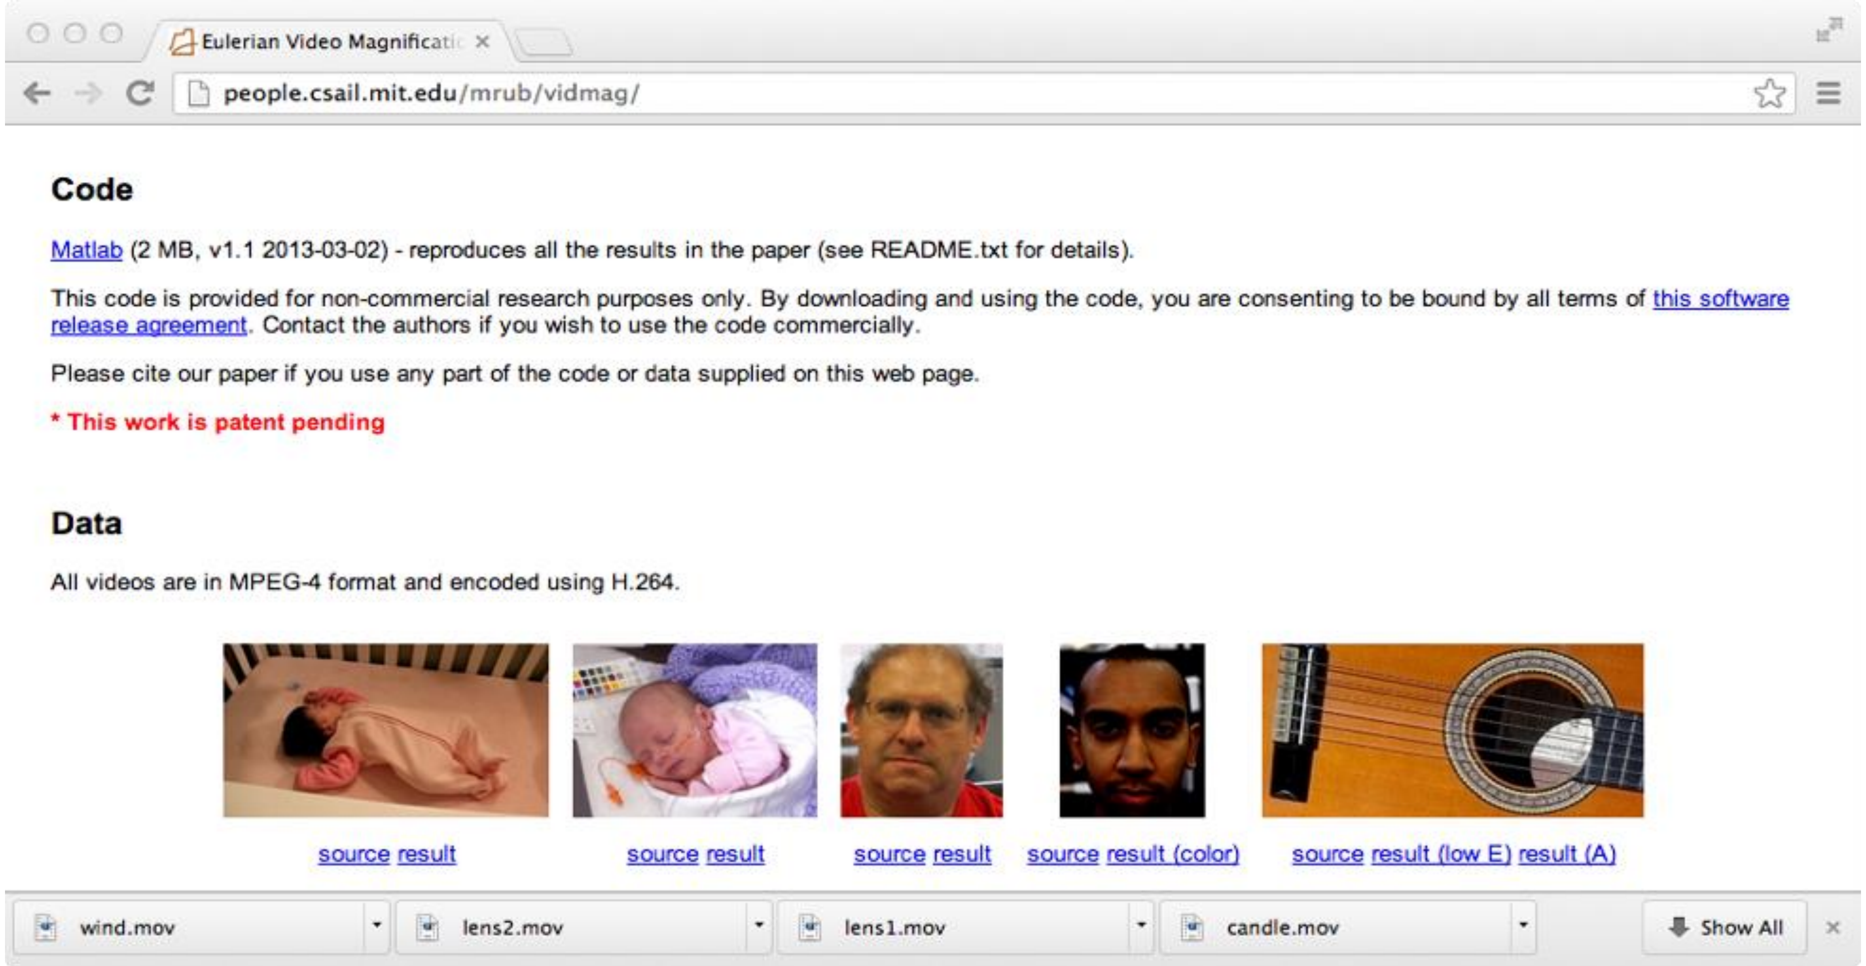

Code

Matlab (2 MB, v1.1 2013-03-02) - reproduces all the results in the paper (see README.txt for details).

This code is provided for non-commercial research purposes only. By downloading and using the code, you are consenting to be bound by all terms of this software release agreement. Contact the authors if you wish to use the code commercially.

Please cite our paper if you use any part of the code or data supplied on this web page.

* This work is patent pending

Data

All videos are in MPEG-4 format and encoded using H.264.

source result | source result | source result | source result (color) | source result (low E) result (A)

wind.mov | lens2.mov | lens1.mov | candle.mov | Show All

VideoScope by Quanta Research Cambridge

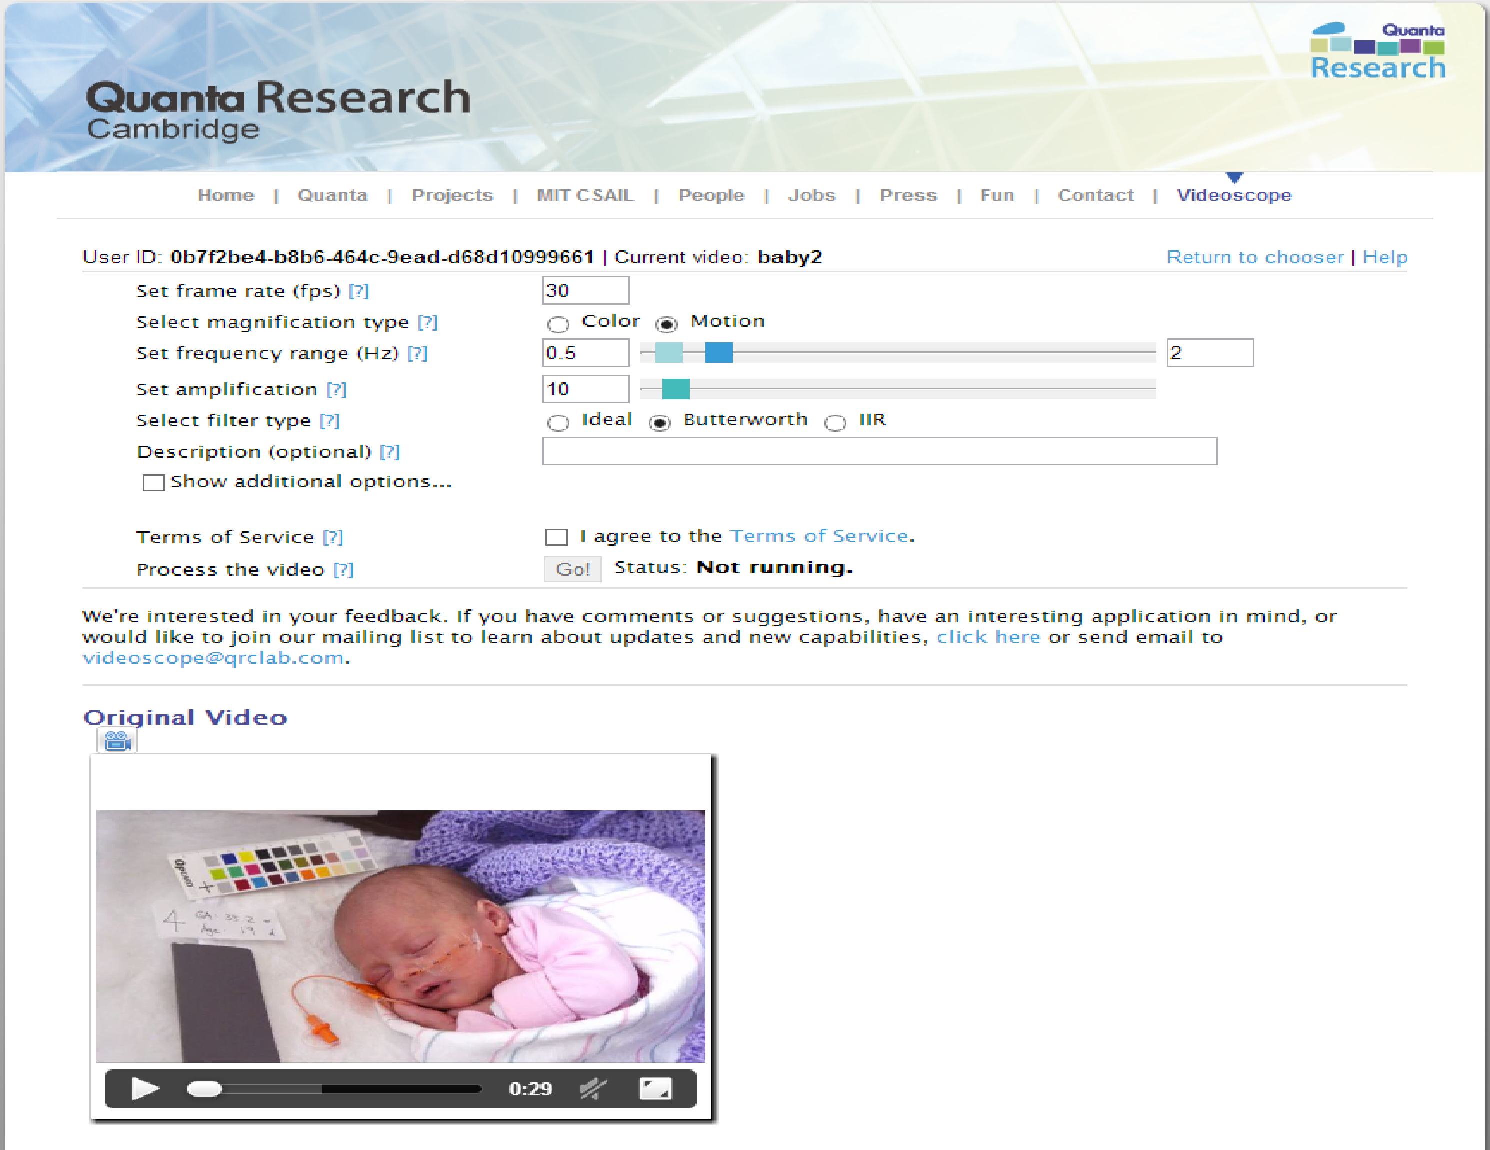

Independent (Real-time) Ports

“webcam-pulse-detector”
(Python + openCV)
+tracking

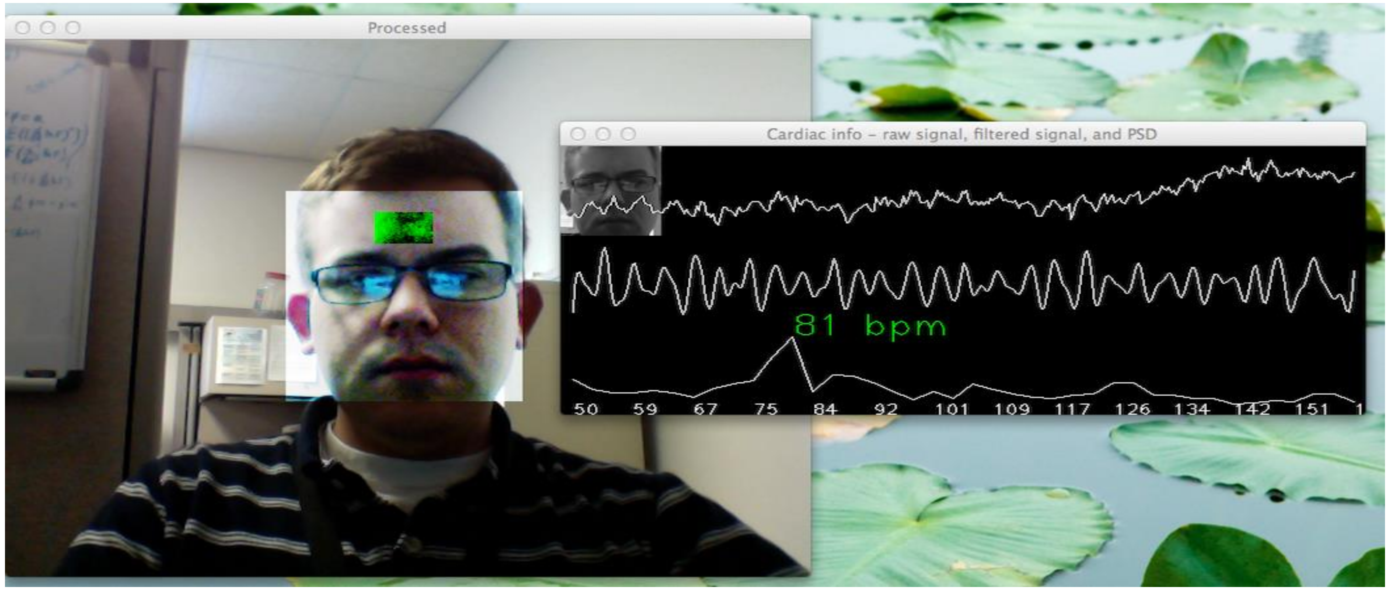

“VAmp - Video Amplifier”
(Java)

EVM in the Wild: Pregnancy

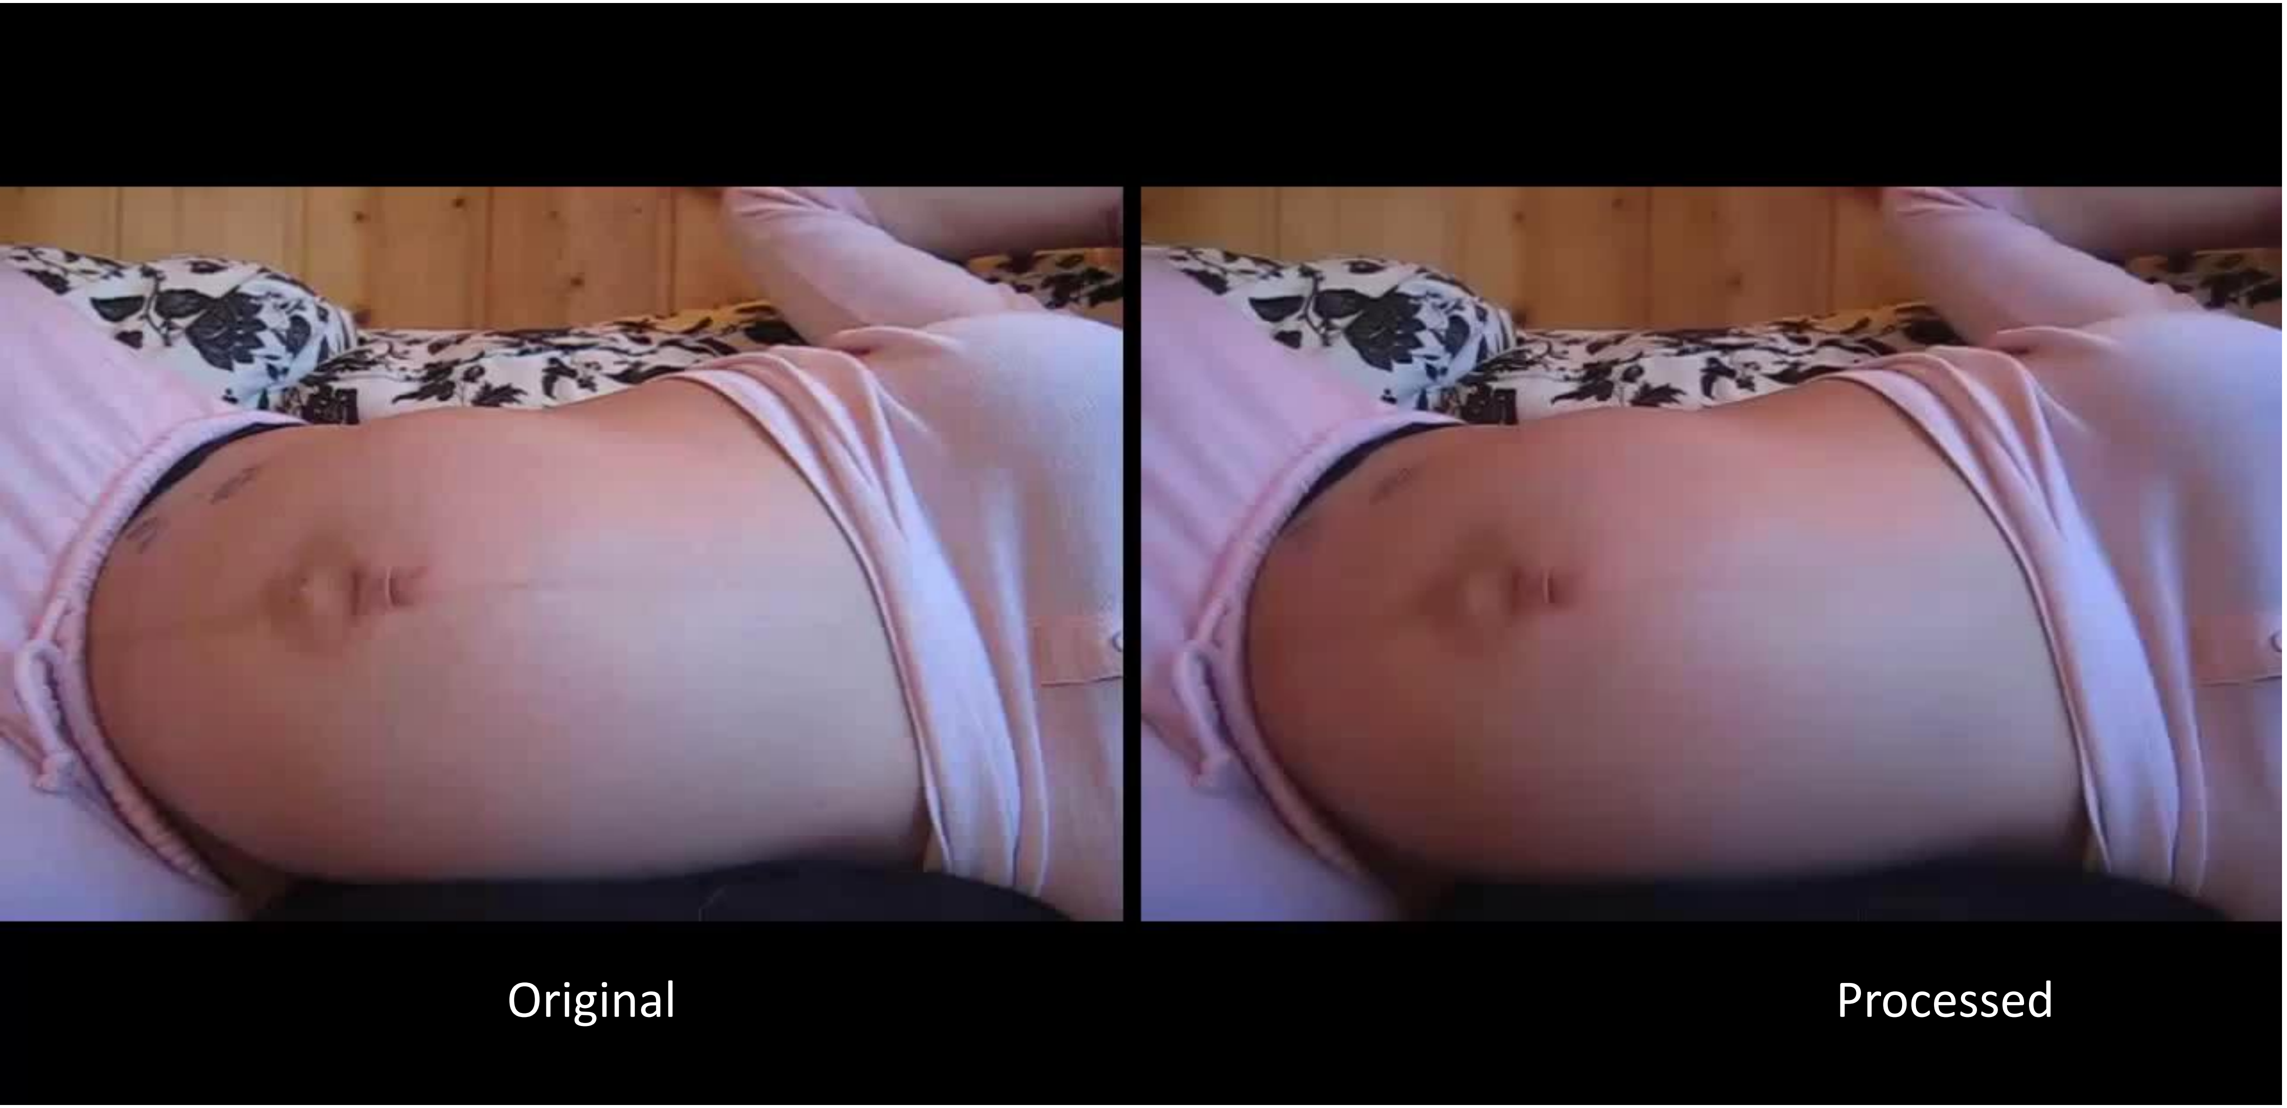

“Tomez85”

EVM in the Wild: Blood flow Visualization

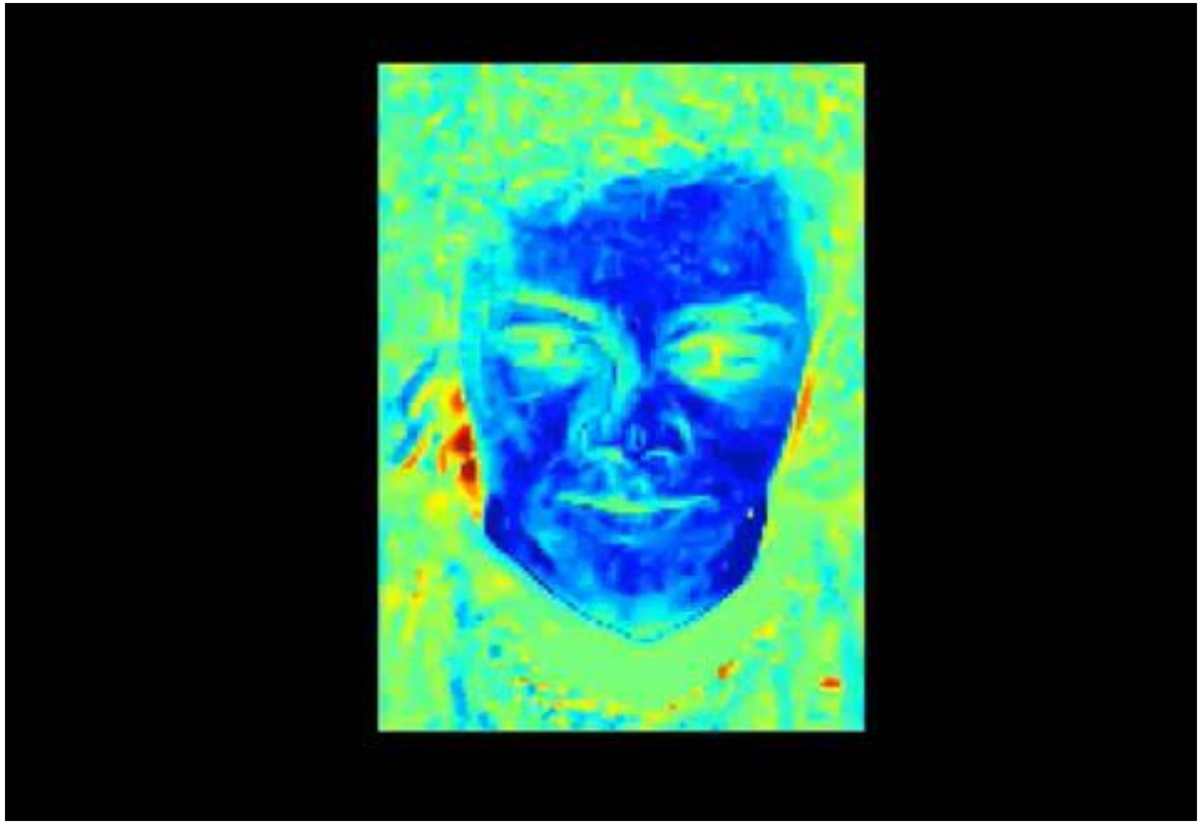

Red = high blood volume
Blue = low blood volume

Institute for Biomedical Engineering, Dresden Germany

EVM in the Wild: Guinea Pig!

Source

Motion-magnified

“SuperCreaturefan”: “Guinea pig Tiffany is the first rodent on Earth to undergo Eulerian Video Magnification.”

Phase-based Pipeline (SIGGRAPH'13)

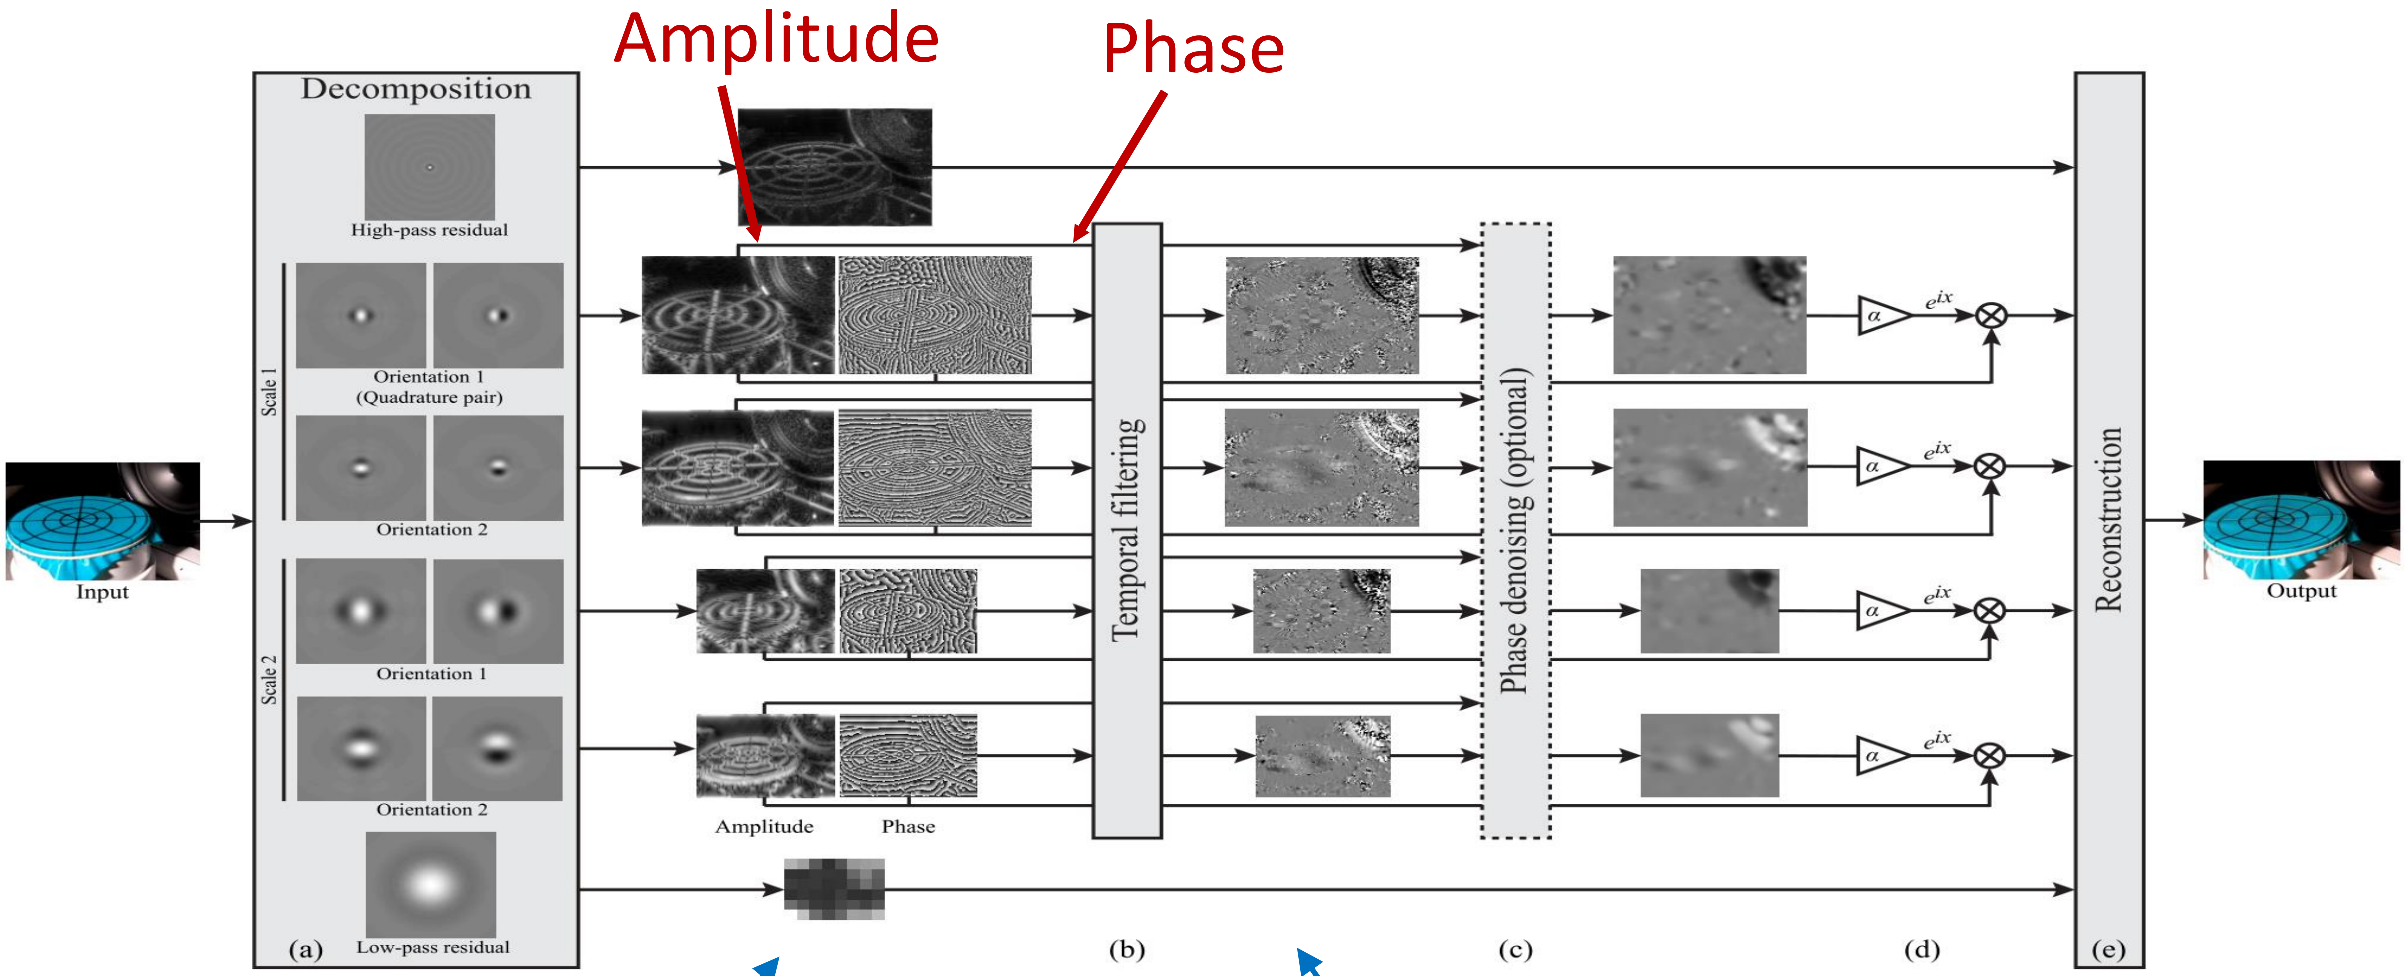

Complex steerable pyramid
[Simoncelli and Freeman 1995]

Temporal filtering on **phases**

Steerable Pyramid (Freeman and Adelson 1991)

- Set of filters to locally decompose image into different spatial bands
- Filters indexed by scale ω and orientation θ
- Transfer functions partition frequency domain
- Apply filters $T_{\omega,\theta}$ to image $I(x, y)$ by doing

$$S_{\omega,\theta} = \underbrace{\mathcal{F}^{-1}}_{\text{To Primal Basis}} \left\{ \underbrace{T_{\omega,\theta}}_{\text{Isolate components}} \times \underbrace{\mathcal{F}\{I\}}_{\text{To Fourier Basis}} \right\}$$

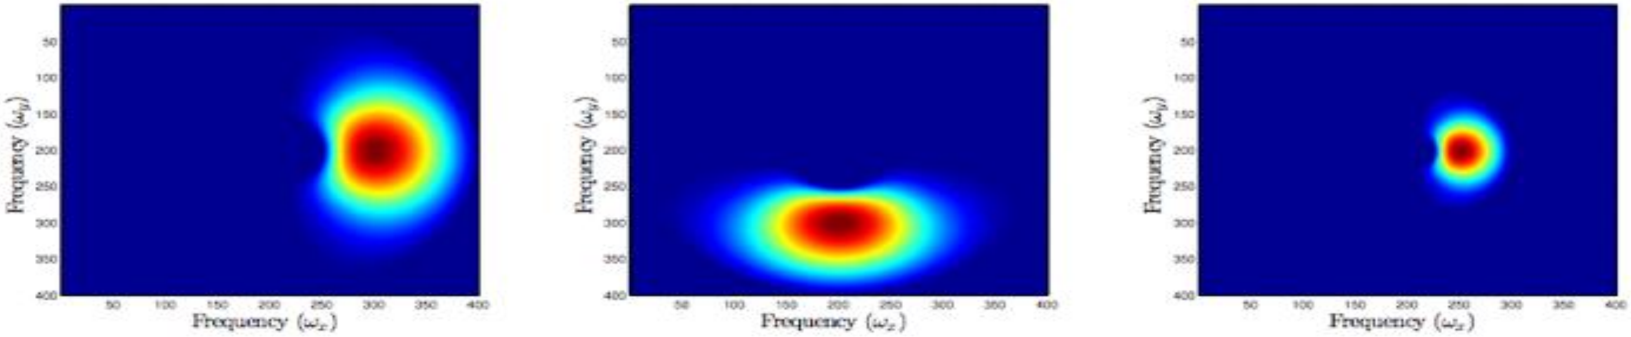

$\theta = 0$, $\omega = 100$ $\quad$ $\theta = \frac{\pi}{2}$, $\omega = 100$ $\quad$ $\theta = 0$, $\omega = 50$

Sample Steerable Pyramid Transfer Functions $T_{\omega,\theta}$

Decomposition into Wavelets

- Locally, image $I(x, y)$ gets decomposed into oriented wavelets

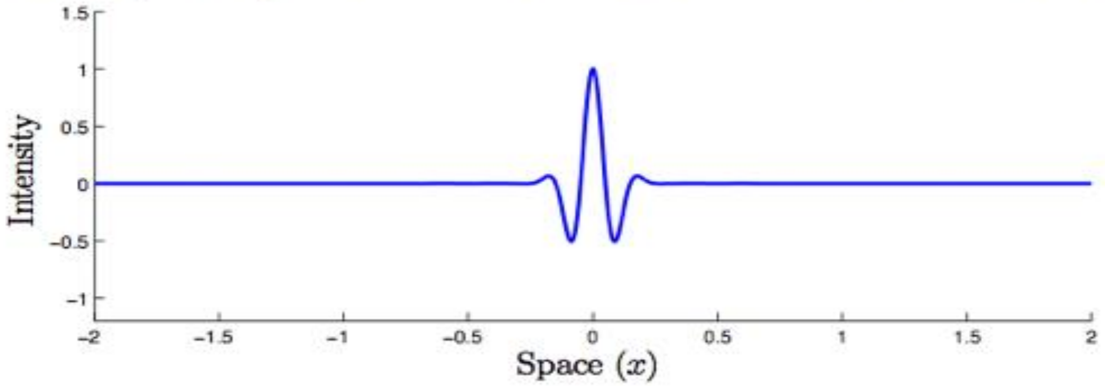

- Technique shifts phase of wavelets in all spatial bands

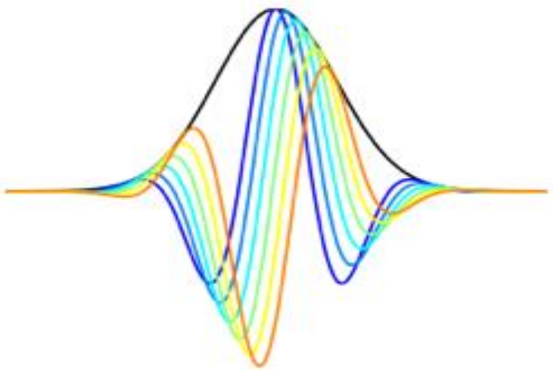

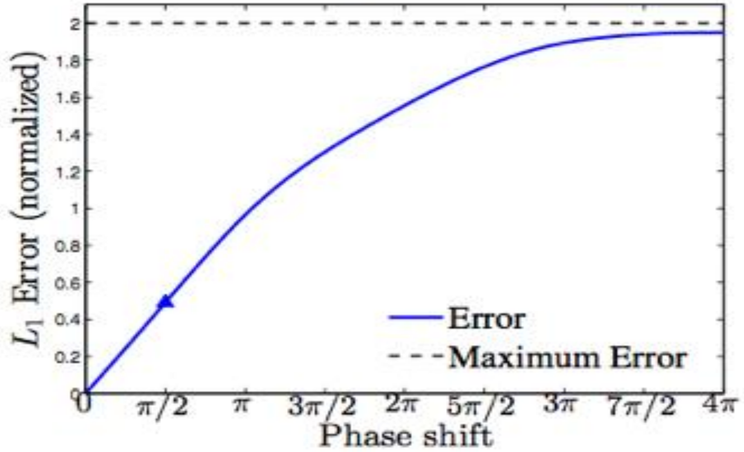

Phase changes from 0 (Blue) to $\frac{3\pi}{4}$ (Orange)

- Phase change less than $\frac{\pi}{2}$ has reasonable error
- Wavelet "wraps around"

Results: Phase-based vs. Linear

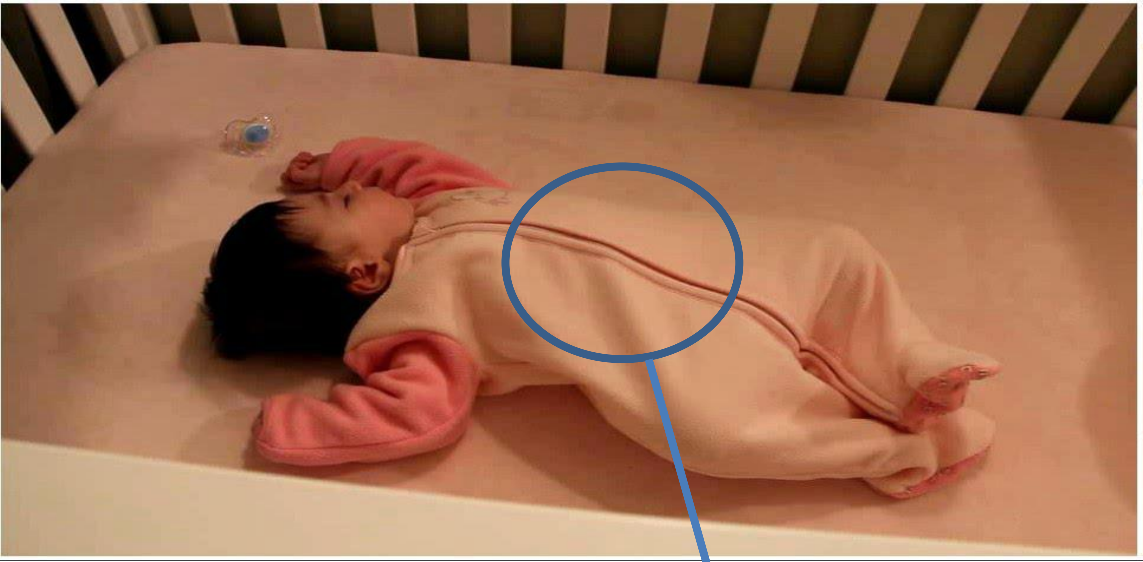

Linear (SIGGRAPH'12)

Phase-based (SIGGRAPH'13)

Clipping artifacts near Sharp edges and larger motions

Results: Phase-based vs. Linear

Linear (SIGGRAPH’12)

Phase-based (SIGGRAPH’13)

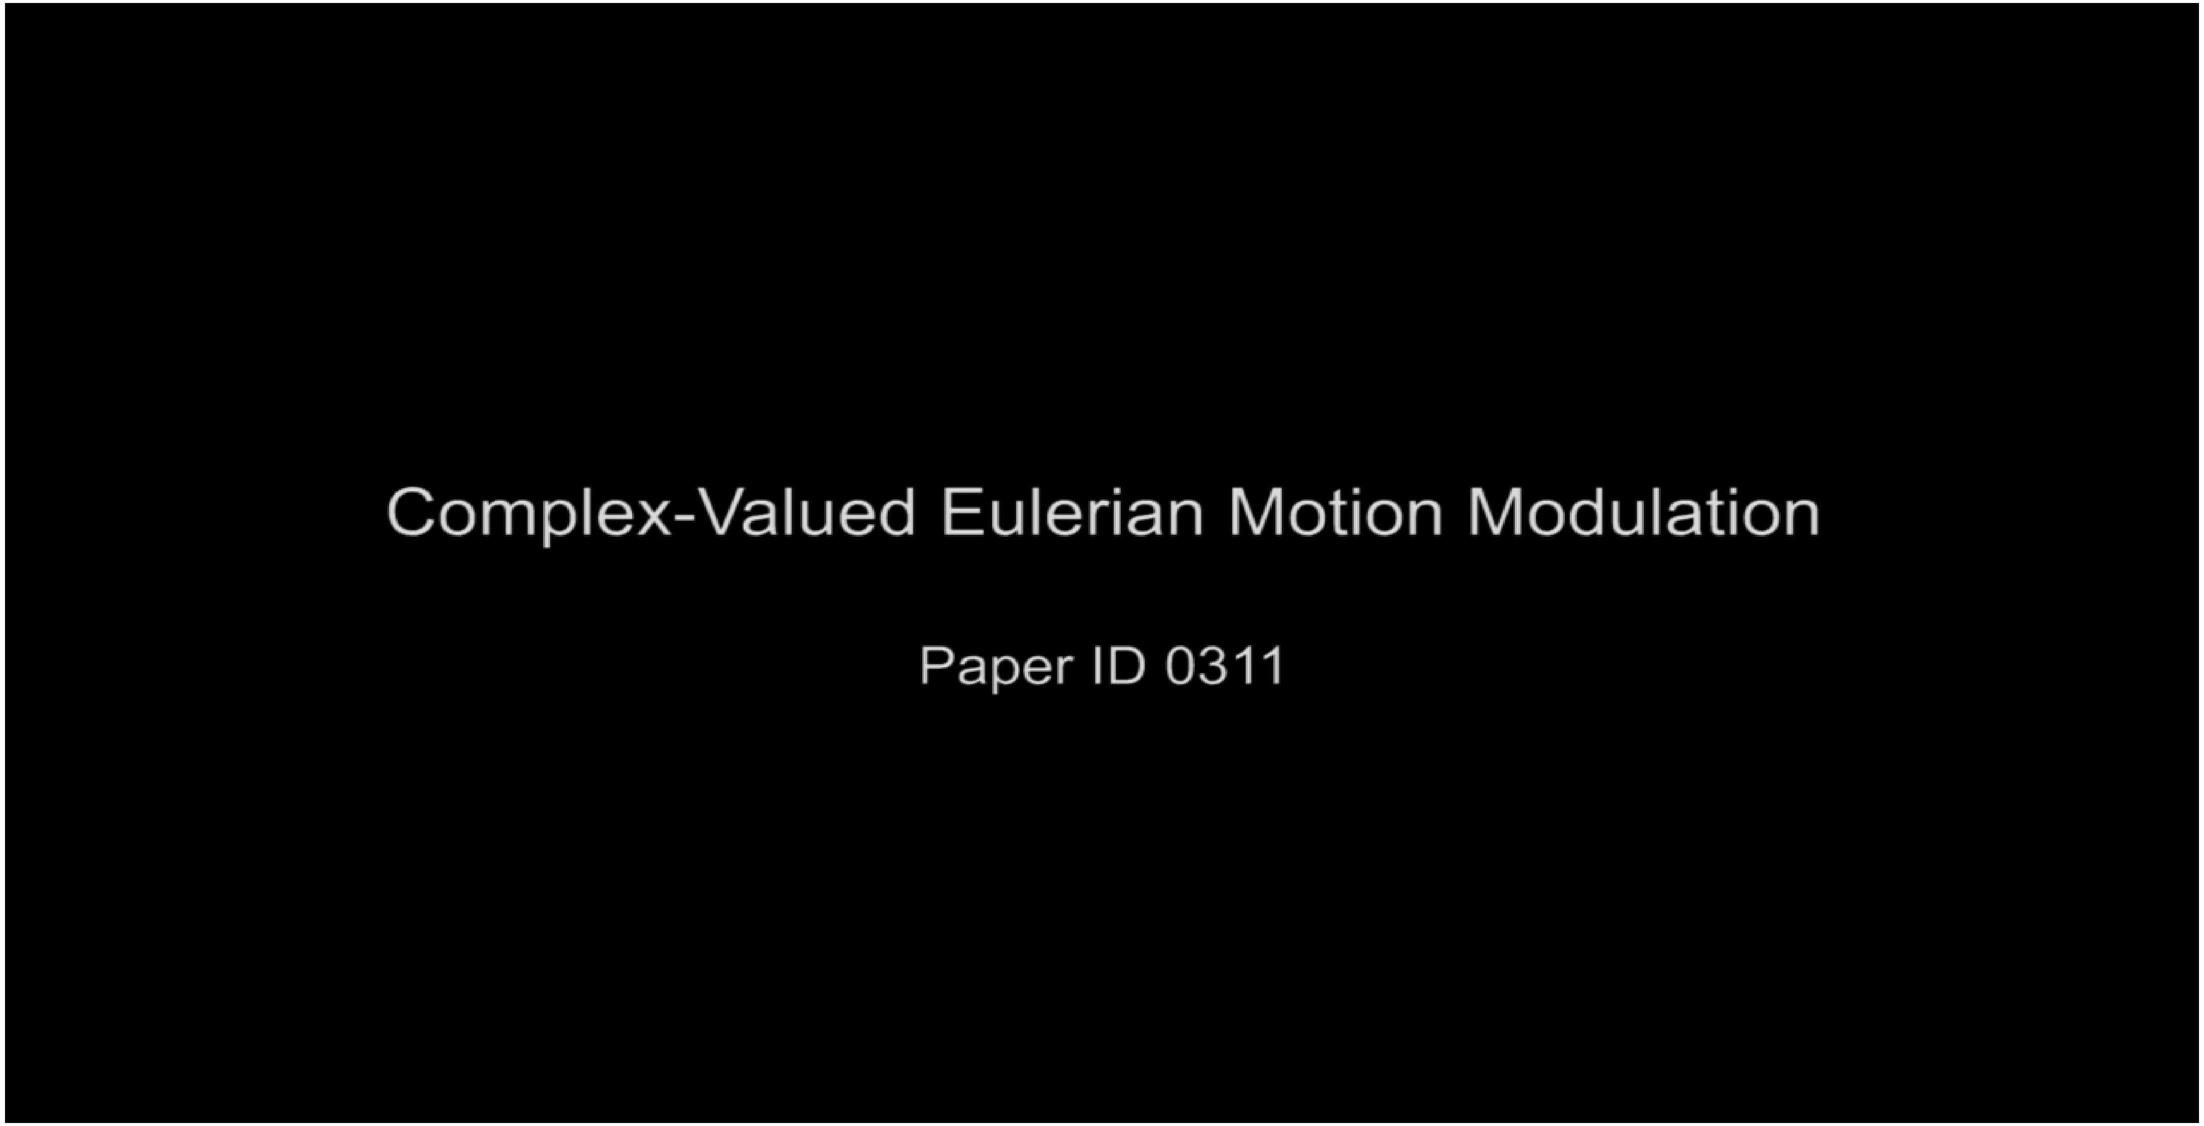
Complex-Valued Eulerian Motion Modulation
Paper ID 0311

Frequency sweep

Frequency sweep

Phase-based Motion Attenuation

Source

Linear

Amplifies color
And motion **jointly**

Motion attenuation +
Color amplification

Amplifies color
Without amplifying motion

Revealing Invisible Changes in the World

- NSF International Science and Engineering Visualization Challenge (SciVis), 2012
- Science Vol. 339 No. 6119 Feb 1 2013

High-speed video, singing single note

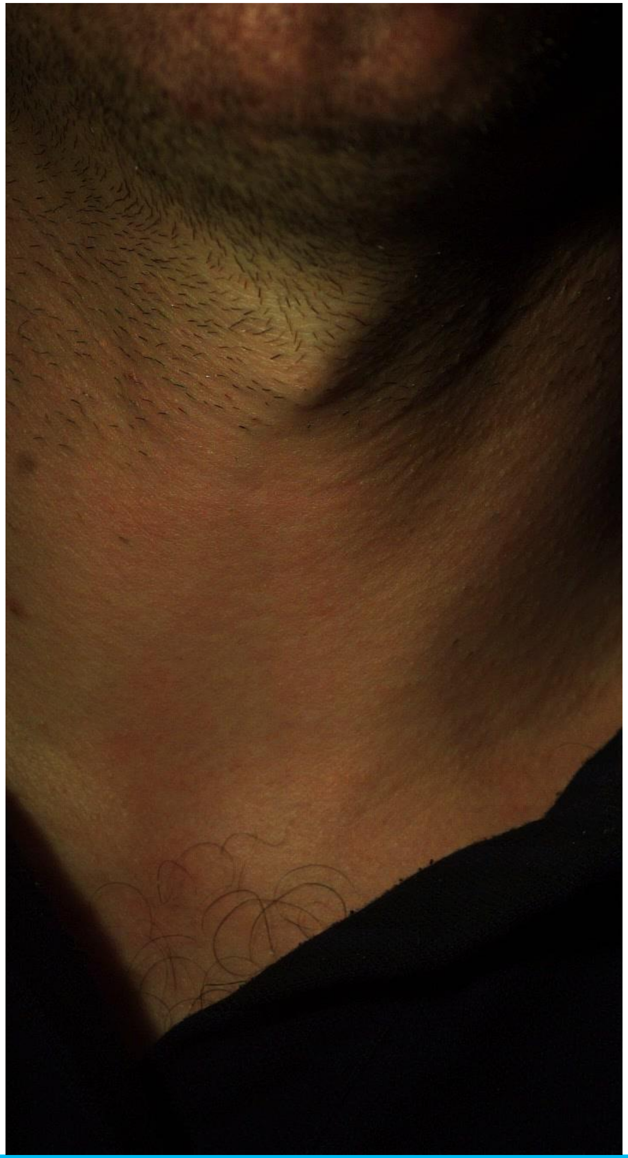

Fundamental note, motion magnified x100

1st harmonic motion magnified x100

Non-harmonic frequency, x100

Motion magnification.
Ce Liu, Antonio Torralba, Bill Freeman,
Fredo Durand, and Edward Adelson

We can register, then amplify, one motion *relative to* another.

empty trunk

full trunk

Original footage courtesy of Paul Robertson, BBN.

Motion magnification.
Ce Liu, Antonio Torralba, Bill Freeman, Fredo Durand, and Edward Adelson

We can register, then amplify, one motion *relative to* another.

empty trunk

full trunk (motion difference amplified)

Original footage courtesy of Paul Robertson, BBN.

Research project pages: "Eulerian video magnification"
"Phase-based video magnification"

Joint work with Michael Rubinstein, Hao Yu, Eugine Hsu, Neal Wadhwa, John Guttag, Fredo Durand

Massachusetts Institute of Technology

Selective Motion Magnification

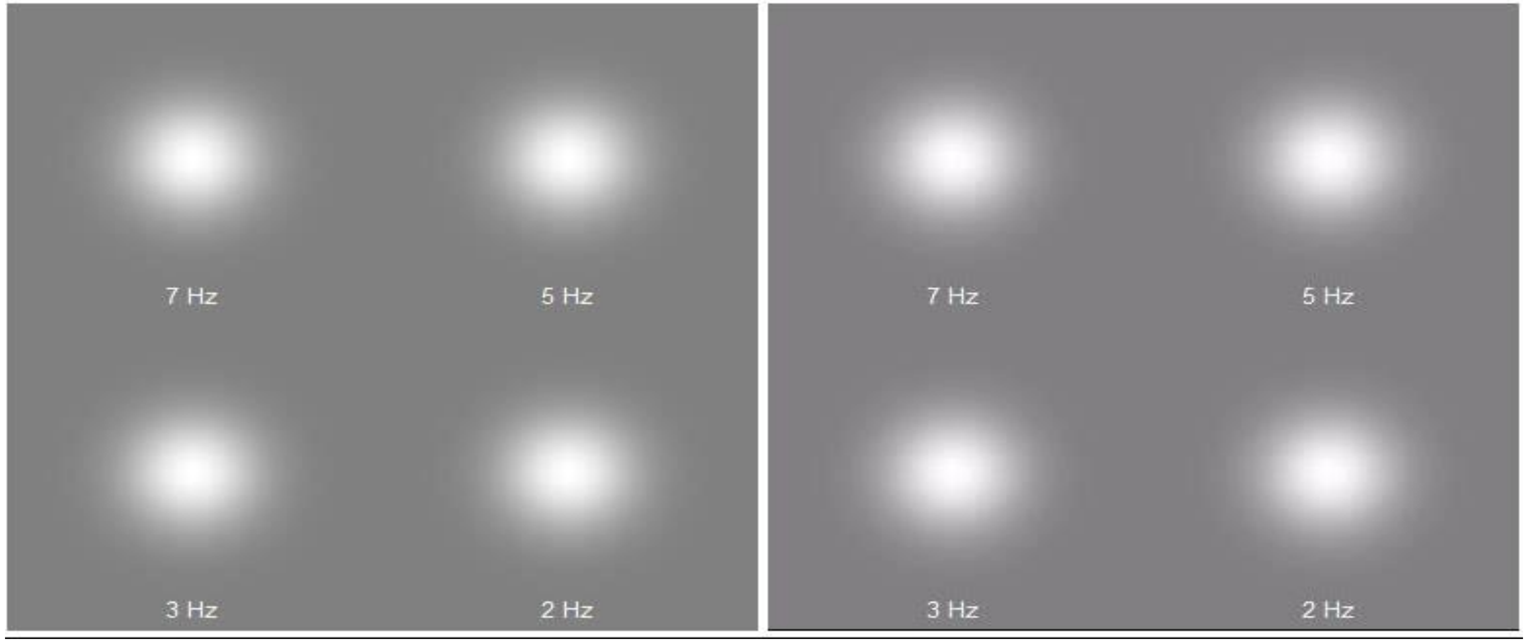

Source

Motion-magnified (2 Hz)

Temporal filter:

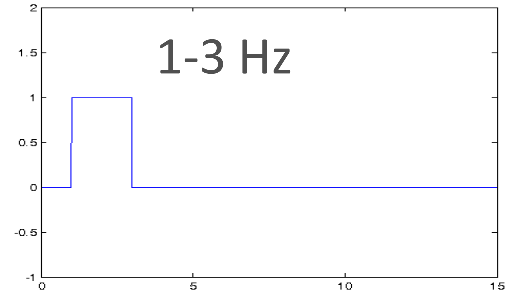

Selective Motion Magnification

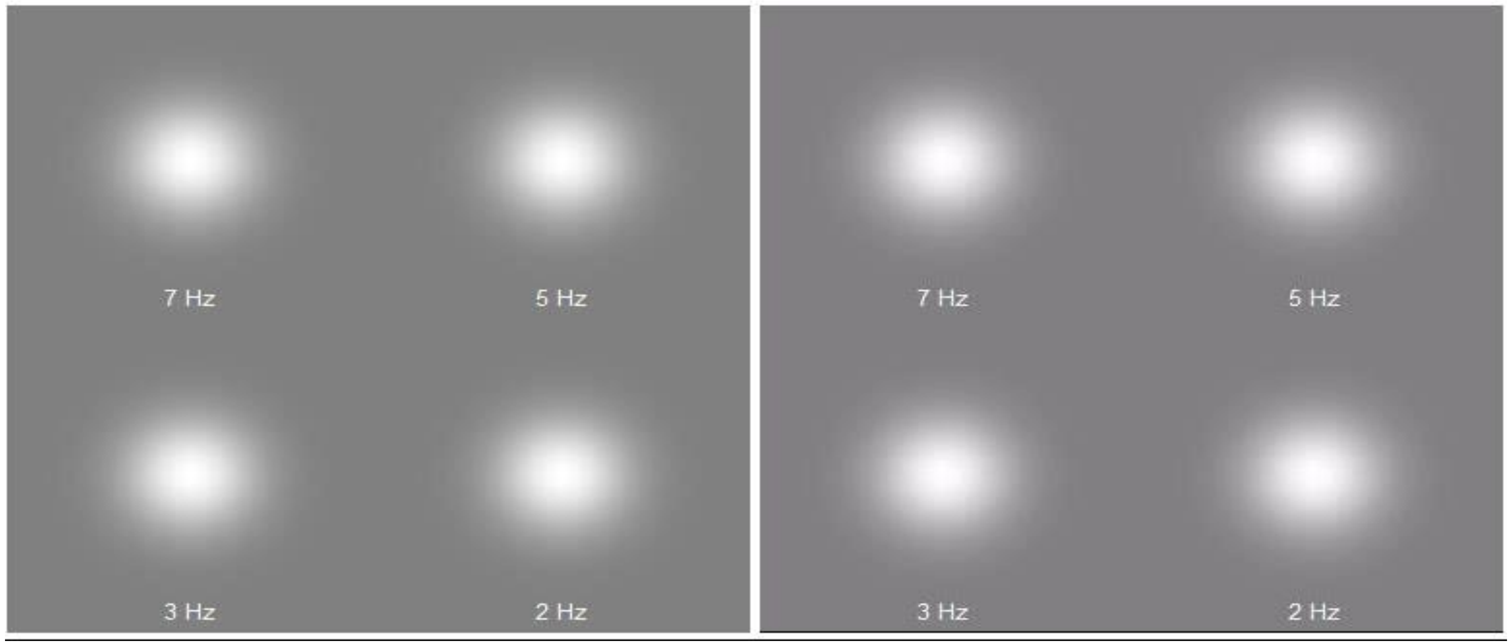

Source

Motion-magnified (3 Hz)

Temporal filter:

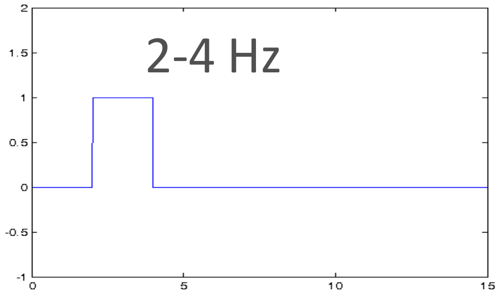

Selective Motion Magnification

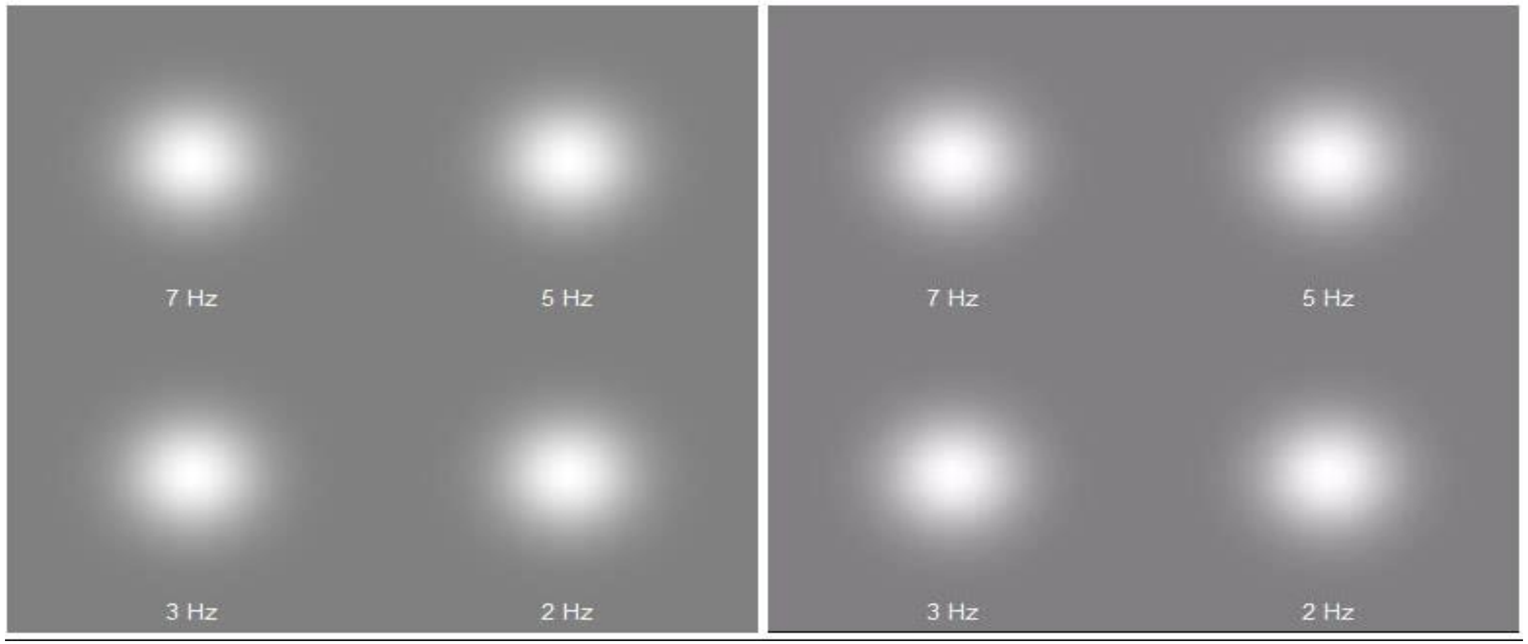

Source

Motion-magnified (5 Hz)

Temporal filter:

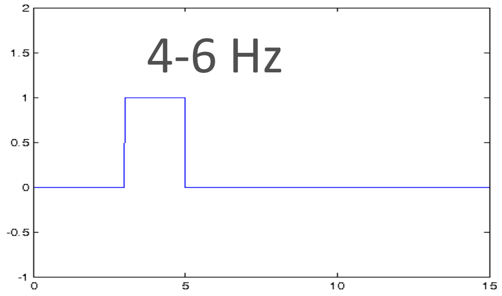

Selective Motion Magnification

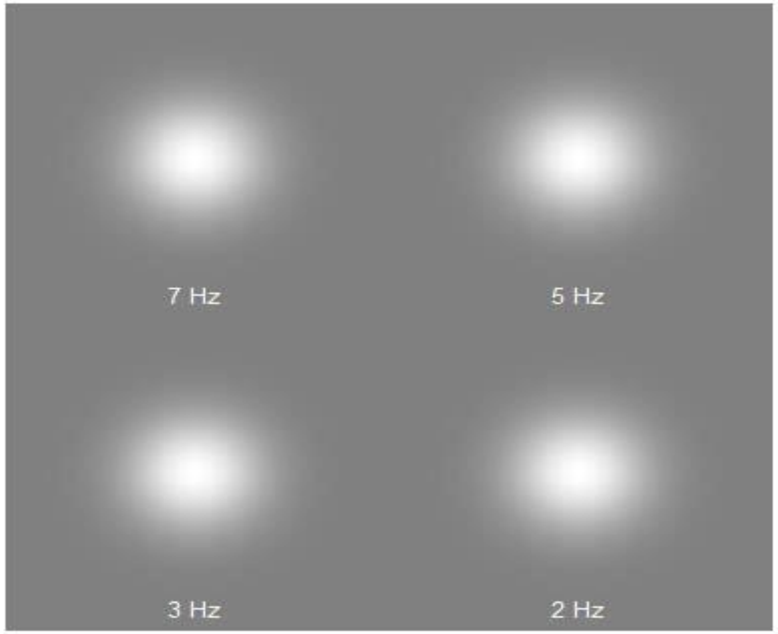

Source

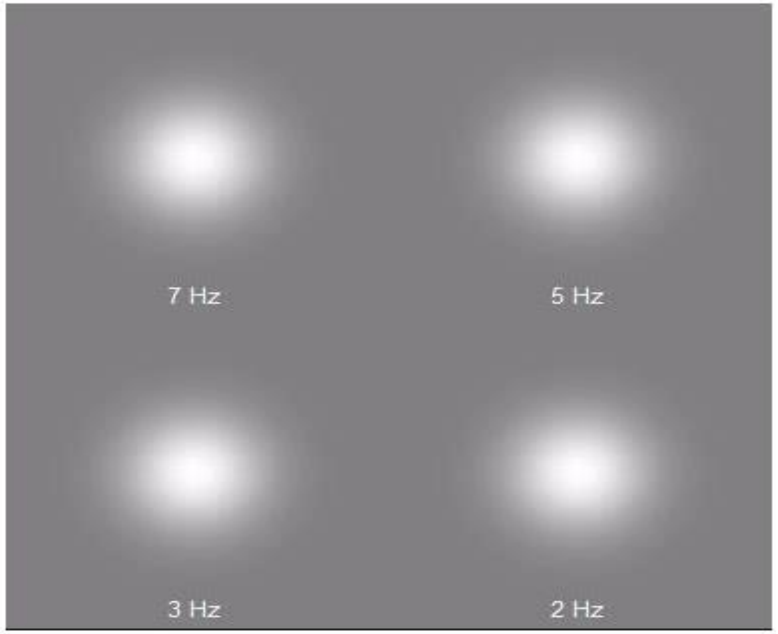

Motion-magnified (7 Hz)

Temporal filter: