How to Give a Good Talk

Vincent Lepetit

université de BORDEAUX

A talk is *not* a paper presented in oral form

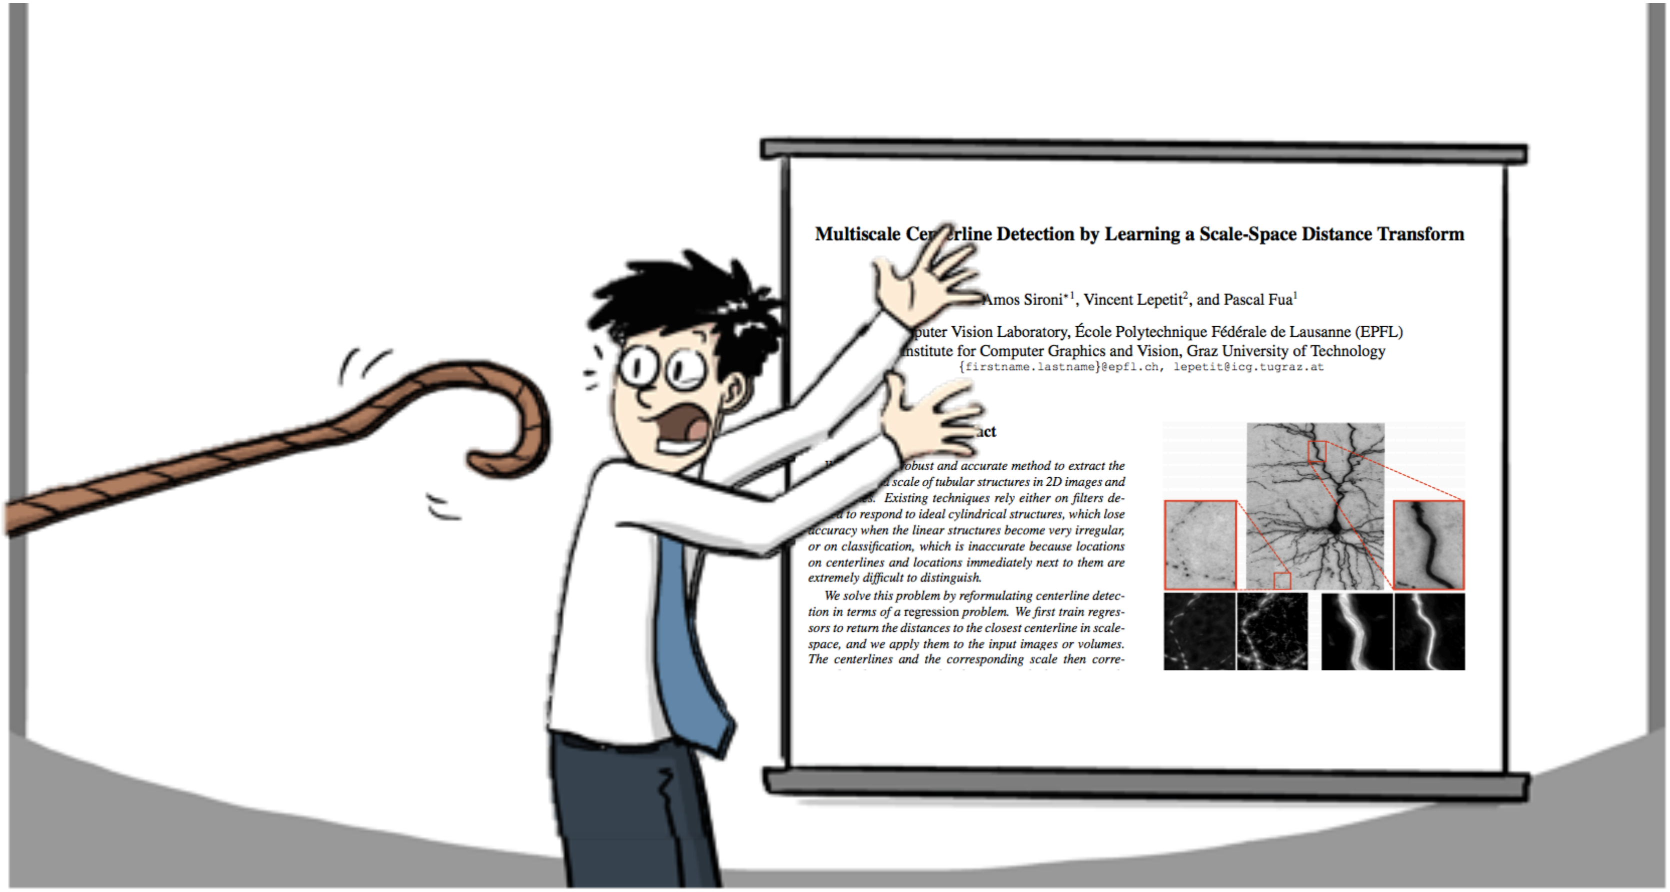

from PhD Comics

université de BORDEAUX

A talk is *not* a paper presented in oral form.

A talk is mostly an *advertisement* for the paper:

Make the audience understand your conclusion and the intuition behind your method, not the technical details.

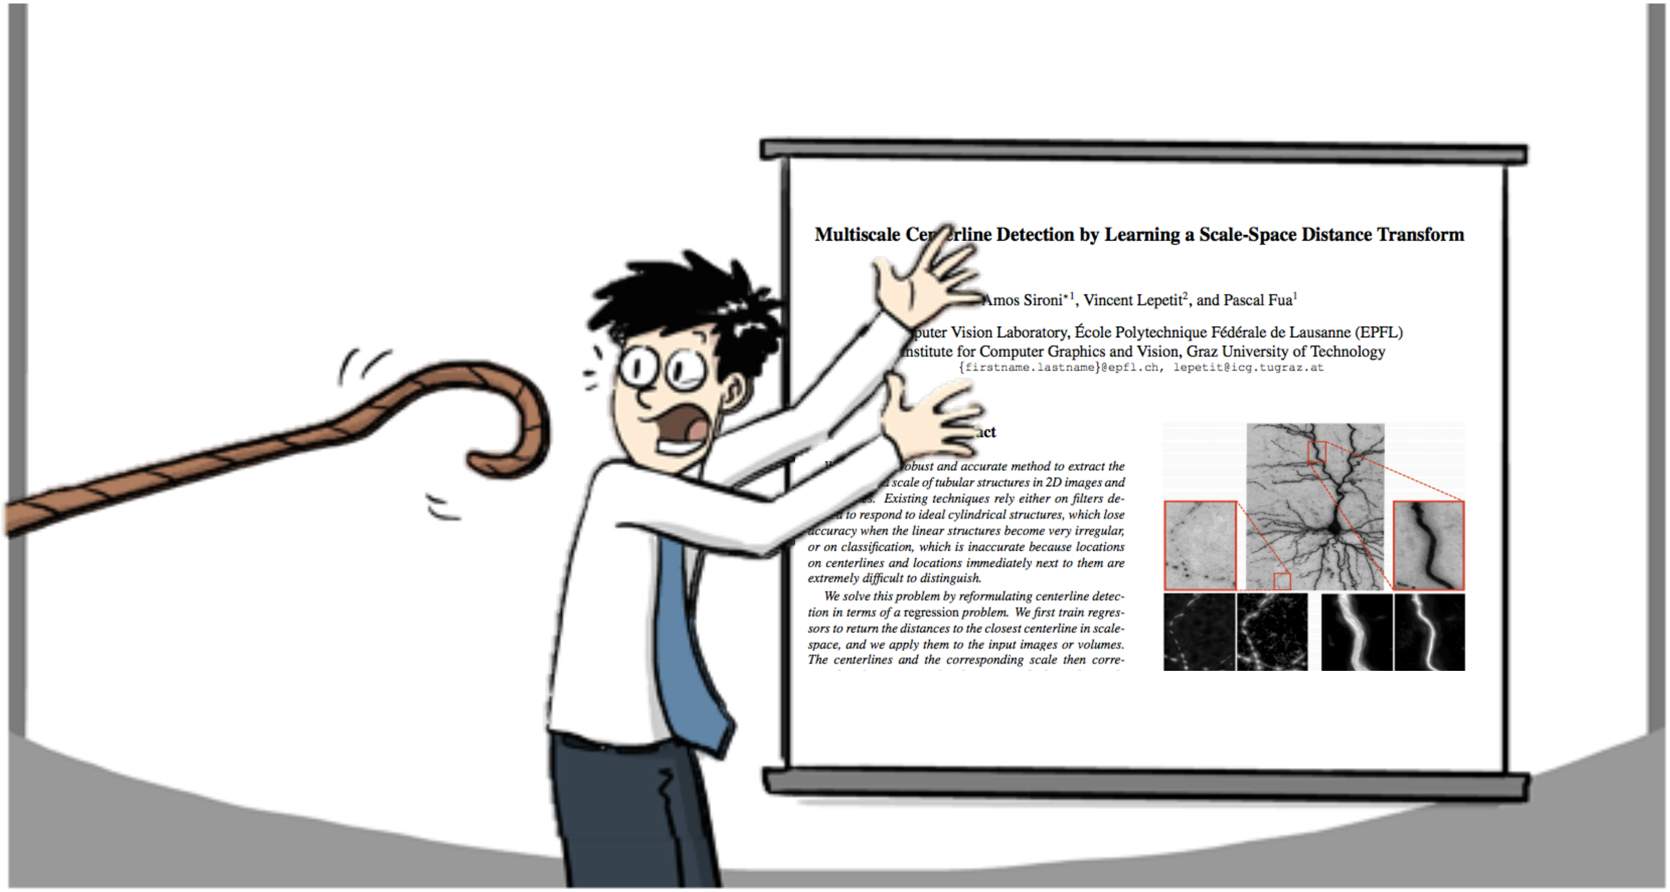

A talk is ***not*** "a paper presented in oral form".

A talk is mostly an ***advertisement*** for the paper: Make the audience understand your conclusion and the intuition behind your method, not the technical details.

Give your audience an intuitive feel for your idea

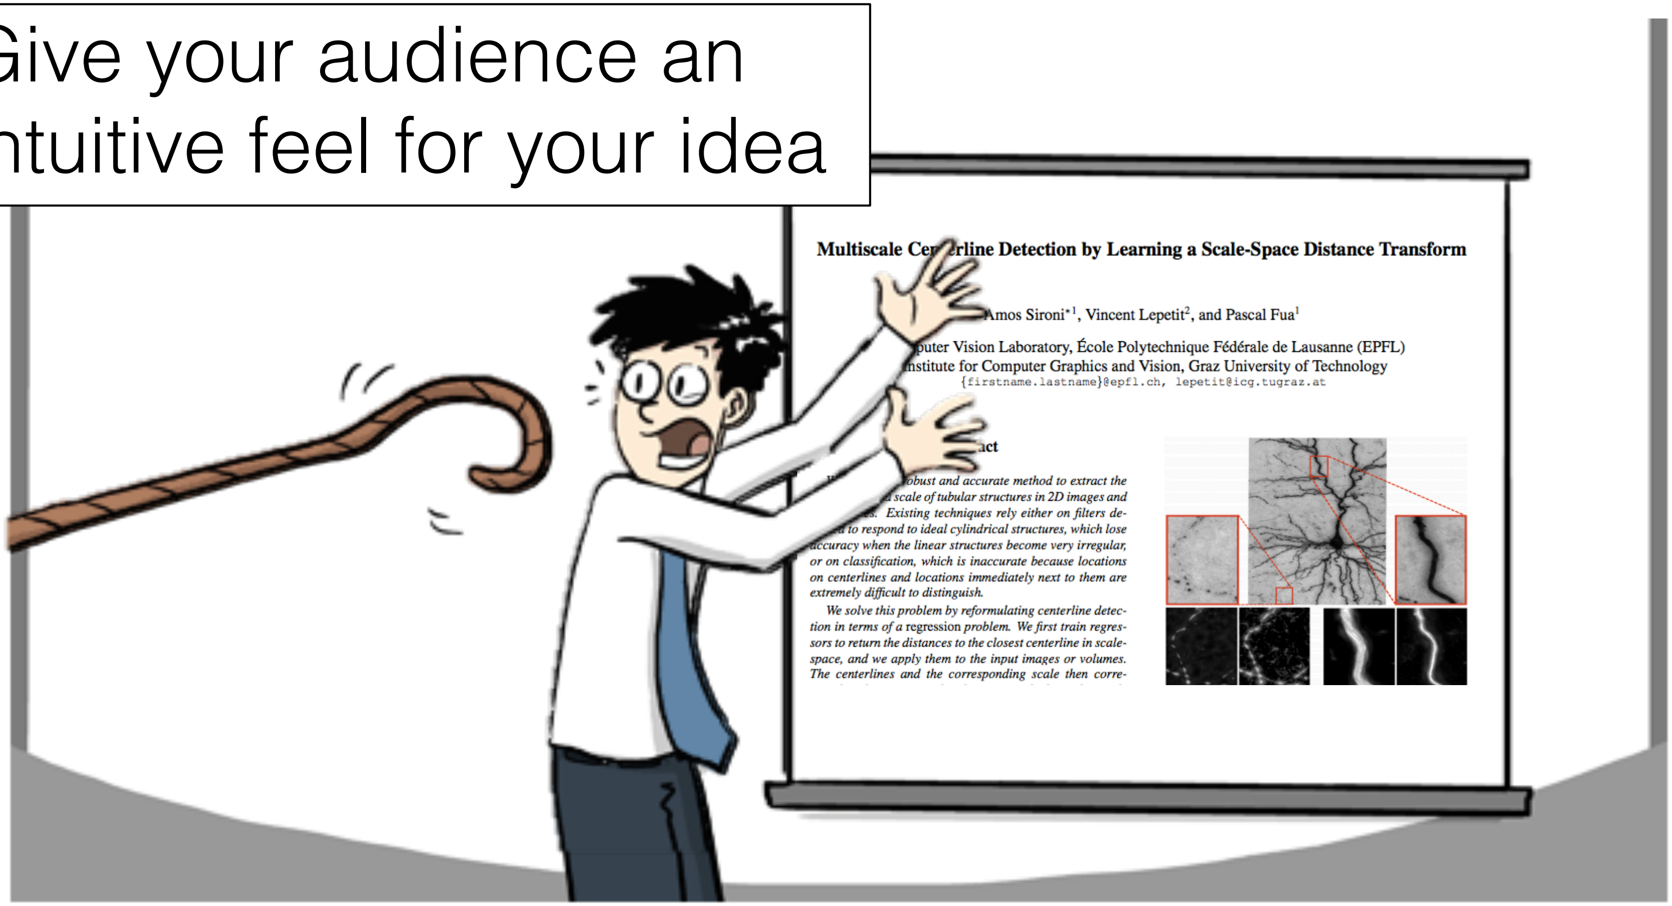

A talk is *not* "a paper presented in oral form".

A talk is mostly an ***advertisement*** for the paper: Make the audience understand your conclusion and the intuition behind your method, not the technical details.

Give your audience an intuitive feel for your idea

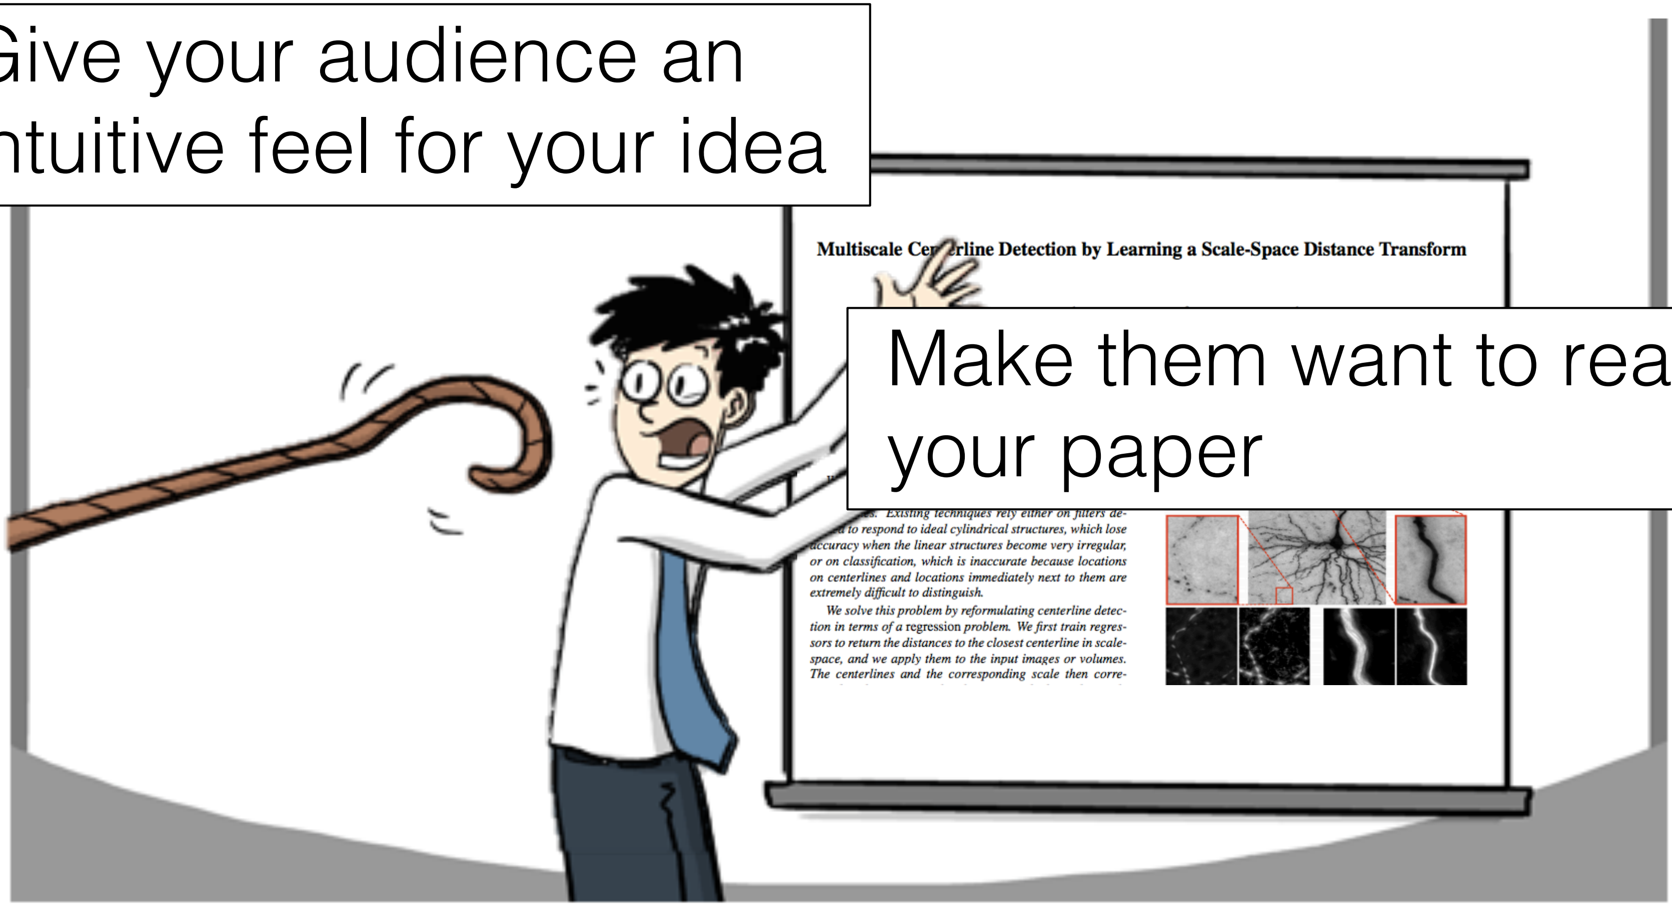

Make them want to read your paper

A talk is *not* "a paper presented in oral form".

A talk is mostly an ***advertisement*** for the paper: Make the audience understand your conclusion and the intuition behind your method, not the technical details.

Give your audience an intuitive feel for your idea

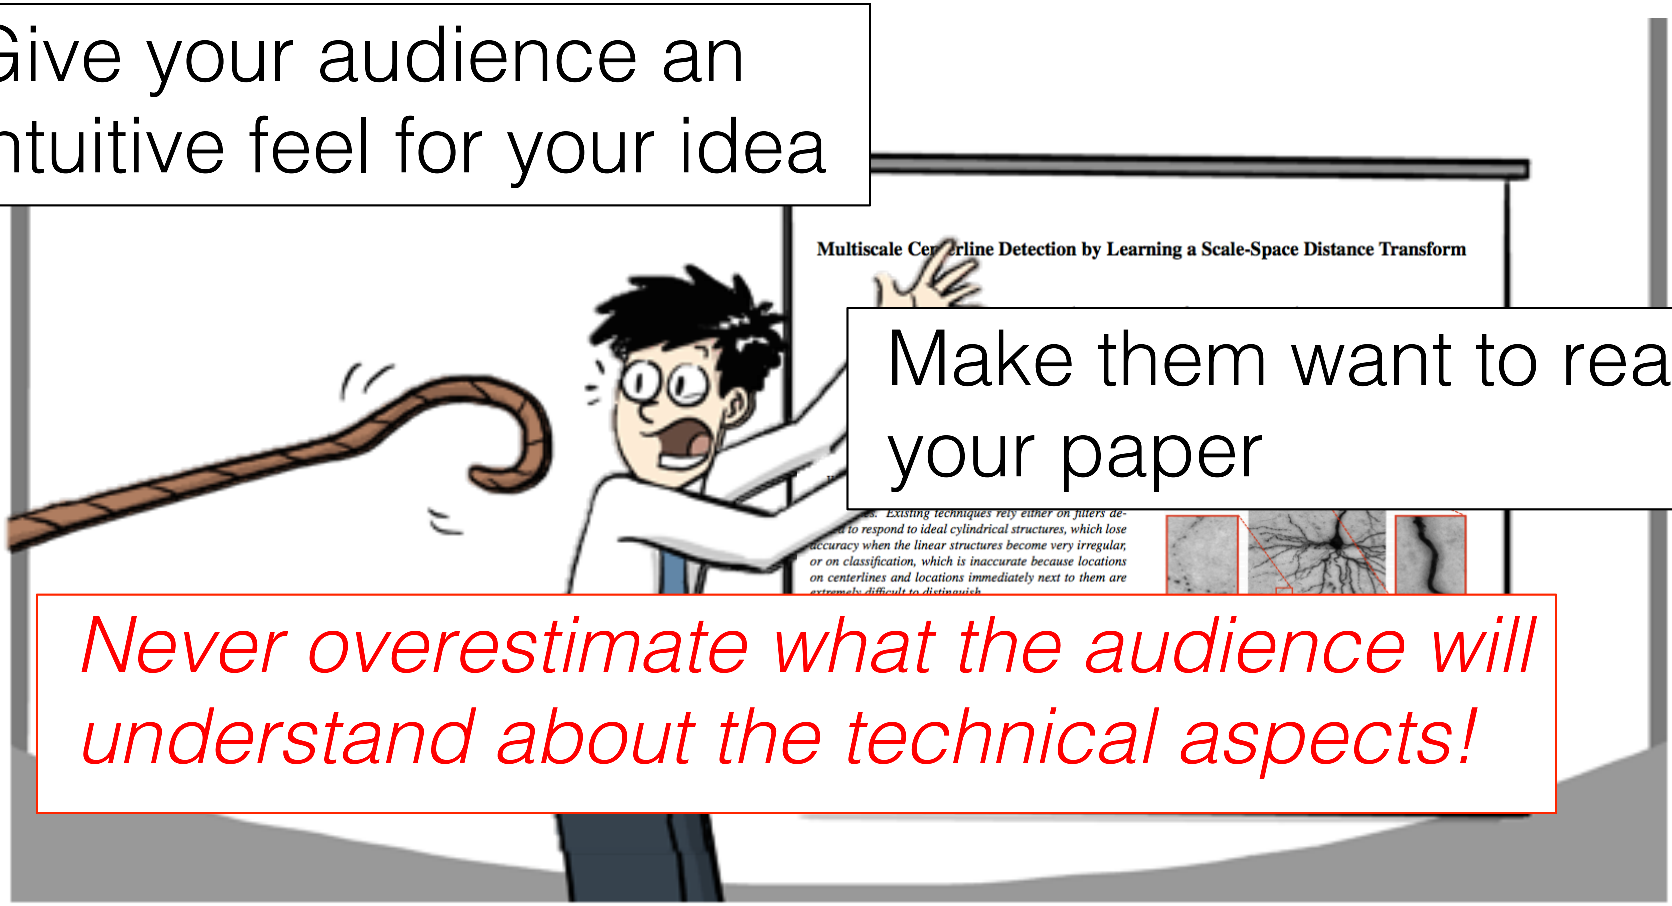

Make them want to read your paper

Never overestimate what the audience will understand about the technical aspects!

About the Style

- Make sure the text is easy to read: example;
- Avoid a blank background, but don't use an overloaded template;
- Show the slide number;
- Use sans-serif fonts (Arial, Calibri, Helvetica Light, not Times);
- Font size at least 24pt (this is 28pt).

université de BORDEAUX

Tell a story:

Put yourself in the shoes of a person that knows nothing about your work.

université de BORDEAUX

You have only a few minutes to get people interested

What is the problem you are aiming to solve?

Why is it interesting to solve it?

Why is it difficult to solve it?

about 20% of the talk { Motivation

{ Contribution, results, comparison

Start with some results!

What is the problem you are aiming to solve?

Why is it interesting to solve it?

Why is it difficult to solve it?

about 20% of the talk { Motivation

Contribution, results, comparison

Scalable Recognition with a Vocabulary Tree (2006)

David Nistér, Henrik Stewénius

Query Images for the Web

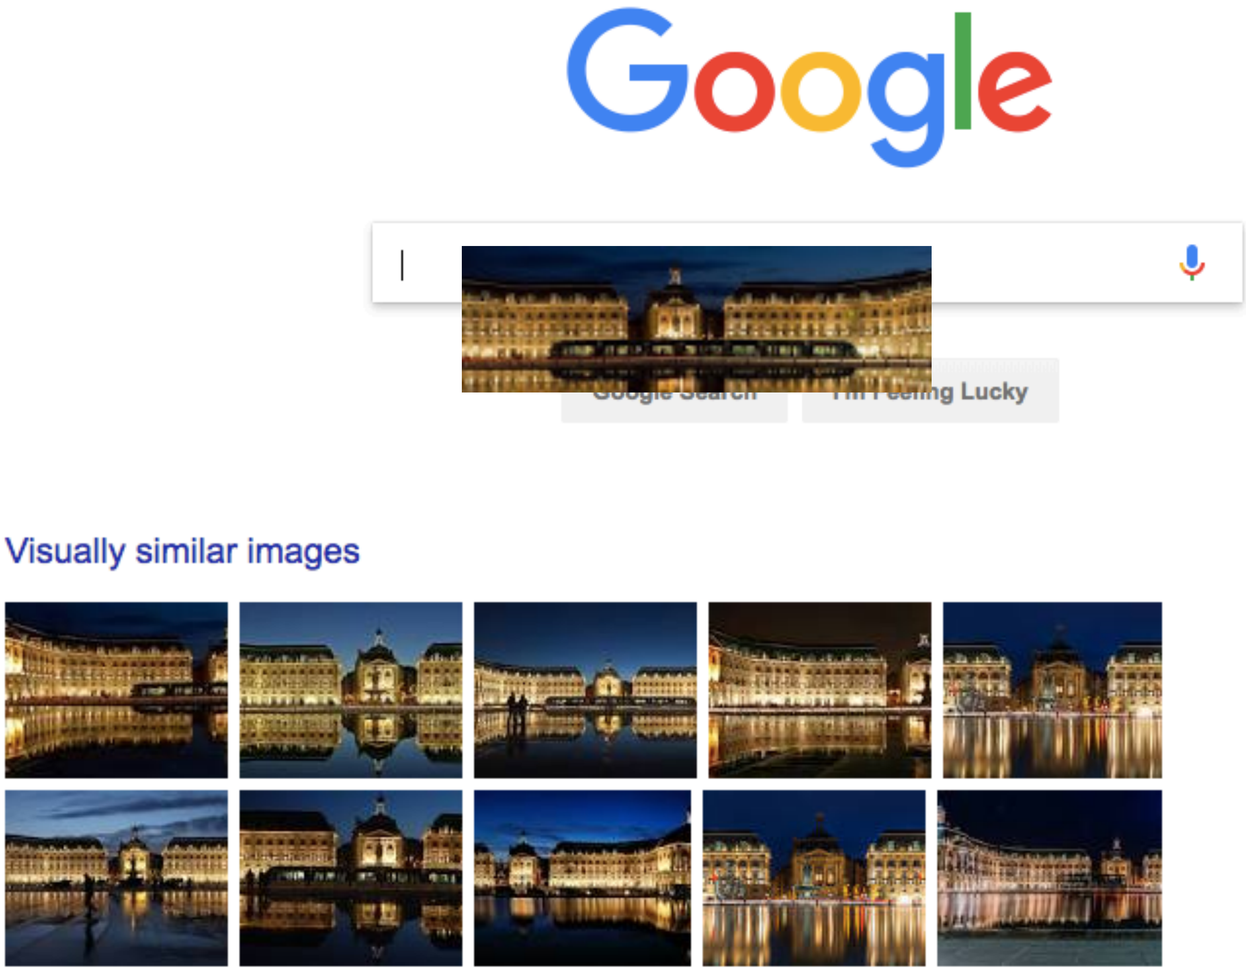



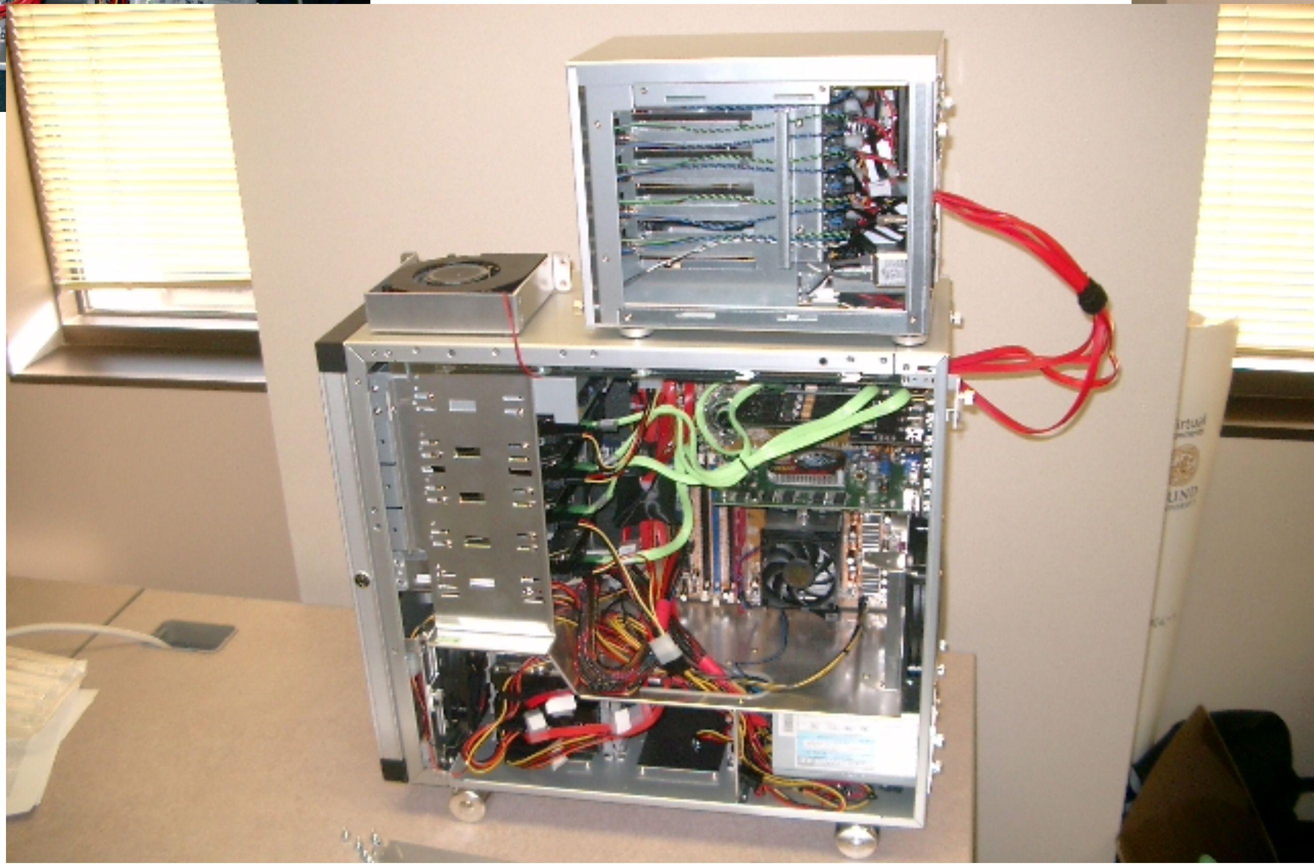

image query in a database of
110,000,000 images in 5.8 Seconds

50 Thousand Images

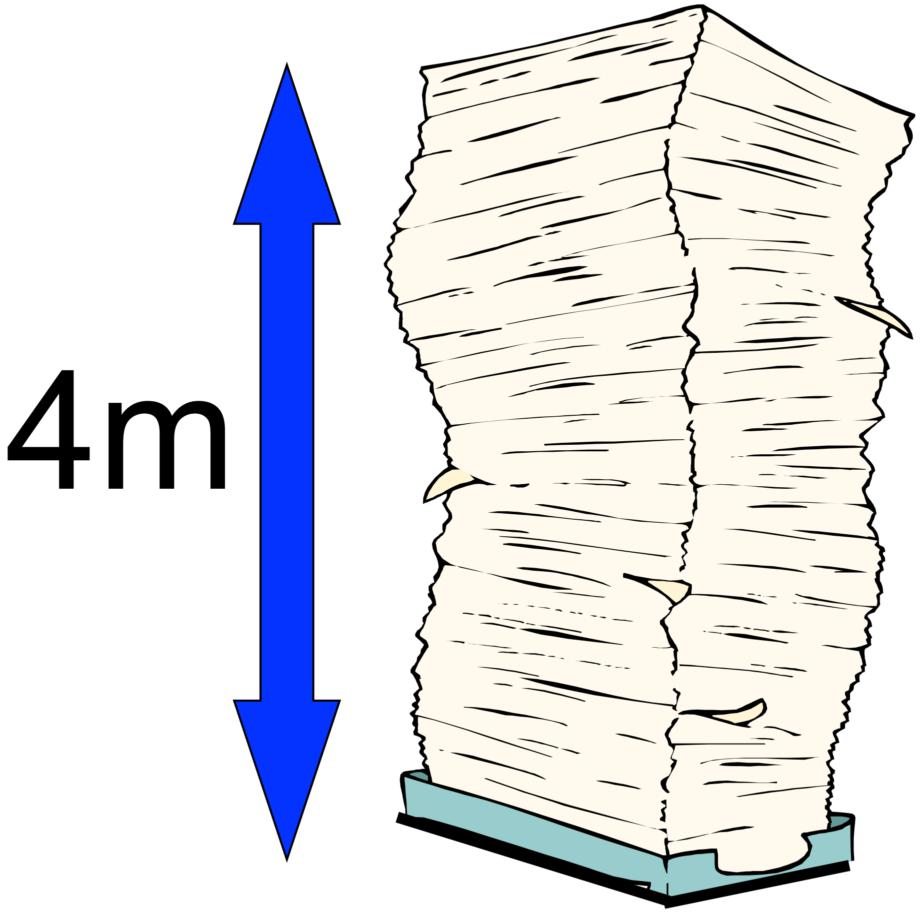

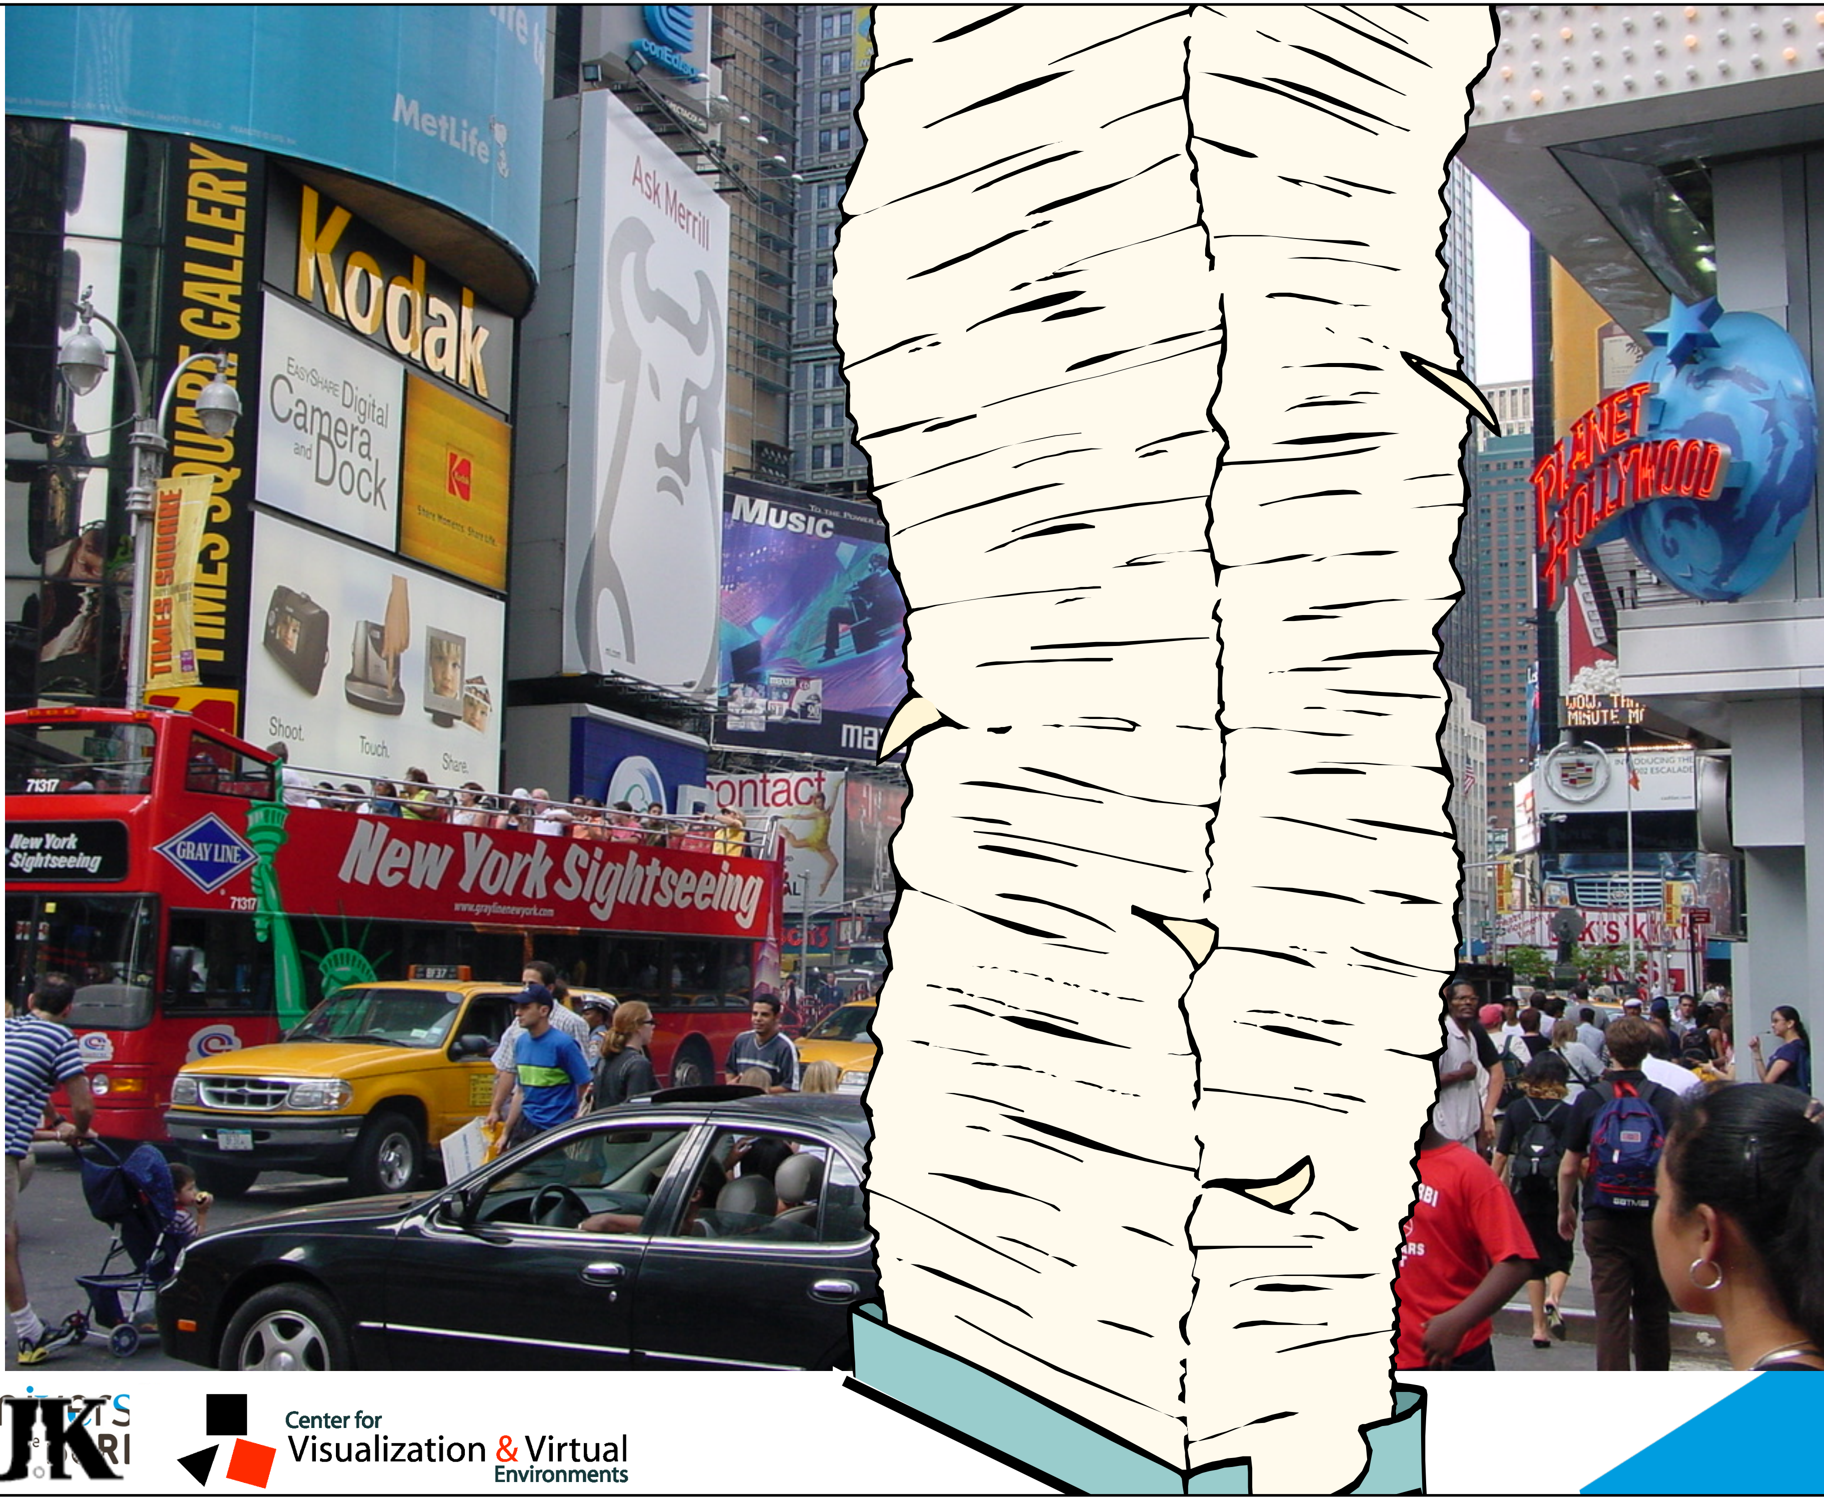
MetLife
TIMES SQUARE GALLERY
Kodak
EasyShare Digital Camera and Dock
Shoot.
Touch.
Share.
Ask Merrill
Music
New York Sightseeing
GRAY LINE
www.graylinenewyork.com
Planet Hollywood

Center for Visualization & Virtual Environments

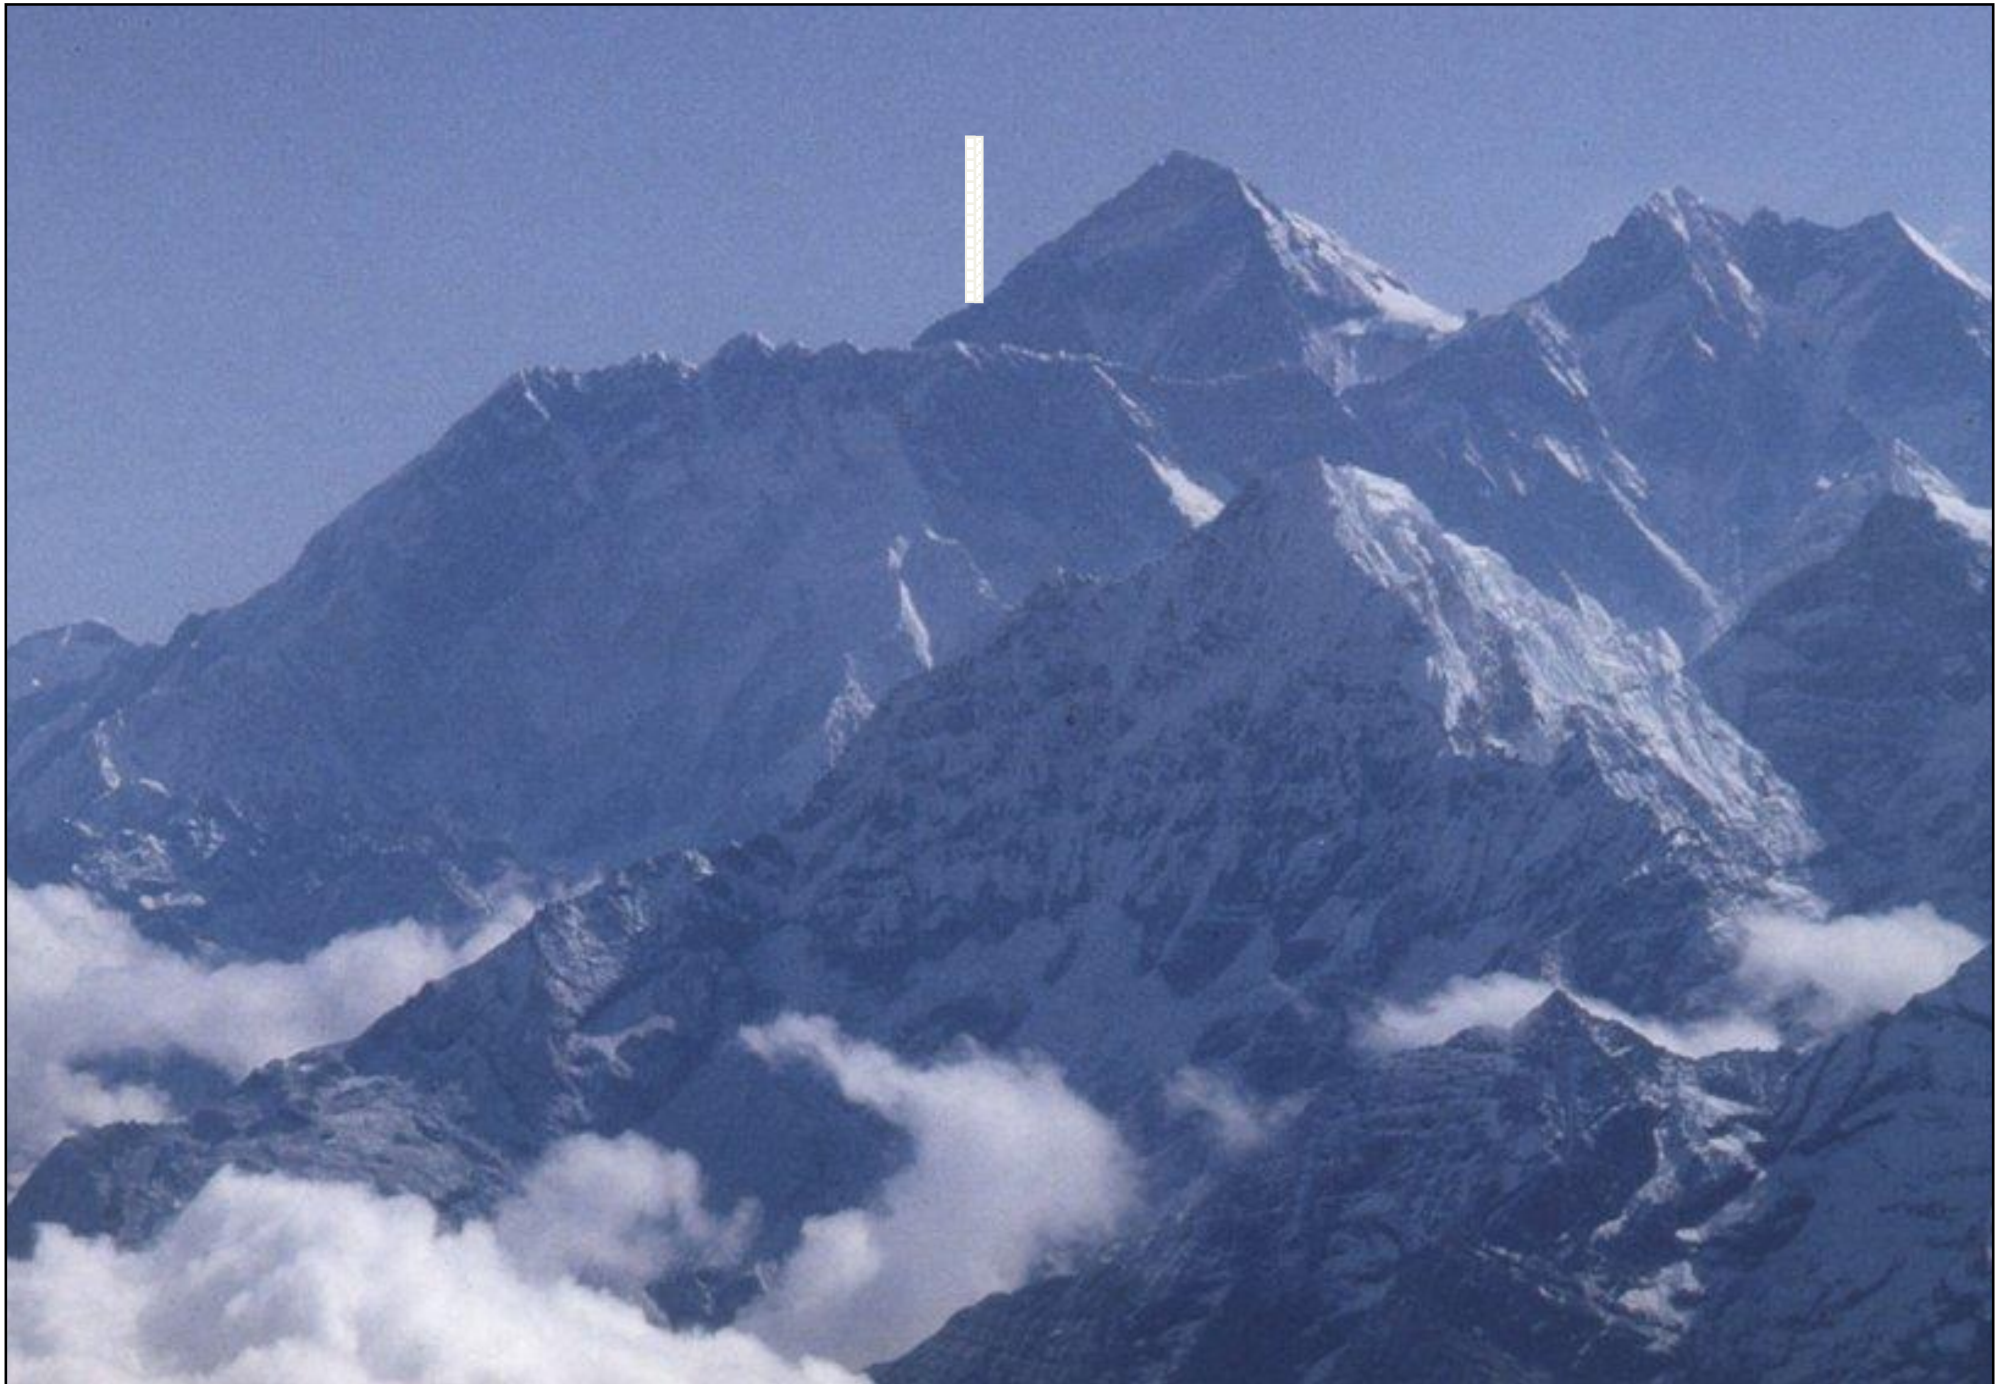

Start with some results!

What is the problem you are aiming to solve?

Why is it interesting to solve it?

Why is it difficult to solve it?

about 20% of the talk { Motivation

If your results are convincing, the audience will listen to your talk

Contribution, results, comparison

Related Work?

Can be useful to show that:

- the problem is important if there are many related papers;
- the problem is new if there is very few references.

- But don't spend too much time on the related work, your talk is about your work.

université de BORDEAUX

Feature Point Pipeline

input image

feature point detection

orientation estimation

descriptor computation

- Harris, C., Stephens, M., "A Combined Corner and Edge Detector,"AVC, 1988
- Mikolajczyk, K., Schmid, C., "Scale and Affine Invariant Interest Point Detectors," IJCV, 2004
- Förstner, W., Dickscheid, T., Schindler, F., "Detecting Interpretable and Accurate Scale-Invariant Keypoints," ICCV, 2009
- Rosten, E., Porter, R., Drummond, T., "Faster and Better: A Machine Learning Approach to Corner Detection," TPAMI, 2010
- Zitnick, C., Ramnath, K., "Edge Foci Interest Points", ICCV, 2011
- Mainali, P., Lafruit, G., Tack, K., Van Gool, L., Lauwereins, R., "Derivative-Based Scale Invariant Image Feature Detector with Error Resilience", TPAMI, 2014
- ...

input image

- Winder, S., Brown, M., "Learning Local Image Descriptors," CVPR, 2007
- Tola, E., Lepetit, V., Fua, P., "A Fast Local Descriptor for Dense Matching," CVPR, 2008
- Fan, B., Wu, F., Hu., Z., "Aggregating Gradient Distributions into Intensity Orders: A Novel Local Image Descriptor," CVPR, 2011
- Alahi, A., Ortiz, R., Vandergheynst, P., "FREAK: Fast Retina Keypoint," CVPR, 2012
- Simonyan, K., Vedaldi, A., Zisserman, A., "Learning Local Feature Descriptors Using Convex Optimisation," TPAMI, 2014
- Zagoruyko, S., Komodakis, N., "Learning to Compare Image Patches via Convolutional Neural Networks,", CVPR, 2015
- ...

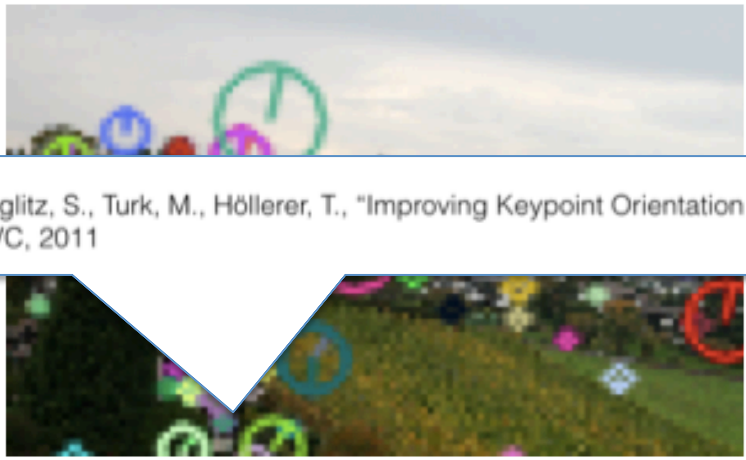

- Gauglitz, S., Turk, M., Höllerer, T., "Improving Keypoint Orientation Assignment,", BMVC, 2011

orientation estimation

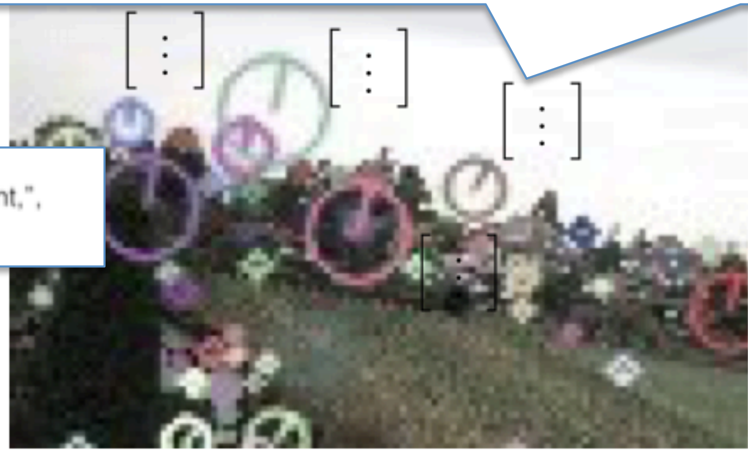

descriptor computation

Outline

- Introduction;
- Related work;
- *some concepts related to your work*
- *your work*
- Results and comparison;
- Conclusion and future work.

université de BORDEAUX

No Outline!

- Introduction;
- Related work;
- *some concepts related to your work*
- *your work*
- Results and comparison;
- Conclusion and future work

université de BORDEAUX

No Outline!

Possibly a list of the topics of your talk if it is in several parts;

Have transition slides between the different parts.

université de BORDEAUX

- 3D hand pose estimation and tracking.

- Feature point detection and description.

- 3D hand pose estimation and tracking.

Training a Feedback Loop for Hand Pose Estimation.
Markus Oberweger, Paul Wohlhart, and Vincent Lepetit.
ICCV'15. Oral

- Feature point detection and description.

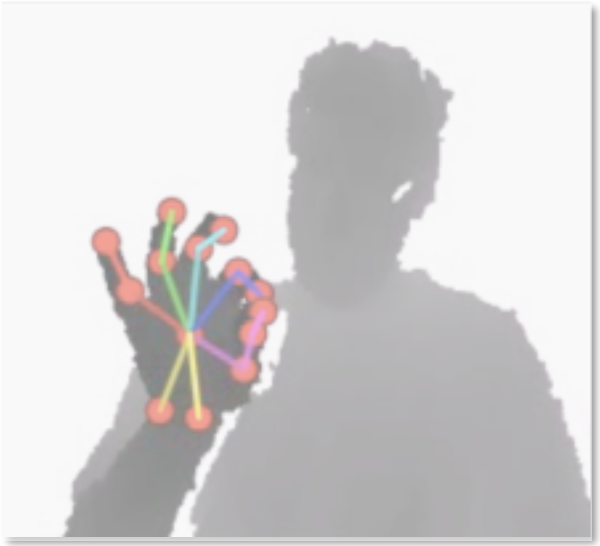

• 3D hand pose estimation and tracking.

• Feature point detection and description.

LIFT: Learned Invariant Feature Transform.
Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua
ECCV'16. Spotlight

Give the Intuition of your Contribution

université de BORDEAUX

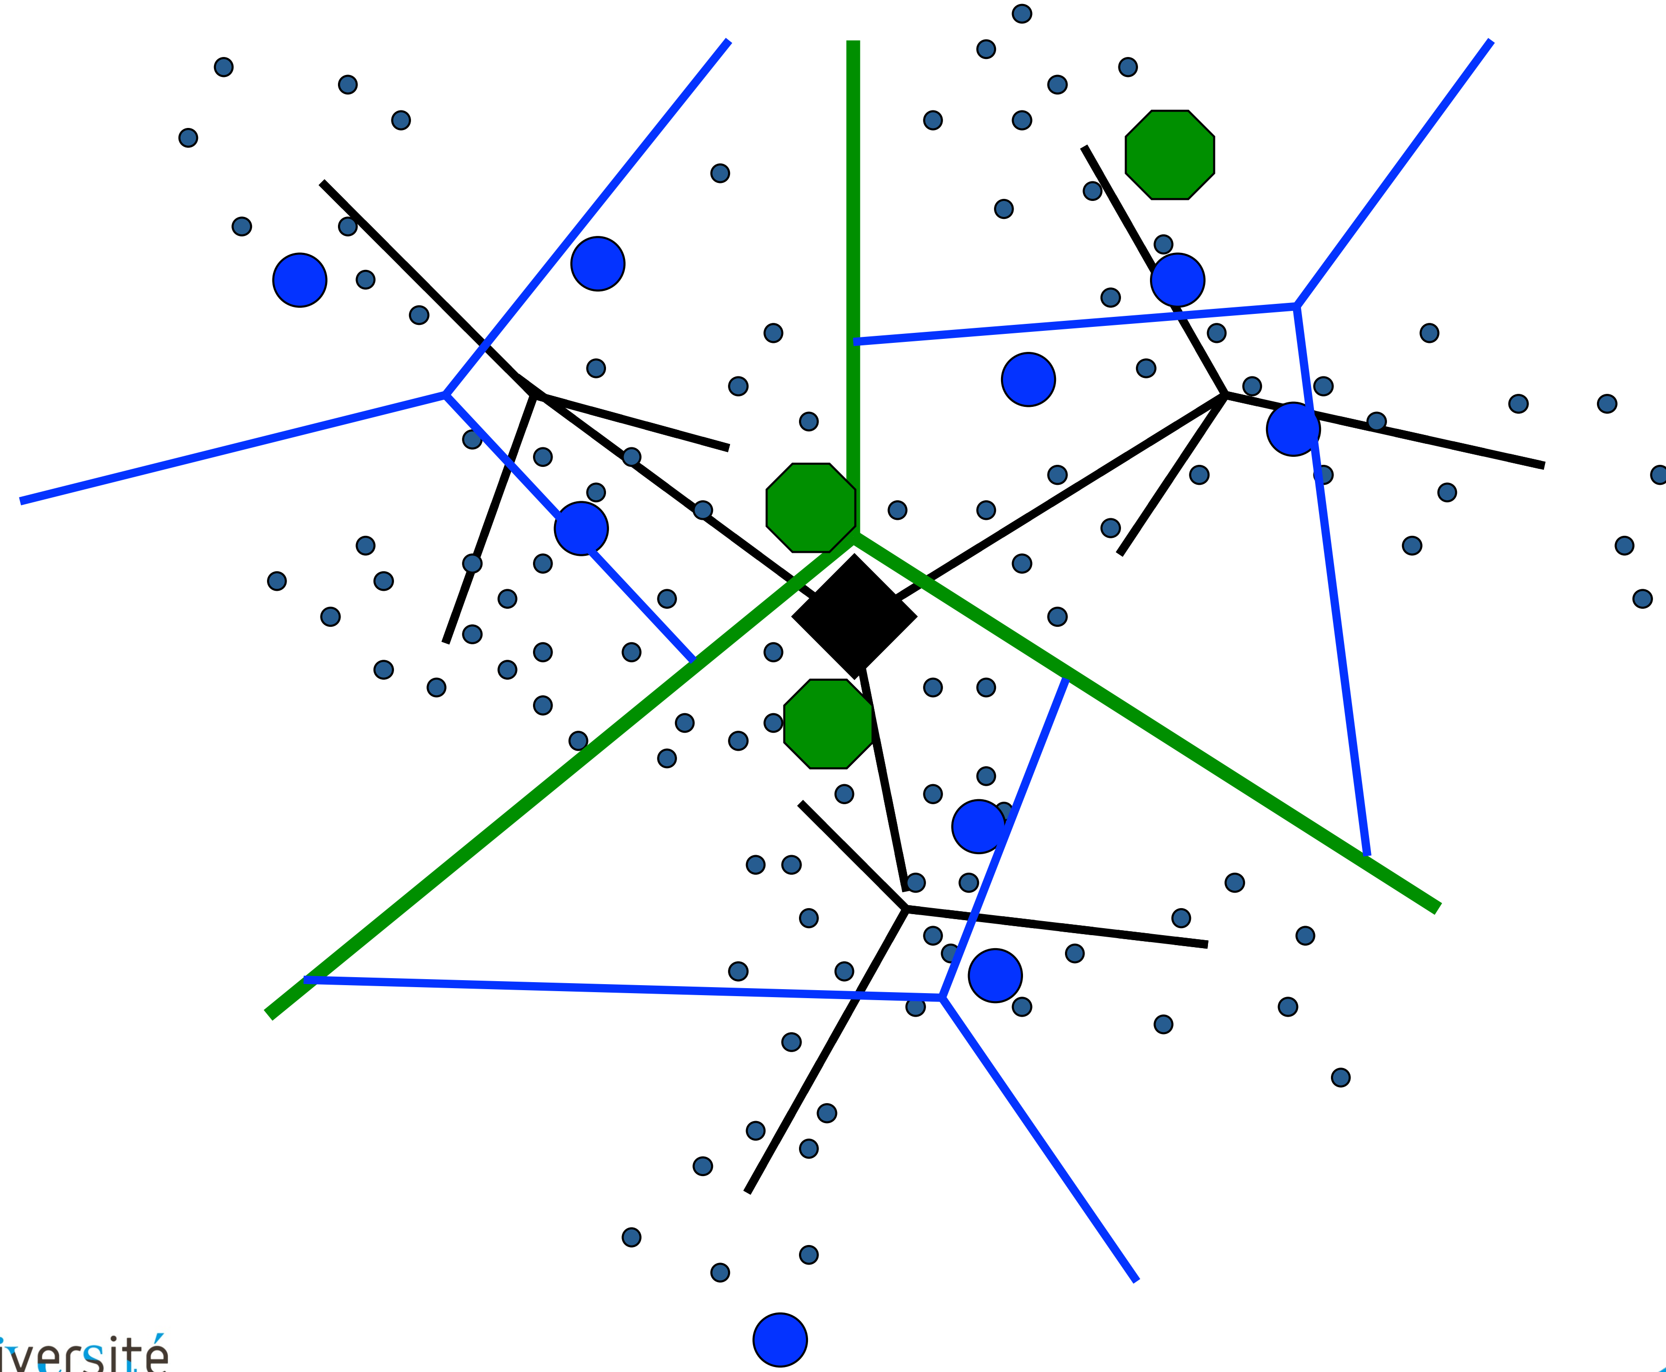

université de BORDEAUX

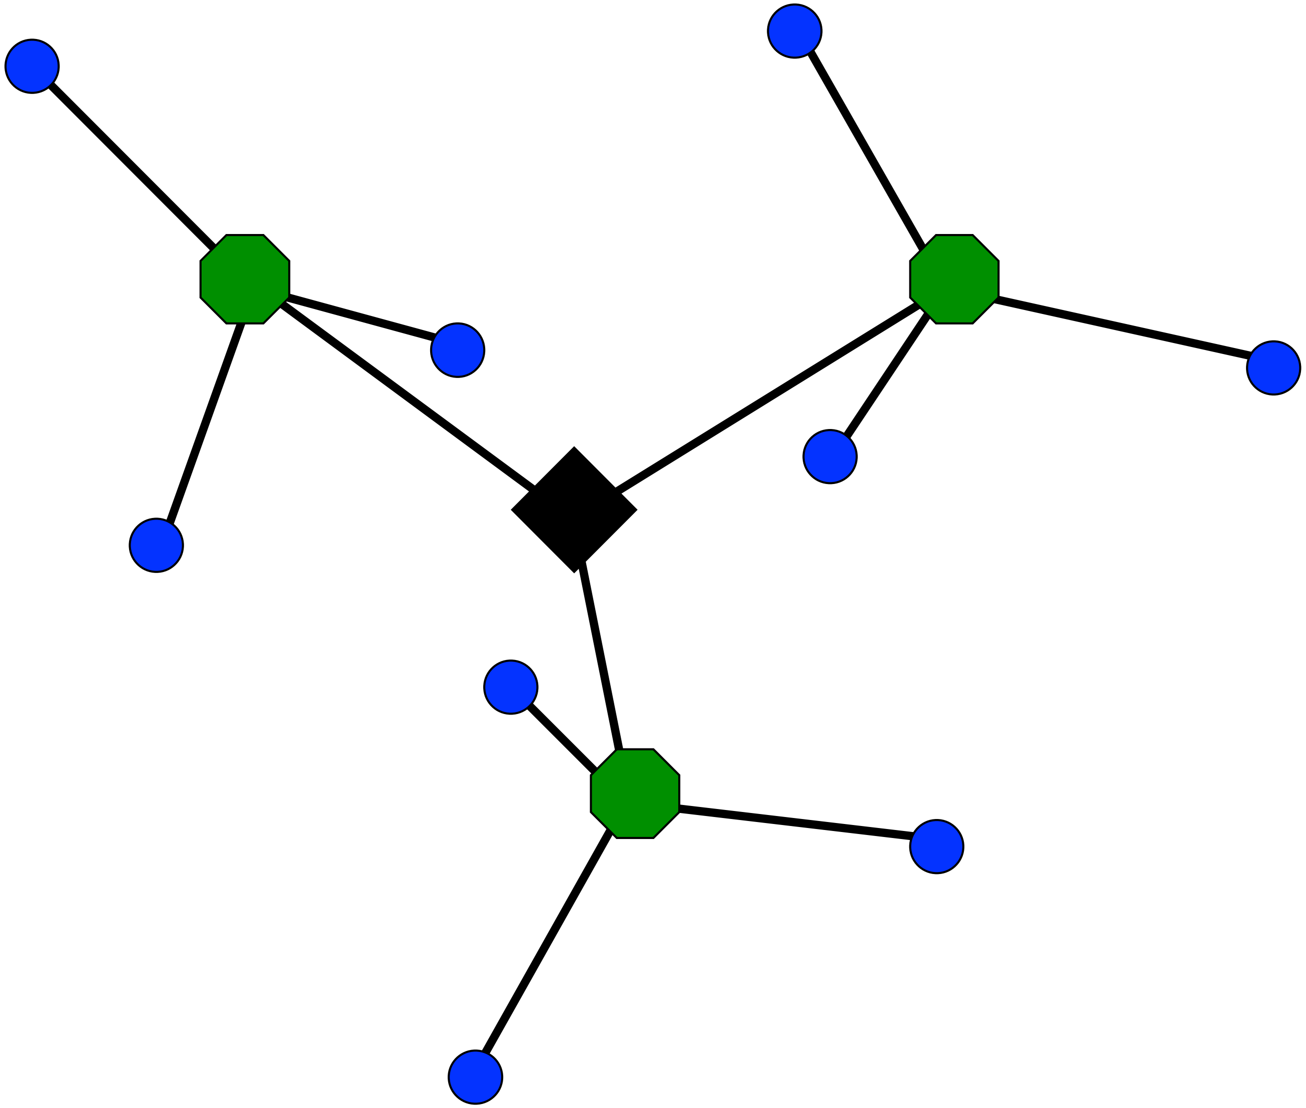

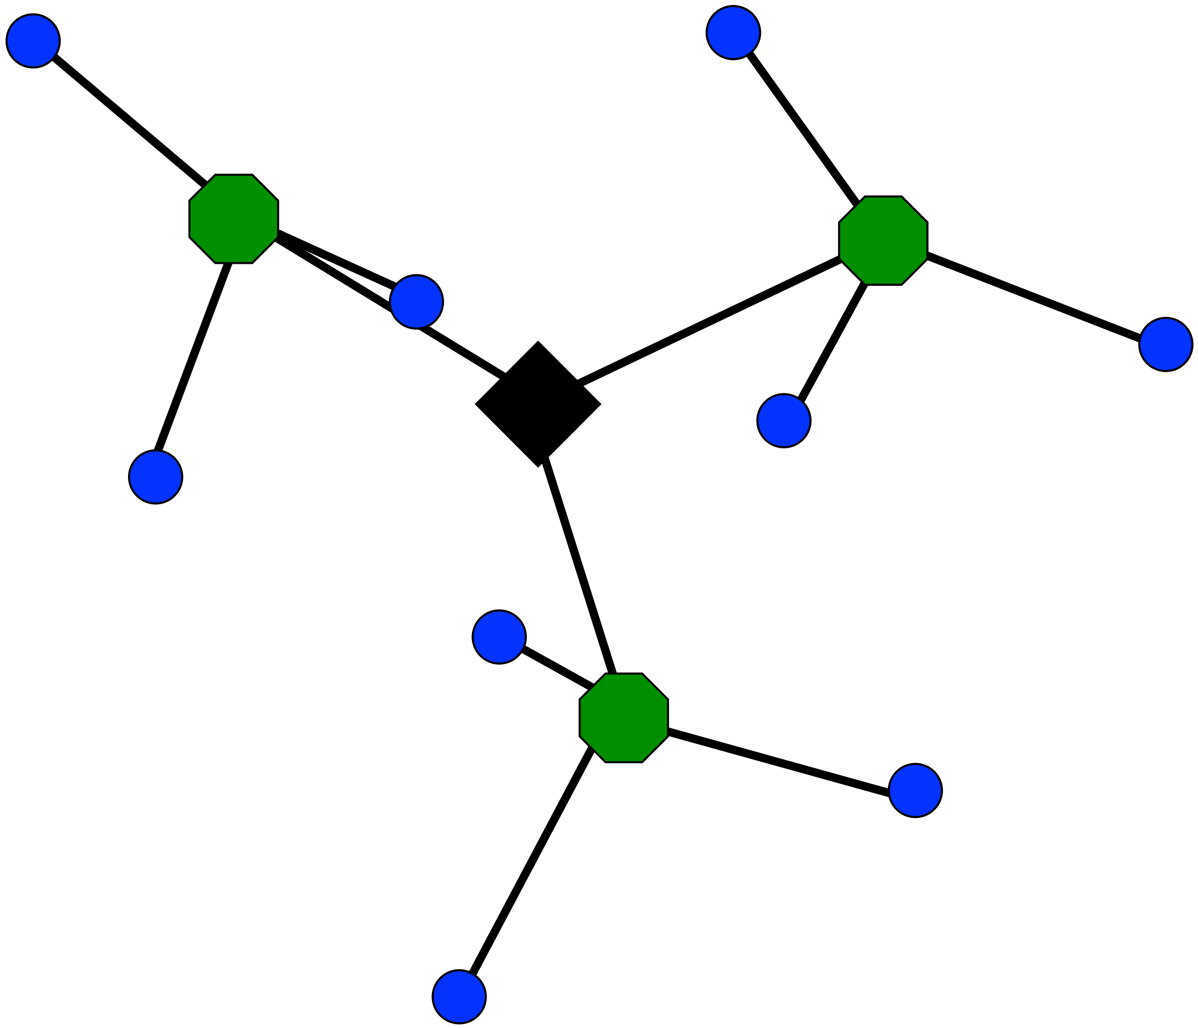

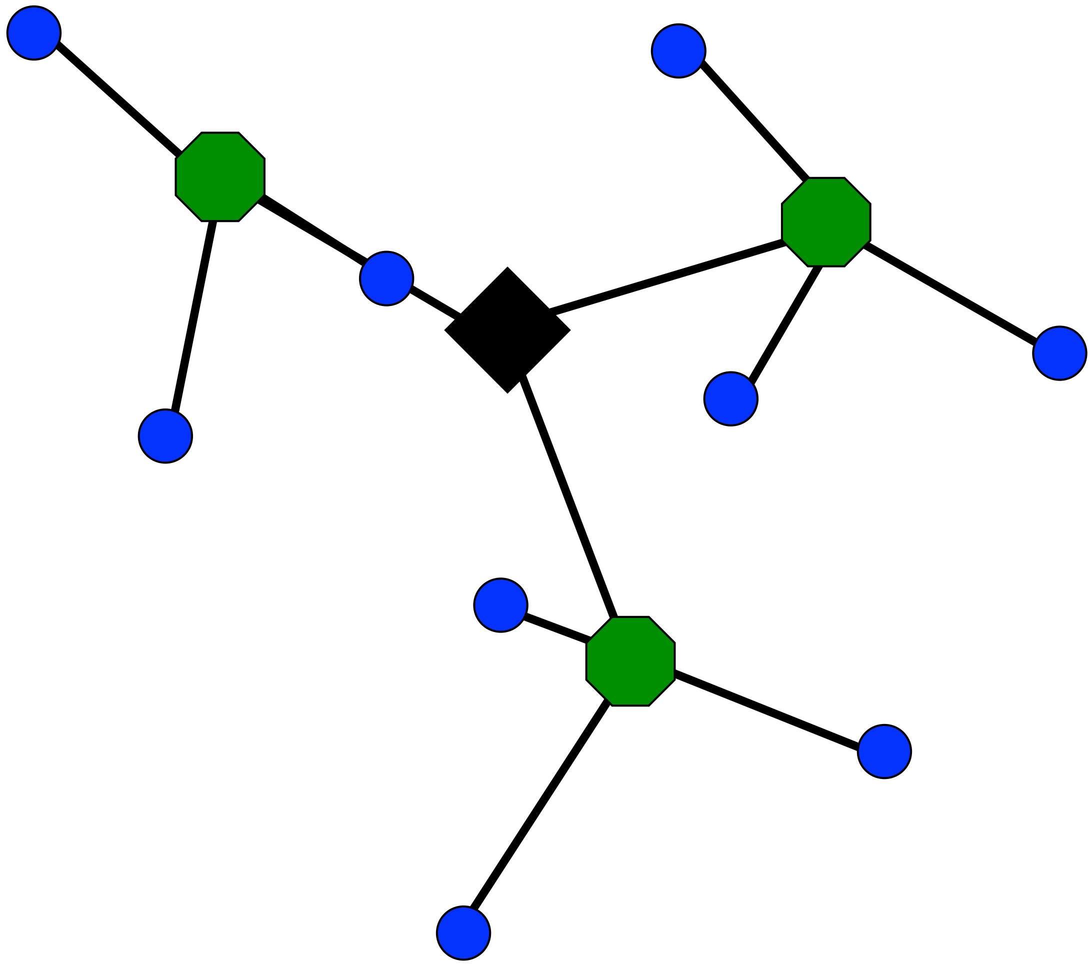

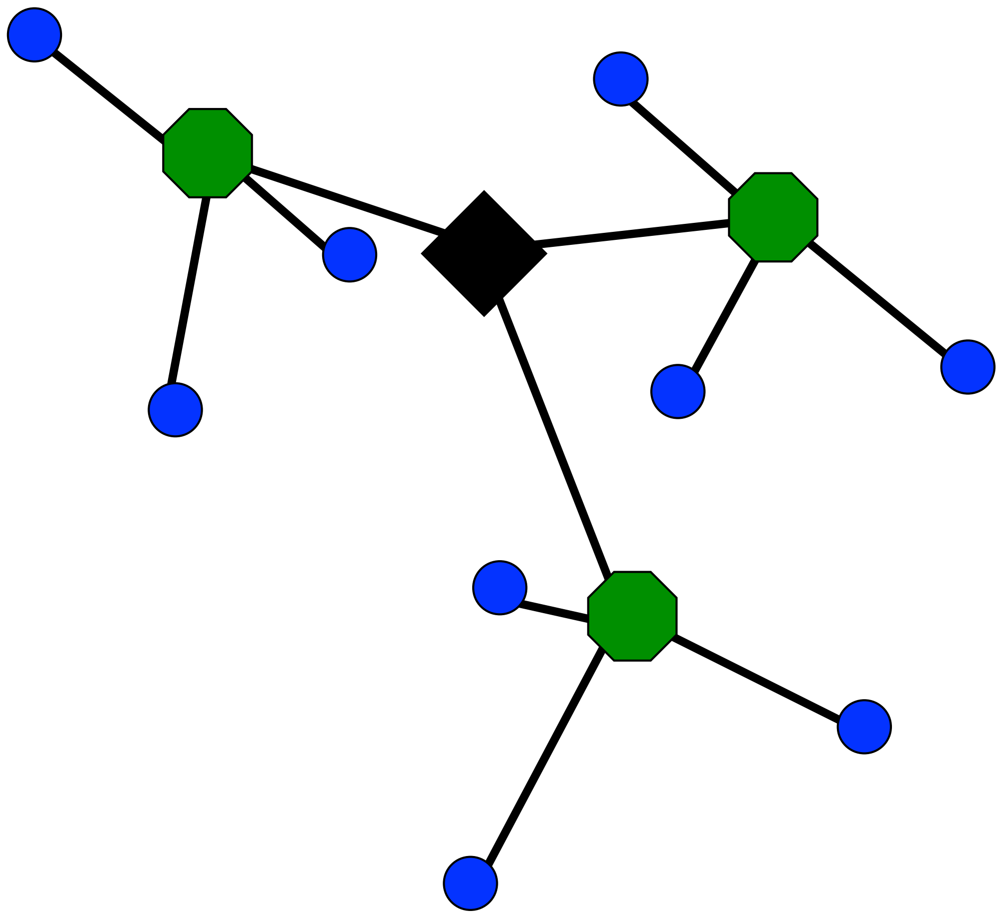

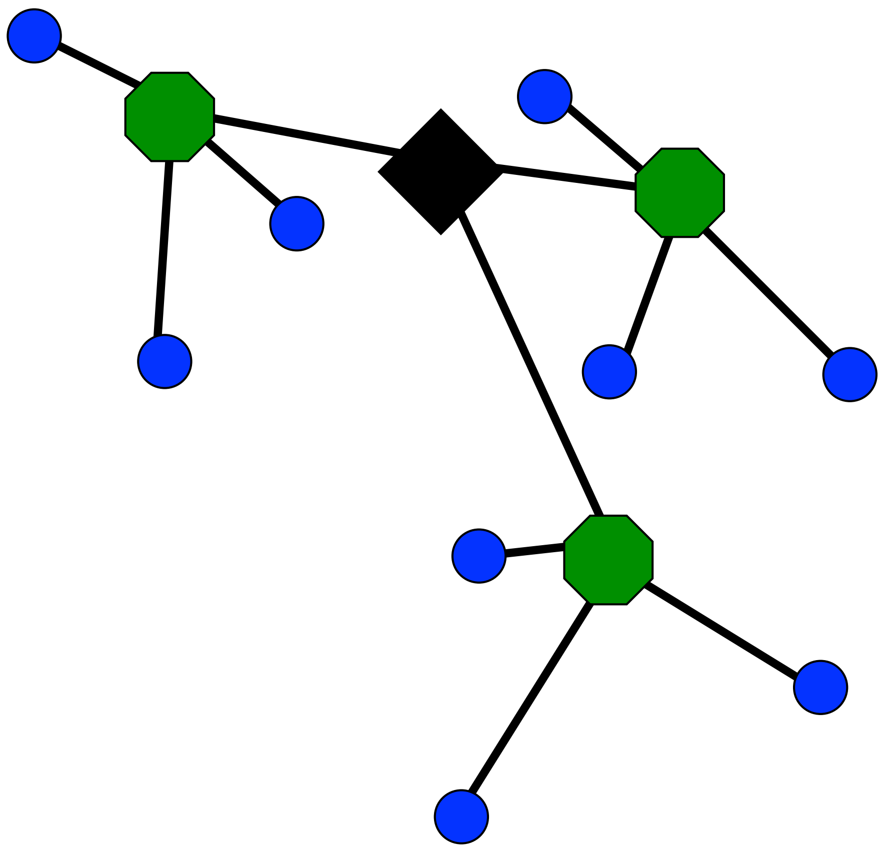

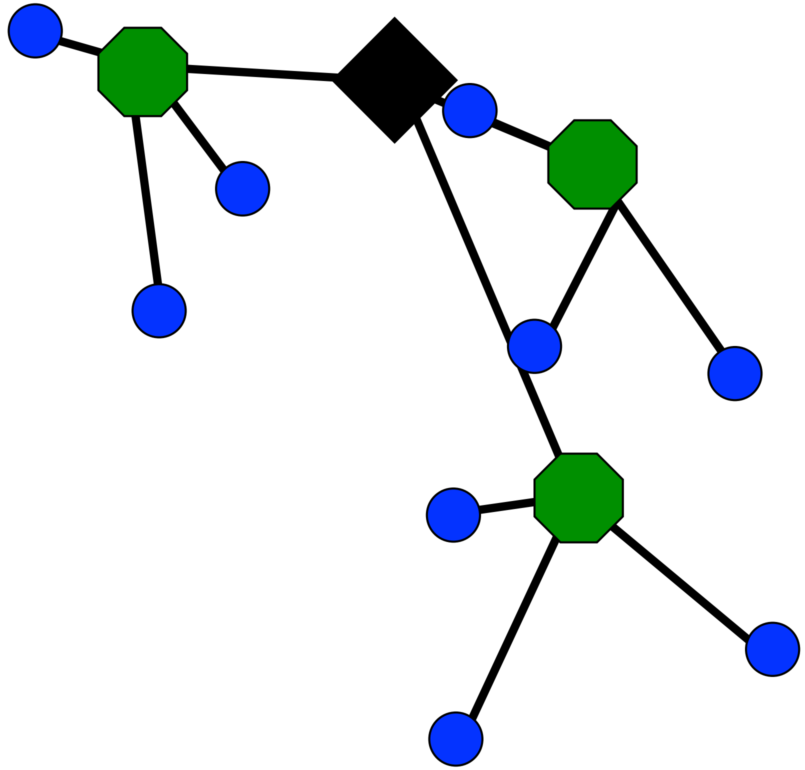

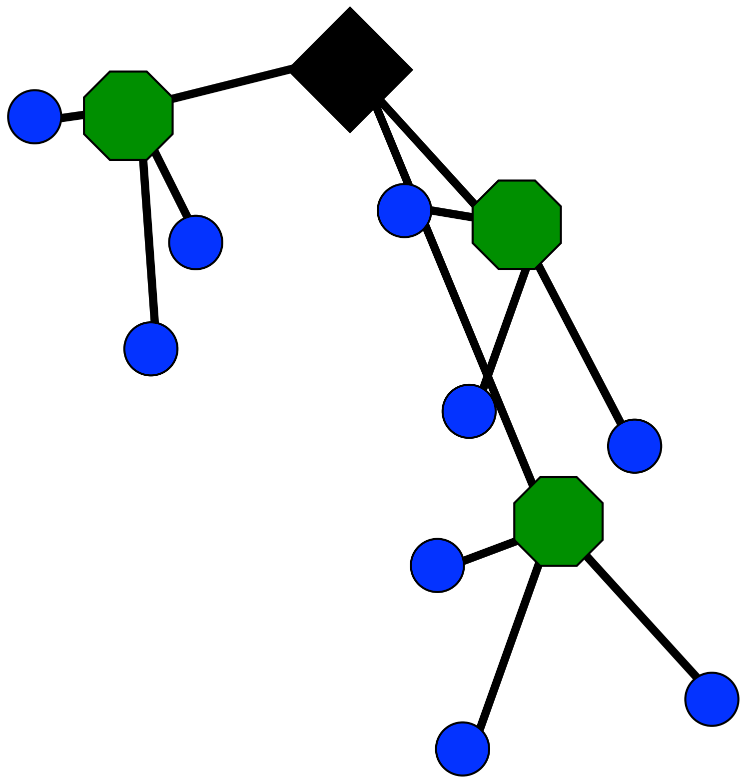

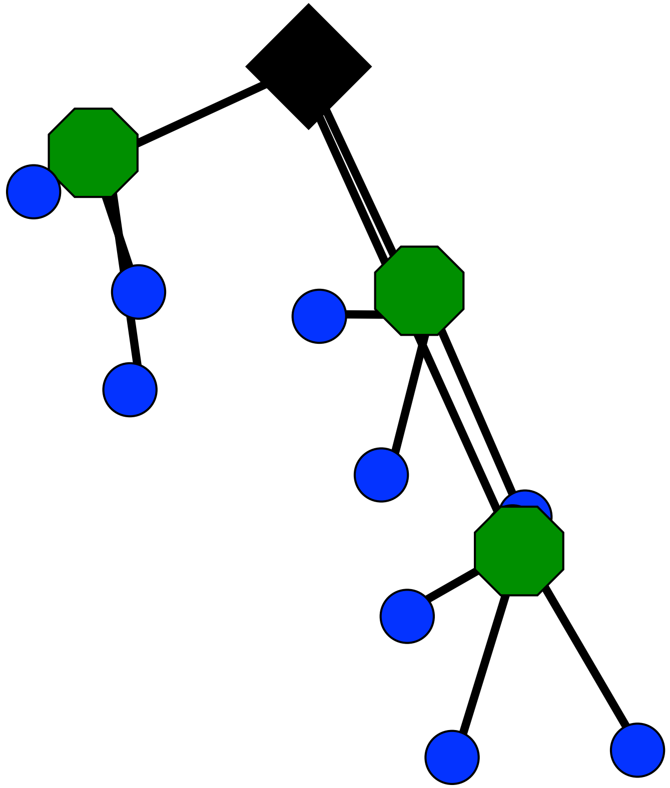

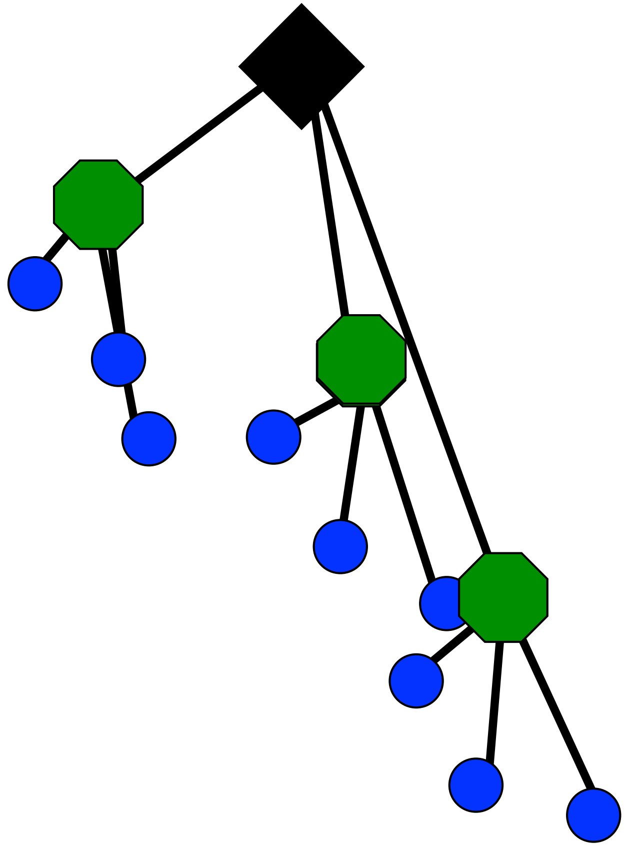

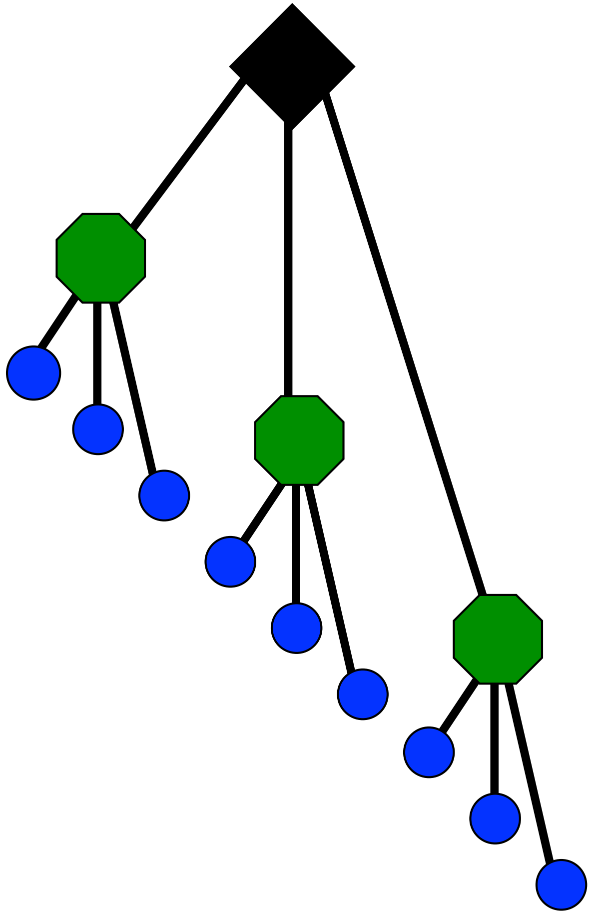

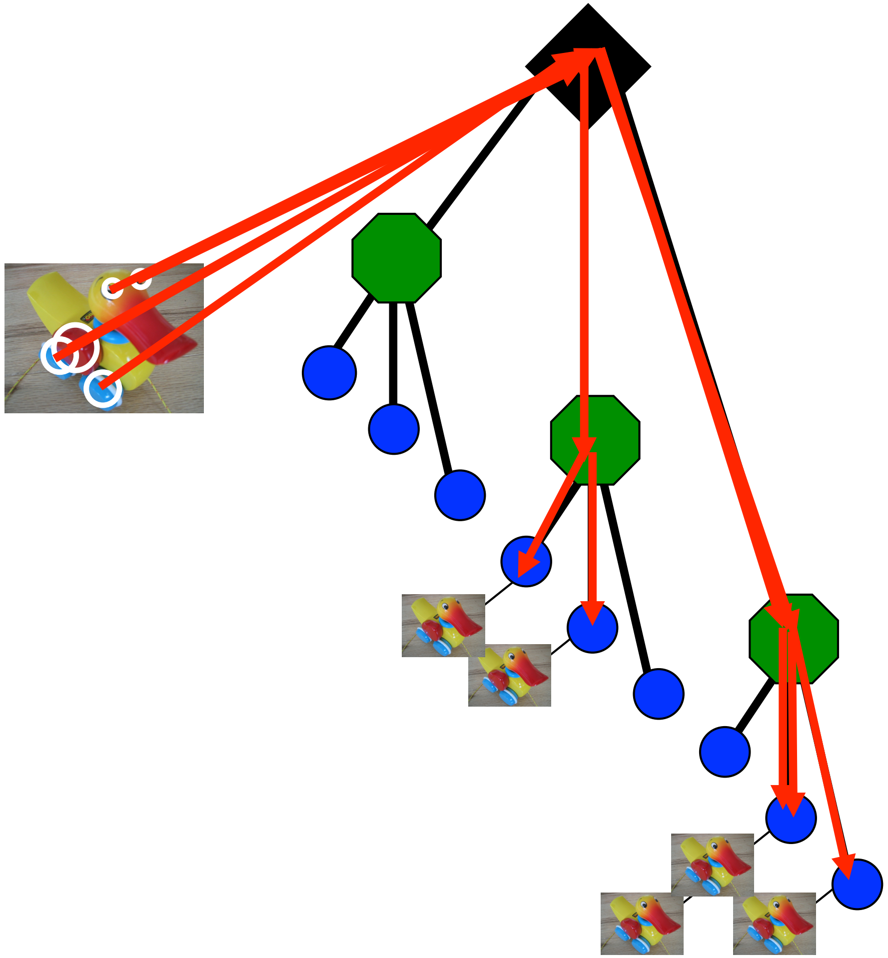

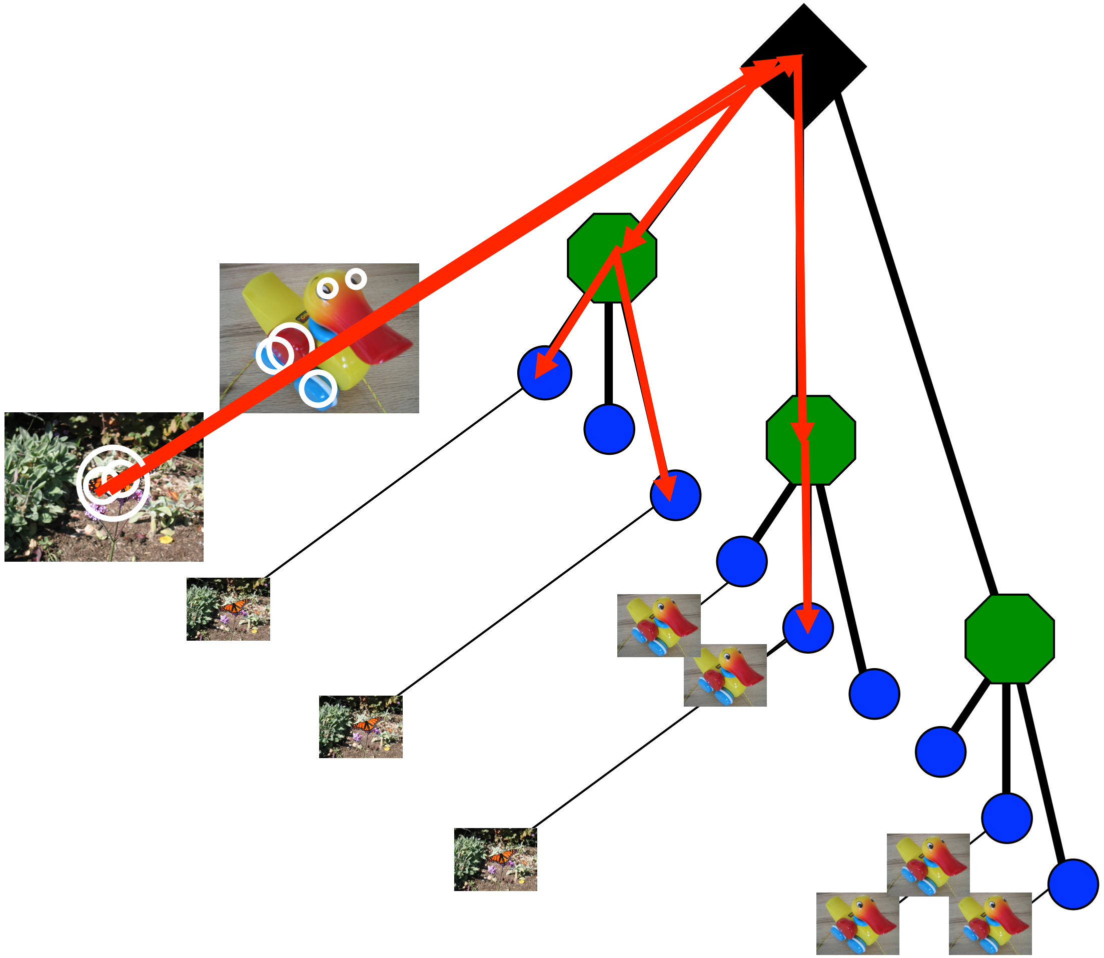

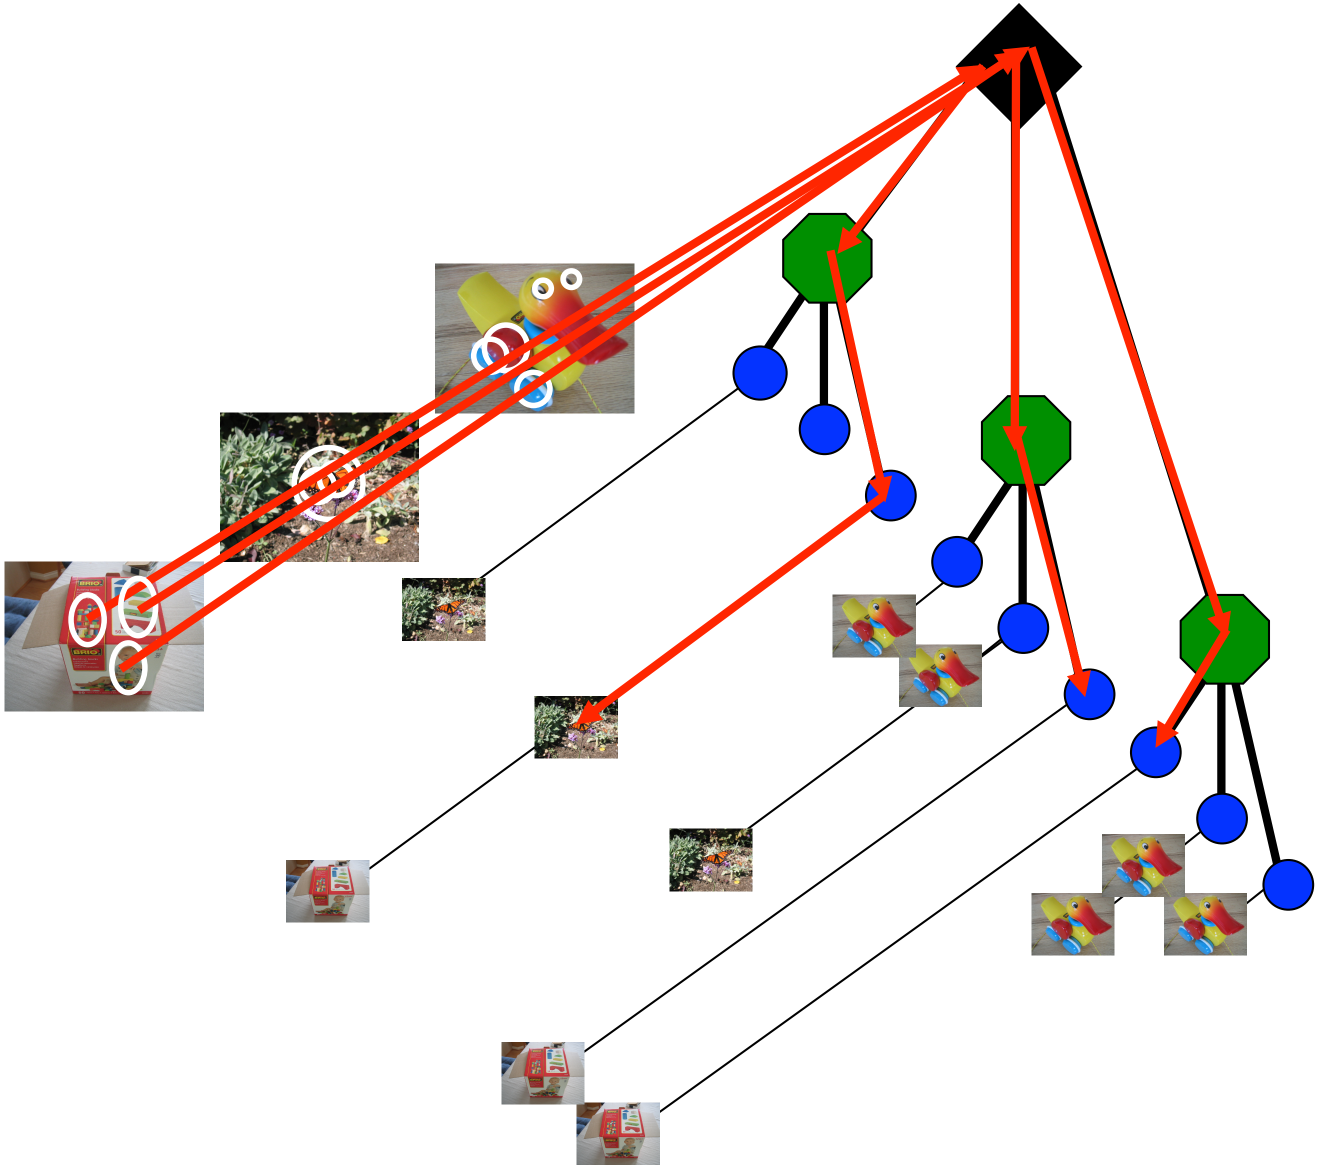

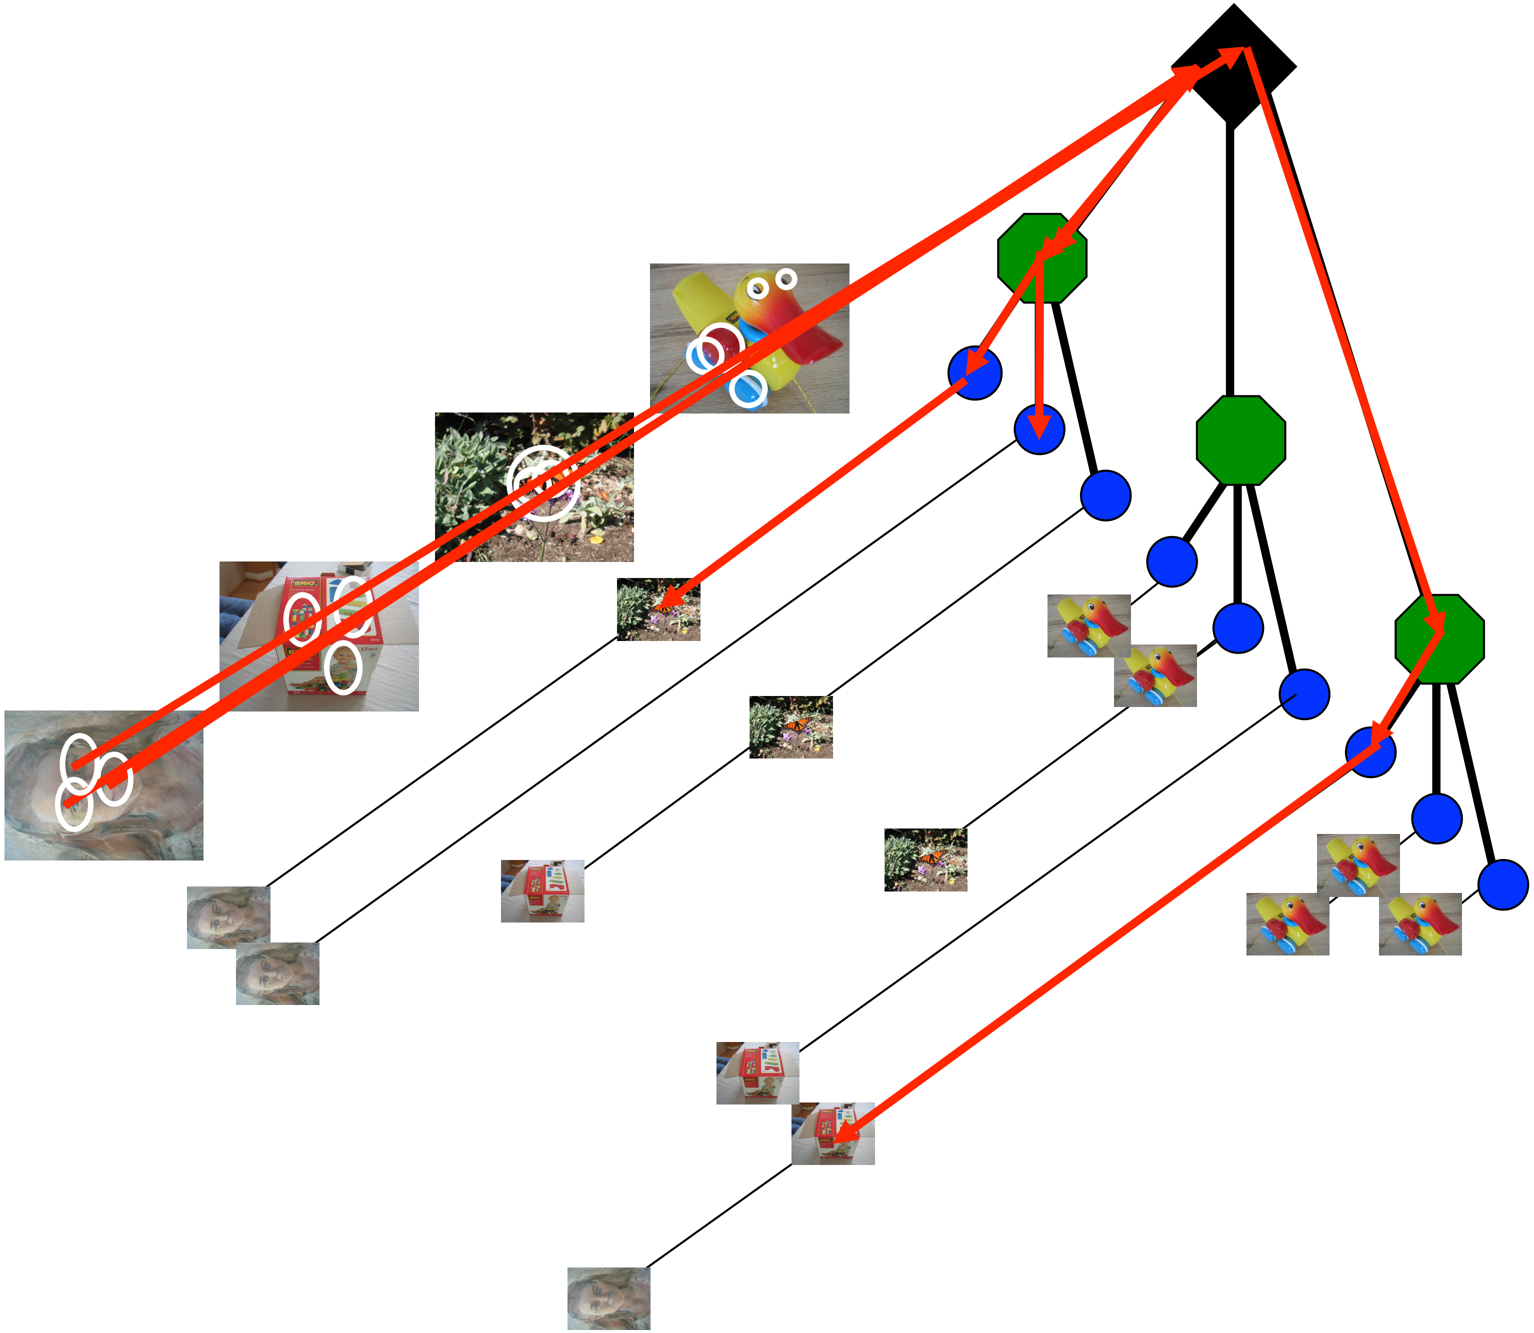

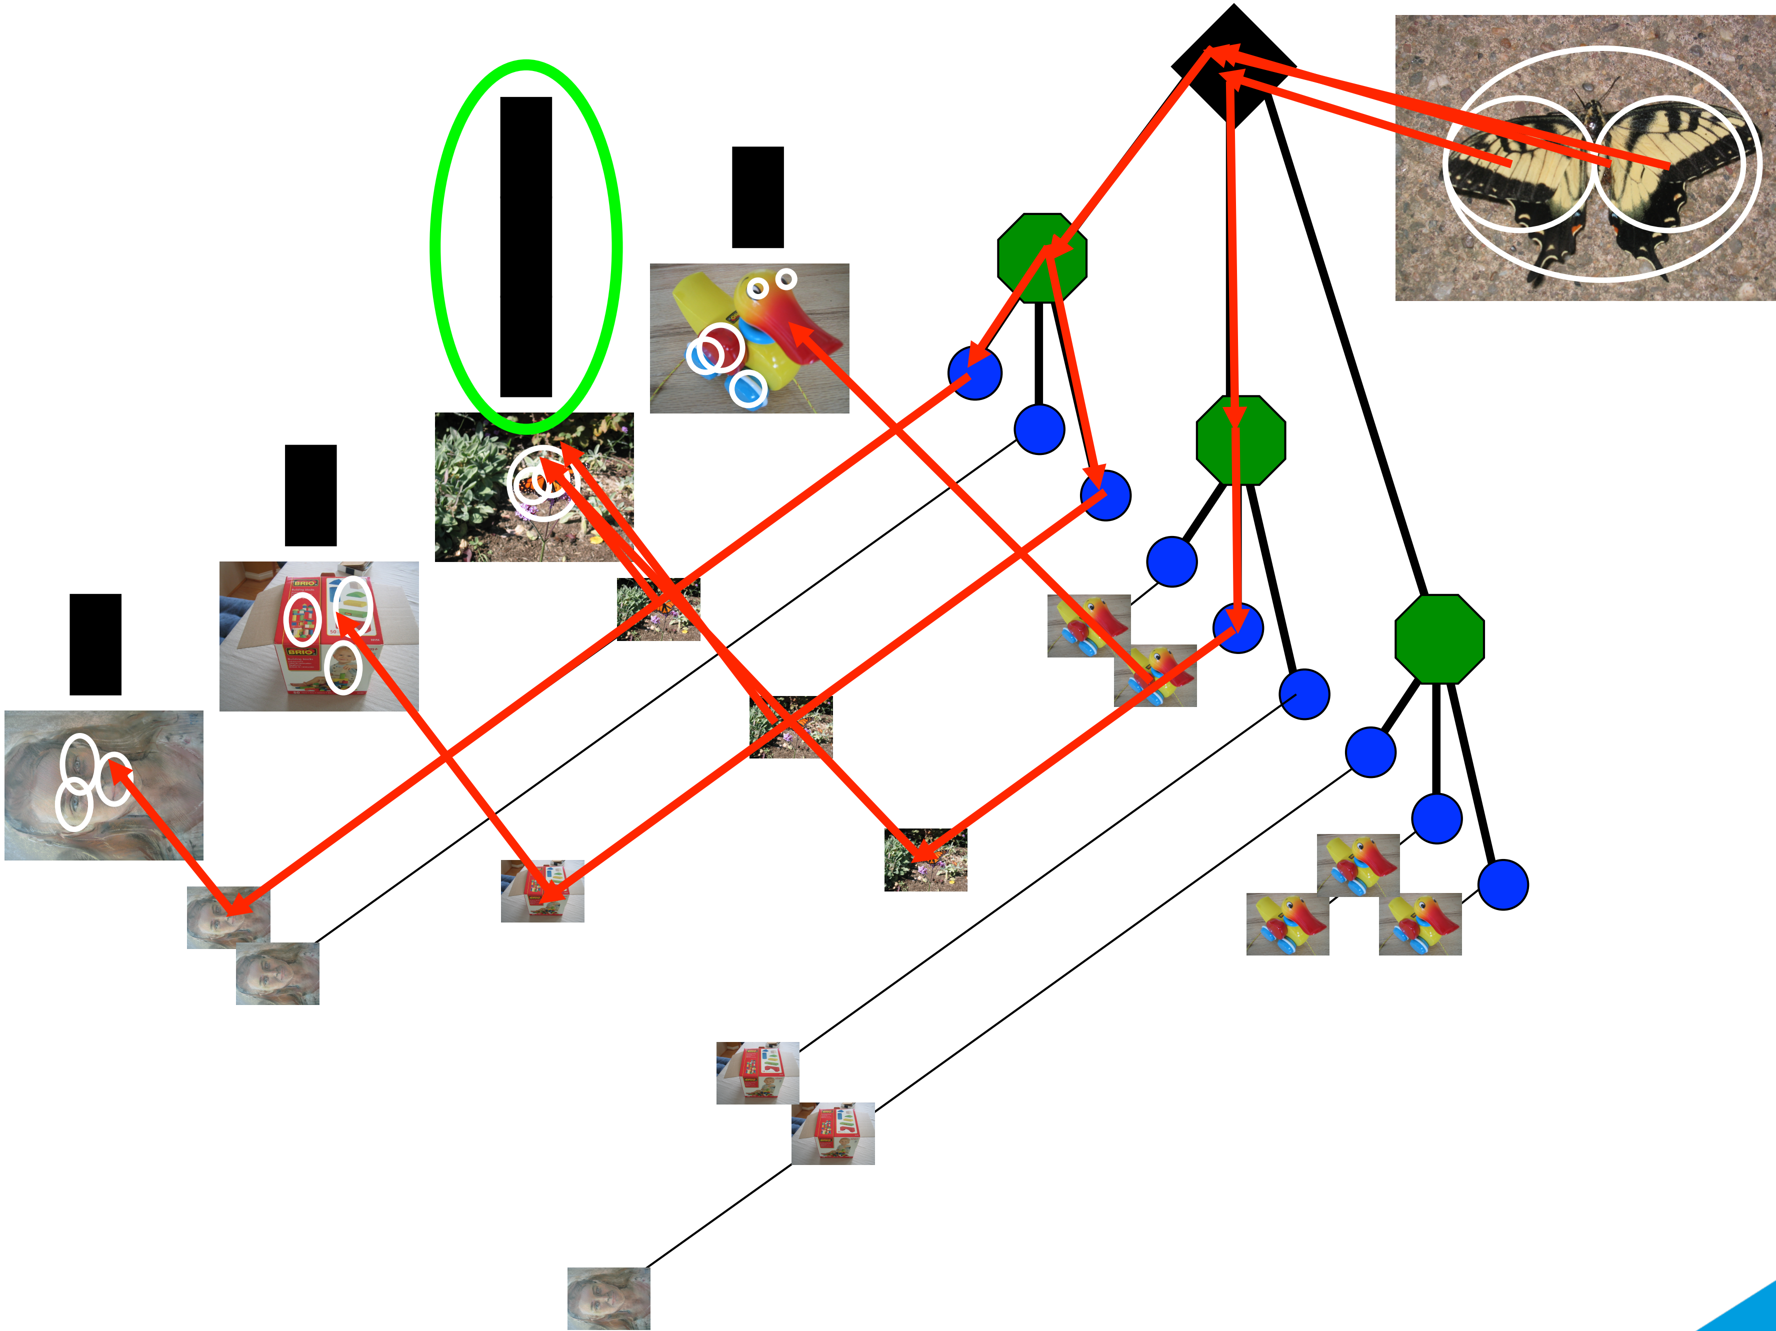

Timing

One idea / item per slide;

Rule of thumb: about 1 minute per slide.

Formulas

Be very careful!

$$\mathcal{L}_{D_0} = \mathbb{E}_{(I_0,t)\sim p_{data}}[\log D_0(I_0, \varphi_t)] + \mathbb{E}_{z\sim p_z, t\sim p_{data}}[\log(1 - D_0(G_0(z, \hat{c}_0), \varphi_t))],$$

$$\mathcal{L}_{G_0} = \mathbb{E}_{z\sim p_z, t\sim p_{data}}[\log(1 - D_0(G_0(z, \hat{c}_0), \varphi_t))] + \lambda D_{KL}(\mathcal{N}(\mu_0(\varphi_t), \Sigma_0(\varphi_t)) \,||\, \mathcal{N}(0, I))$$

université de BORDEAUX

Multilayer Perceptron: Optimization

$$\mathbf{h} = g(\mathbf{W}\mathbf{x} + \mathbf{b})$$
$$\mathbf{y}(\mathbf{x}) = \mathbf{W}_2\mathbf{h} + \mathbf{b}_2$$

How can we find $\mathbf{W}$, $\mathbf{b}$, $\mathbf{W}_2$, and $\mathbf{b}_2$?

[Rumelhart, Hinton, Williams] introduces an objective (or loss) function:

$$\mathcal{L}(\mathcal{T}) = \sum_{(\mathbf{x},\mathbf{d})\in\mathcal{T}} \|\mathbf{y}(\mathbf{x}) - \mathbf{d}\|^2$$

Multilayer Perceptron: Optimization

$$\mathbf{h} = g(\mathbf{W}\mathbf{x} + \mathbf{b})$$
$$\mathbf{y}(\mathbf{x}) = \mathbf{W}_2\mathbf{h} + \mathbf{b}_2$$

How can we find $\mathbf{W}$, $\mathbf{b}$, $\mathbf{W}_2$, and $\mathbf{b}_2$?

[Rumelhart, Hinton, Williams] introduces an objective (or loss) function:

$$\mathcal{L}(\mathcal{T}) = \sum_{(\mathbf{x},\mathbf{d})\in\mathcal{T}} \|\mathbf{y}(\mathbf{x}) - \mathbf{d}\|^2$$

$\mathcal{T}$: training set

Multilayer Perceptron: Optimization

$$\mathbf{h} = g(\mathbf{W}\mathbf{x} + \mathbf{b})$$
$$\mathbf{y}(\mathbf{x}) = \mathbf{W}_2\mathbf{h} + \mathbf{b}_2$$

How can we find $\mathbf{W}$, $\mathbf{b}$, $\mathbf{W}_2$, and $\mathbf{b}_2$?

[Rumelhart, Hinton, Williams] introduces an objective (or loss) function:

$$\mathcal{L}(\mathcal{T}) = \sum_{(\mathbf{x},\mathbf{d})\in\mathcal{T}} \|\mathbf{y}(\mathbf{x}) - \mathbf{d}\|^2$$

$\mathcal{T}$: training set

$\mathbf{x}$: 1 training example;
$\mathbf{d}$: the desired output for $\mathbf{x}$.

université de BORDEAUX

Multilayer Perceptron: Optimization

$$\mathbf{h} = g(\mathbf{W}\mathbf{x} + \mathbf{b})$$
$$\mathbf{y}(\mathbf{x}) = \mathbf{W}_2\mathbf{h} + \mathbf{b}_2$$

How can we find $\mathbf{W}$, $\mathbf{b}$, $\mathbf{W}_2$, and $\mathbf{b}_2$?

[Rumelhart, Hinton, Williams] introduces an objective (or loss) function:

$$\mathcal{L}(\mathcal{T}) = \sum_{(\mathbf{x},\mathbf{d})\in\mathcal{T}} \|\mathbf{y}(\mathbf{x}) - \mathbf{d}\|^2$$

$\mathbf{y}(\mathbf{x})$: network output for $\mathbf{x}$

$\mathcal{T}$: training set

$\mathbf{x}$: 1 training example;
$\mathbf{d}$: the desired output for $\mathbf{x}$.

université de BORDEAUX

Formulas

If you use the same terms over several slides, make sure the audience did not forget them.

Descriptor Learning

Training: from a training set $\{((\mathbf{x}_n, \mathbf{y}_n), l_n))\}$, find $k_d(.)$ functions and weights α_d that minimize:

one sample:
a *pair* of image patches

$$\sum_{n=1}^{N} \exp\left(-l_n \sum_{d=1}^{L} \alpha_d c_d(\overbrace{\mathbf{x}_n, \mathbf{y}_n})\right)$$

label for the sample
(+1 or -1)

$$c_d(\mathbf{x}_n, \mathbf{y}_n) = k_d(\mathbf{x}_n) k_d(\mathbf{y}_n)$$

$\mathbf{x}_i =$  , $\mathbf{y}_i =$  , $l_i = +1$

$\mathbf{x}_j =$  , $\mathbf{y}_j =$ 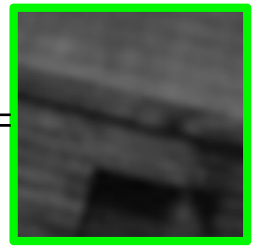 , $l_j =$ -1

université de BORDEAUX

Descriptor Learning

Training: find $k_d(.)$ functions and weights α_d that minimize:

$$\sum_{n=1}^{N} \exp\left(-l_n \sum_{d=1}^{L} \alpha_d c_d(\mathbf{x}_n, \mathbf{y}_n)\right)$$

$$c_d(\mathbf{x}_n, \mathbf{y}_n) = k_d(\mathbf{x}_n) k_d(\mathbf{y}_n)$$

If you use the same terms over several slides, make sure the audience did not forget them.

université de BORDEAUX

Descriptor Learning

Training: find $k_d(.)$ functions and weights α_d that minimize:

$$\sum_{n=1}^{N} \exp\left(-l_n \sum_{d=1}^{L} \alpha_d c_d(\mathbf{x}_n, \mathbf{y}_n)\right)$$

$$c_d(\mathbf{x}_n, \mathbf{y}_n) = k_d(\mathbf{x}_n) k_d(\mathbf{y}_n)$$

At run-time: given two *new* image patches $\mathbf{x}$ and $\mathbf{y}$, first compute

$$\text{descriptor}(\mathbf{x}) = [\sqrt{\alpha_1} k_1(\mathbf{x}), \ldots \sqrt{\alpha_L} k_L(\mathbf{x})]^\top$$

$$\text{descriptor}(\mathbf{y}) = [\sqrt{\alpha_1} k_1(\mathbf{y}), \ldots \sqrt{\alpha_L} k_L(\mathbf{y})]^\top$$

You can use plain names for variables.

Graphs

Graphs

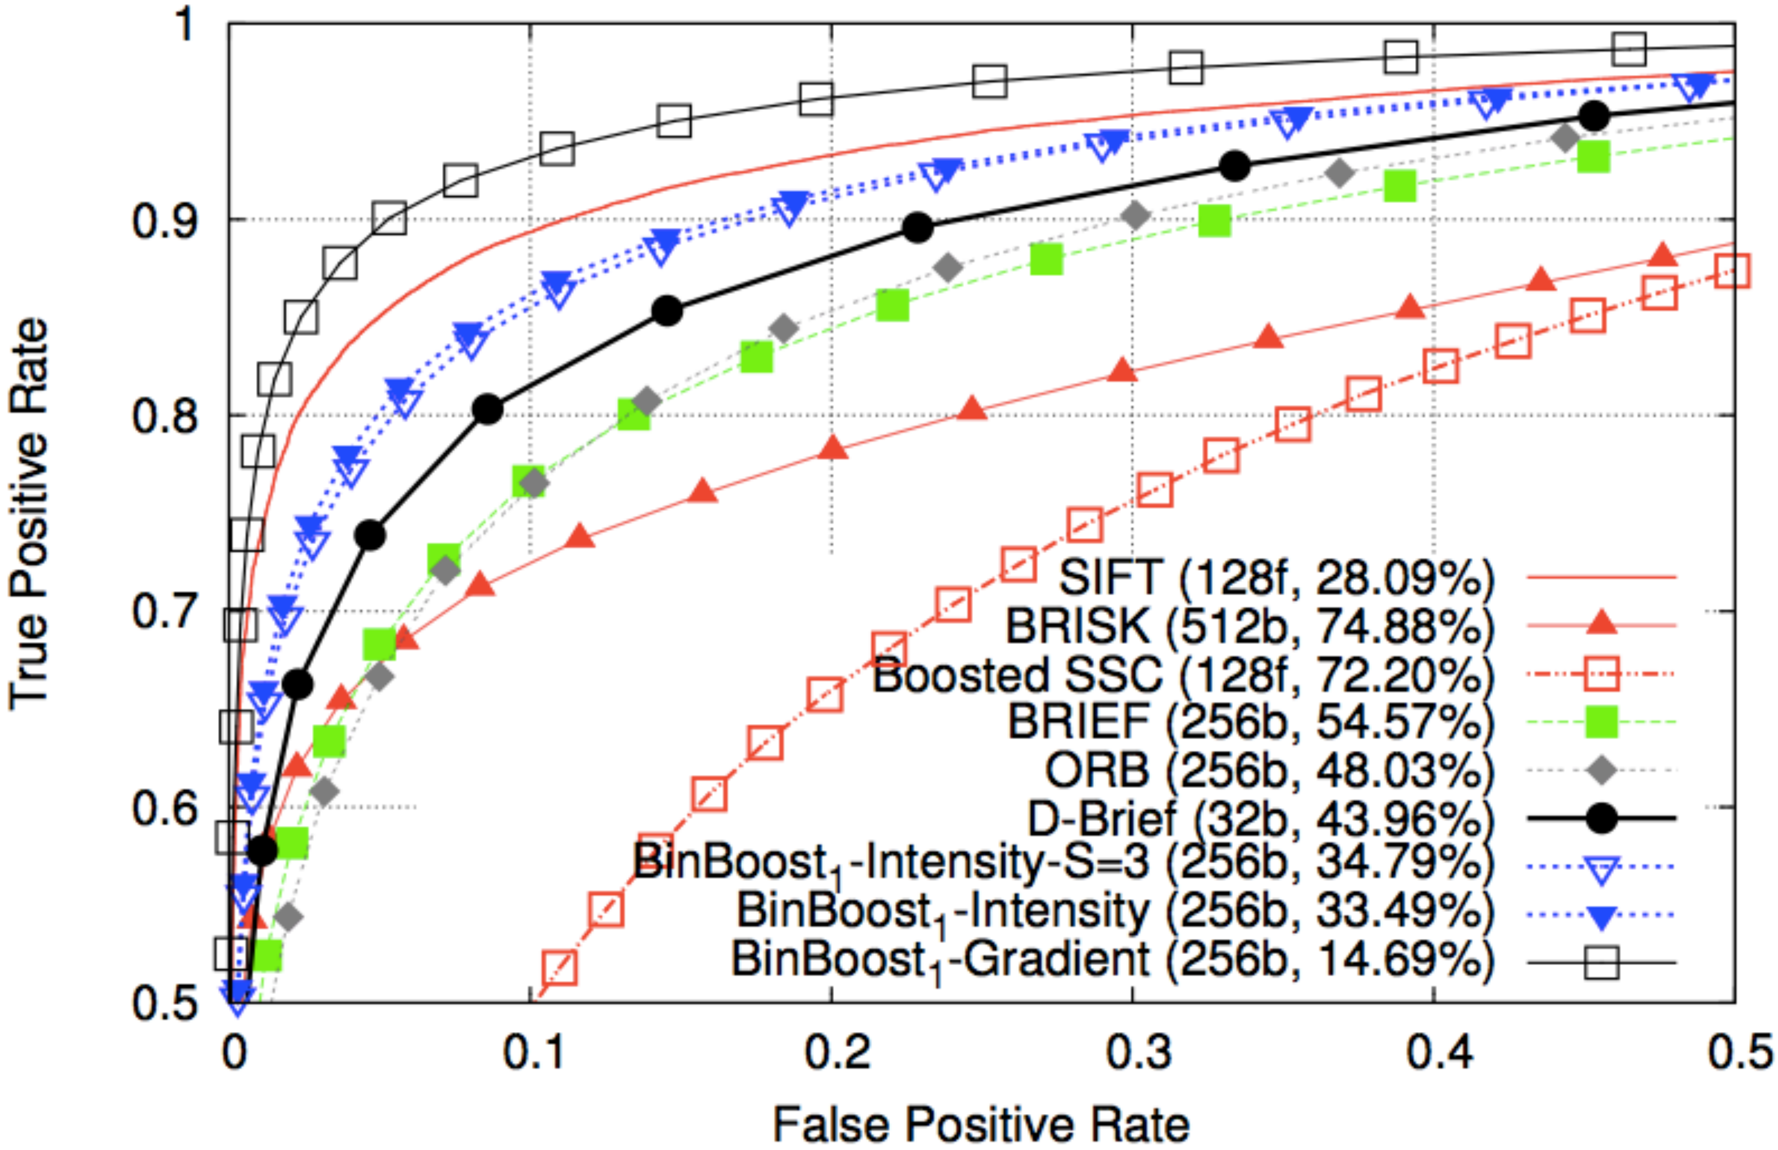

True Positive Rate
False Positive Rate
1
0.9
0.8
0.7
0.6
0.5
0
0.1
0.2
0.3
0.4
0.5
SIFT (128f, 28.09%)
BRISK (512b, 74.88%)
Boosted SSC (128f, 72.20%)
BRIEF (256b, 54.57%)
ORB (256b, 48.03%)
D-Brief (32b, 43.96%)
BinBoost1-Intensity-S=3 (256b, 34.79%)
BinBoost1-Intensity (256b, 33.49%)
BinBoost1-Gradient (256b, 14.69%)

Graphs

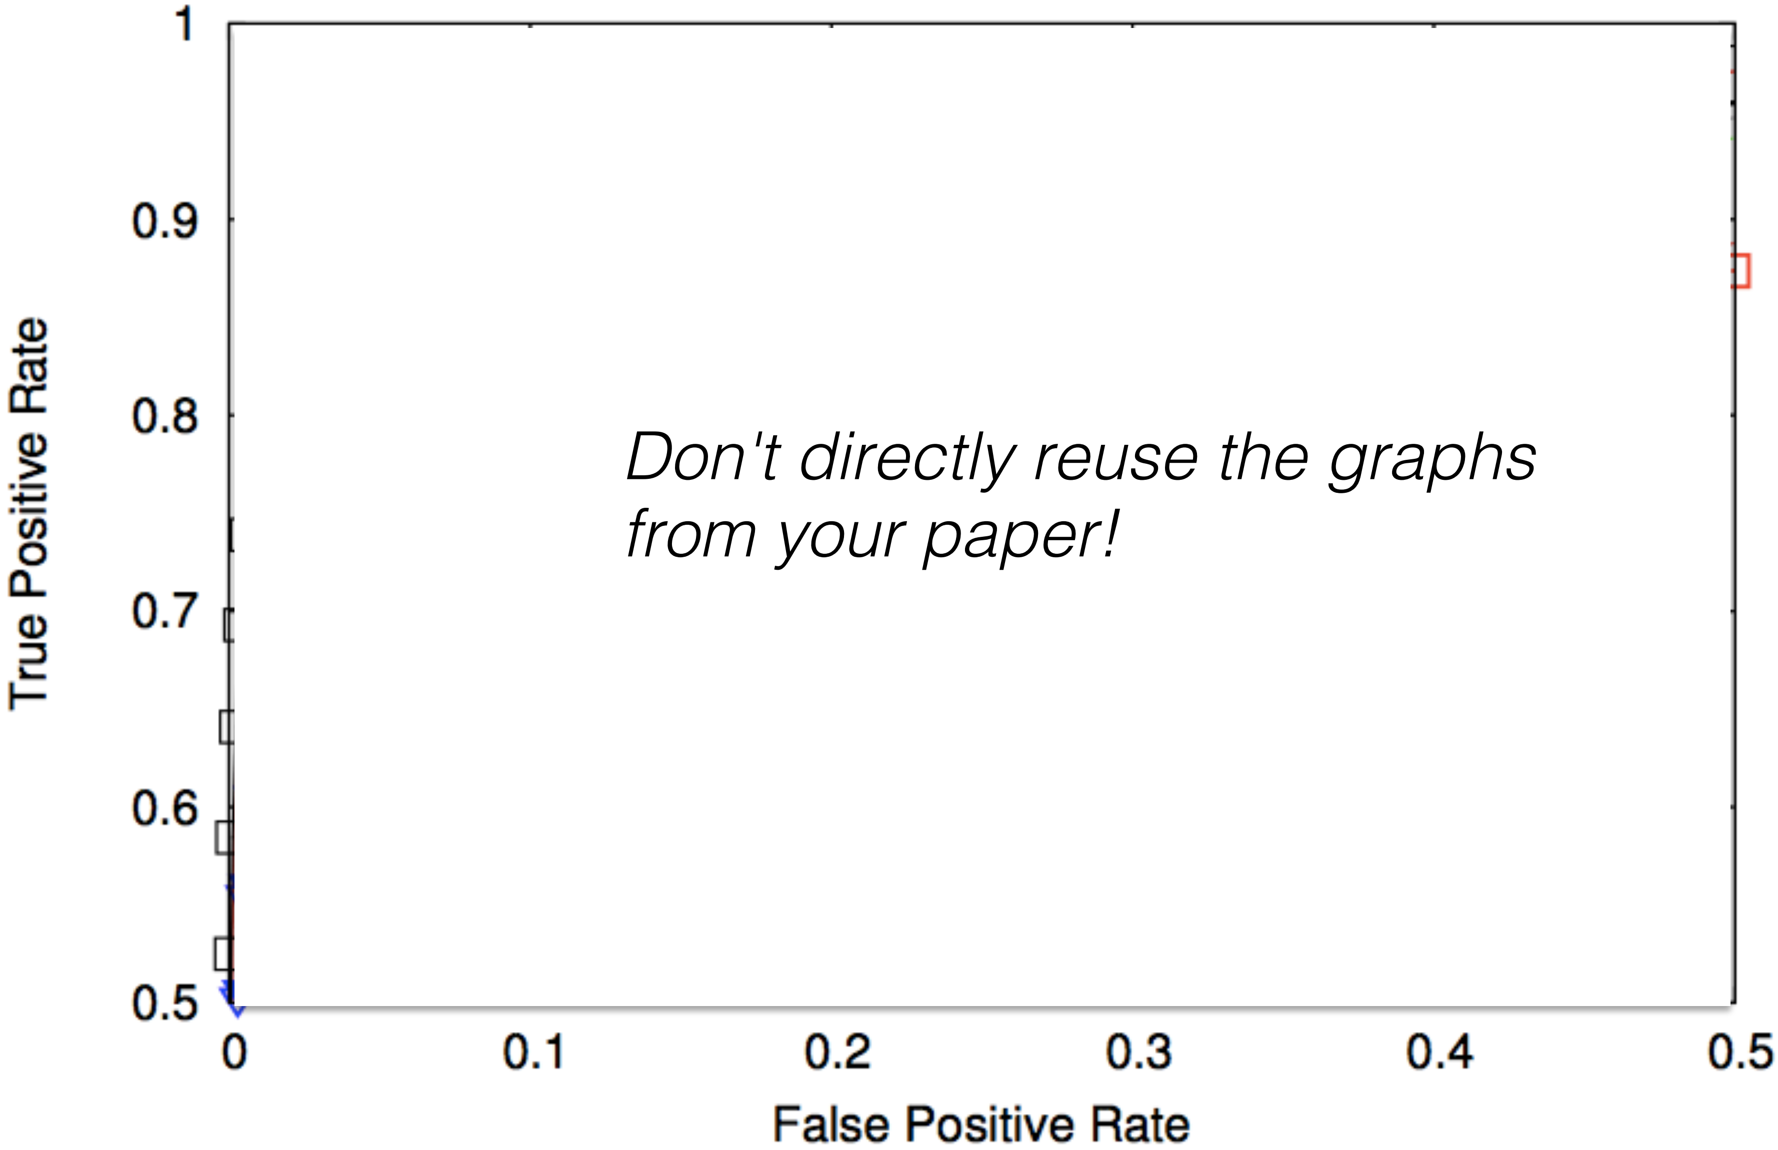

1
0.9
0.8
0.7
0.6
0.5
True Positive Rate
0
0.1
0.2
0.3
0.4
0.5
False Positive Rate
Don't directly reuse the graphs from your paper!

Graphs

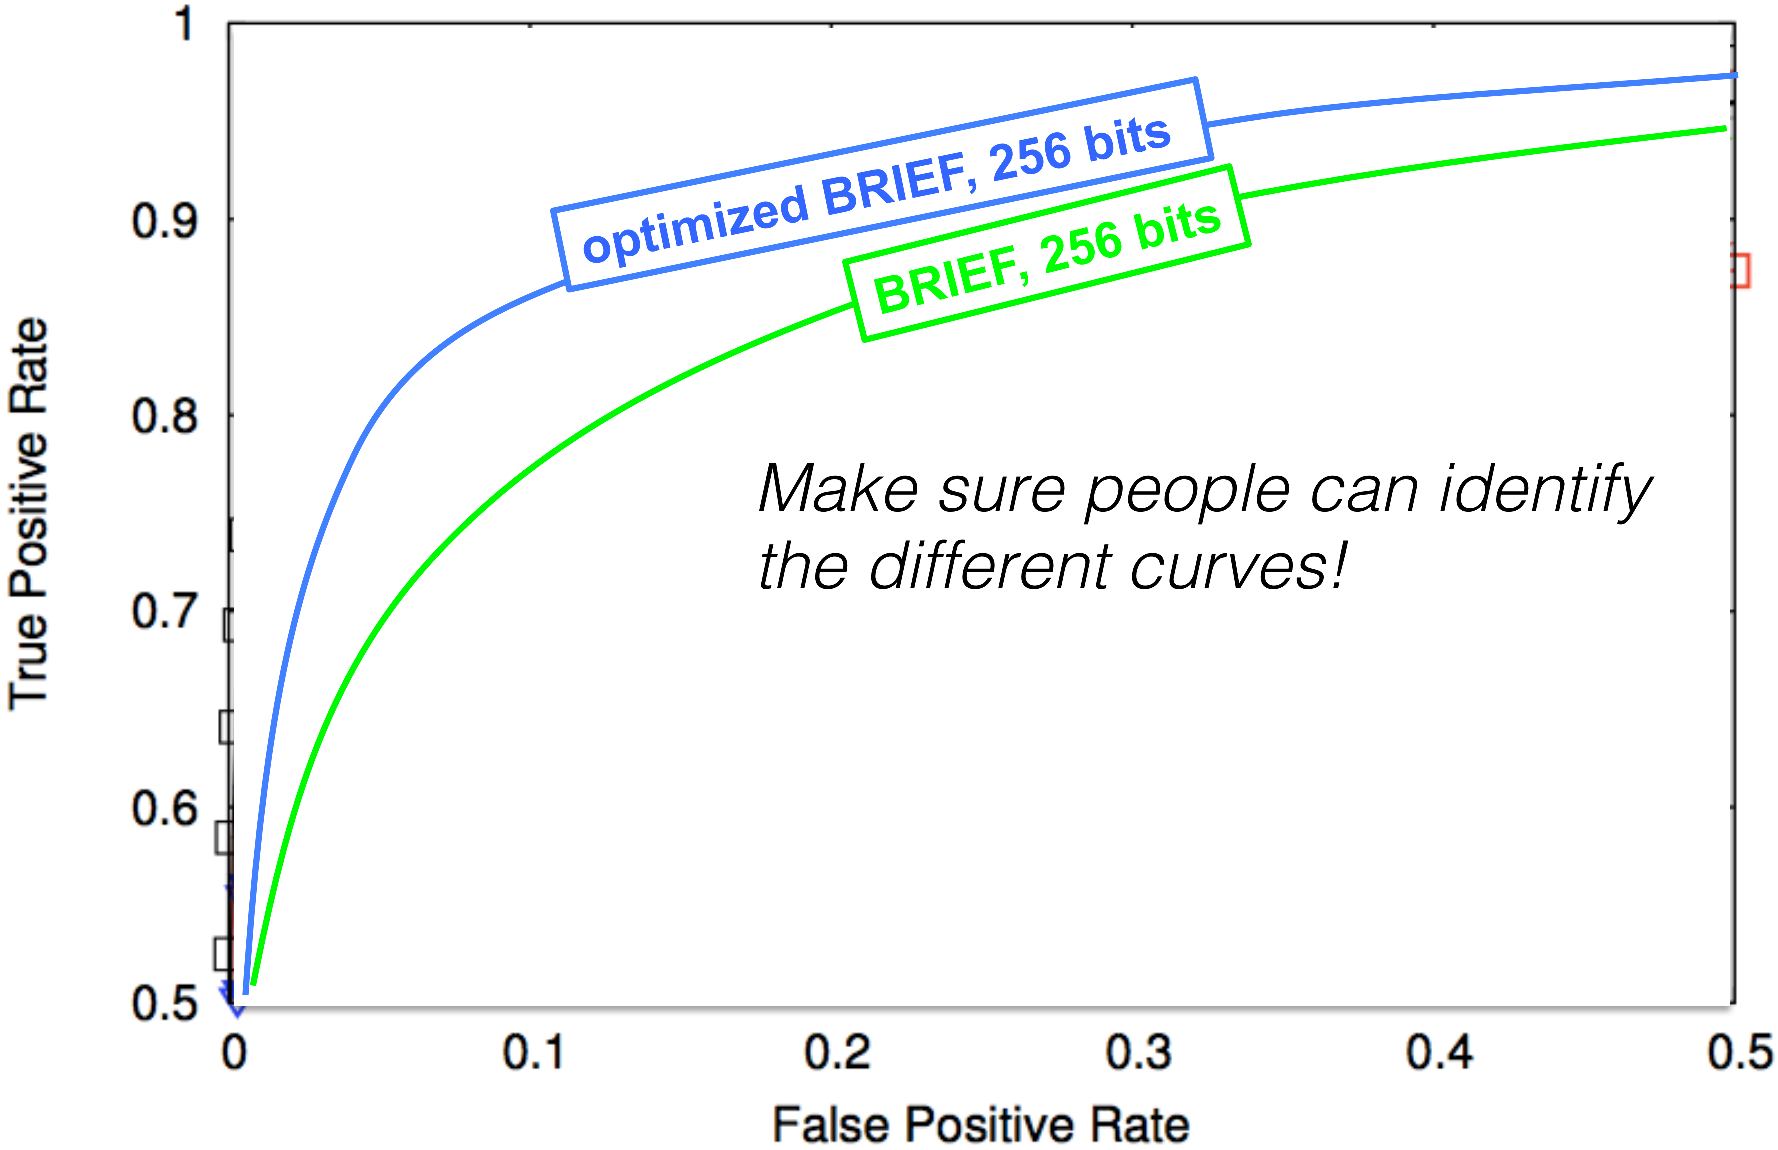

optimized BRIEF, 256 bits
BRIEF, 256 bits
Make sure people can identify the different curves!
True Positive Rate
False Positive Rate
1
0.9
0.8
0.7
0.6
0.5
0
0.1
0.2
0.3
0.4
0.5

Graphs

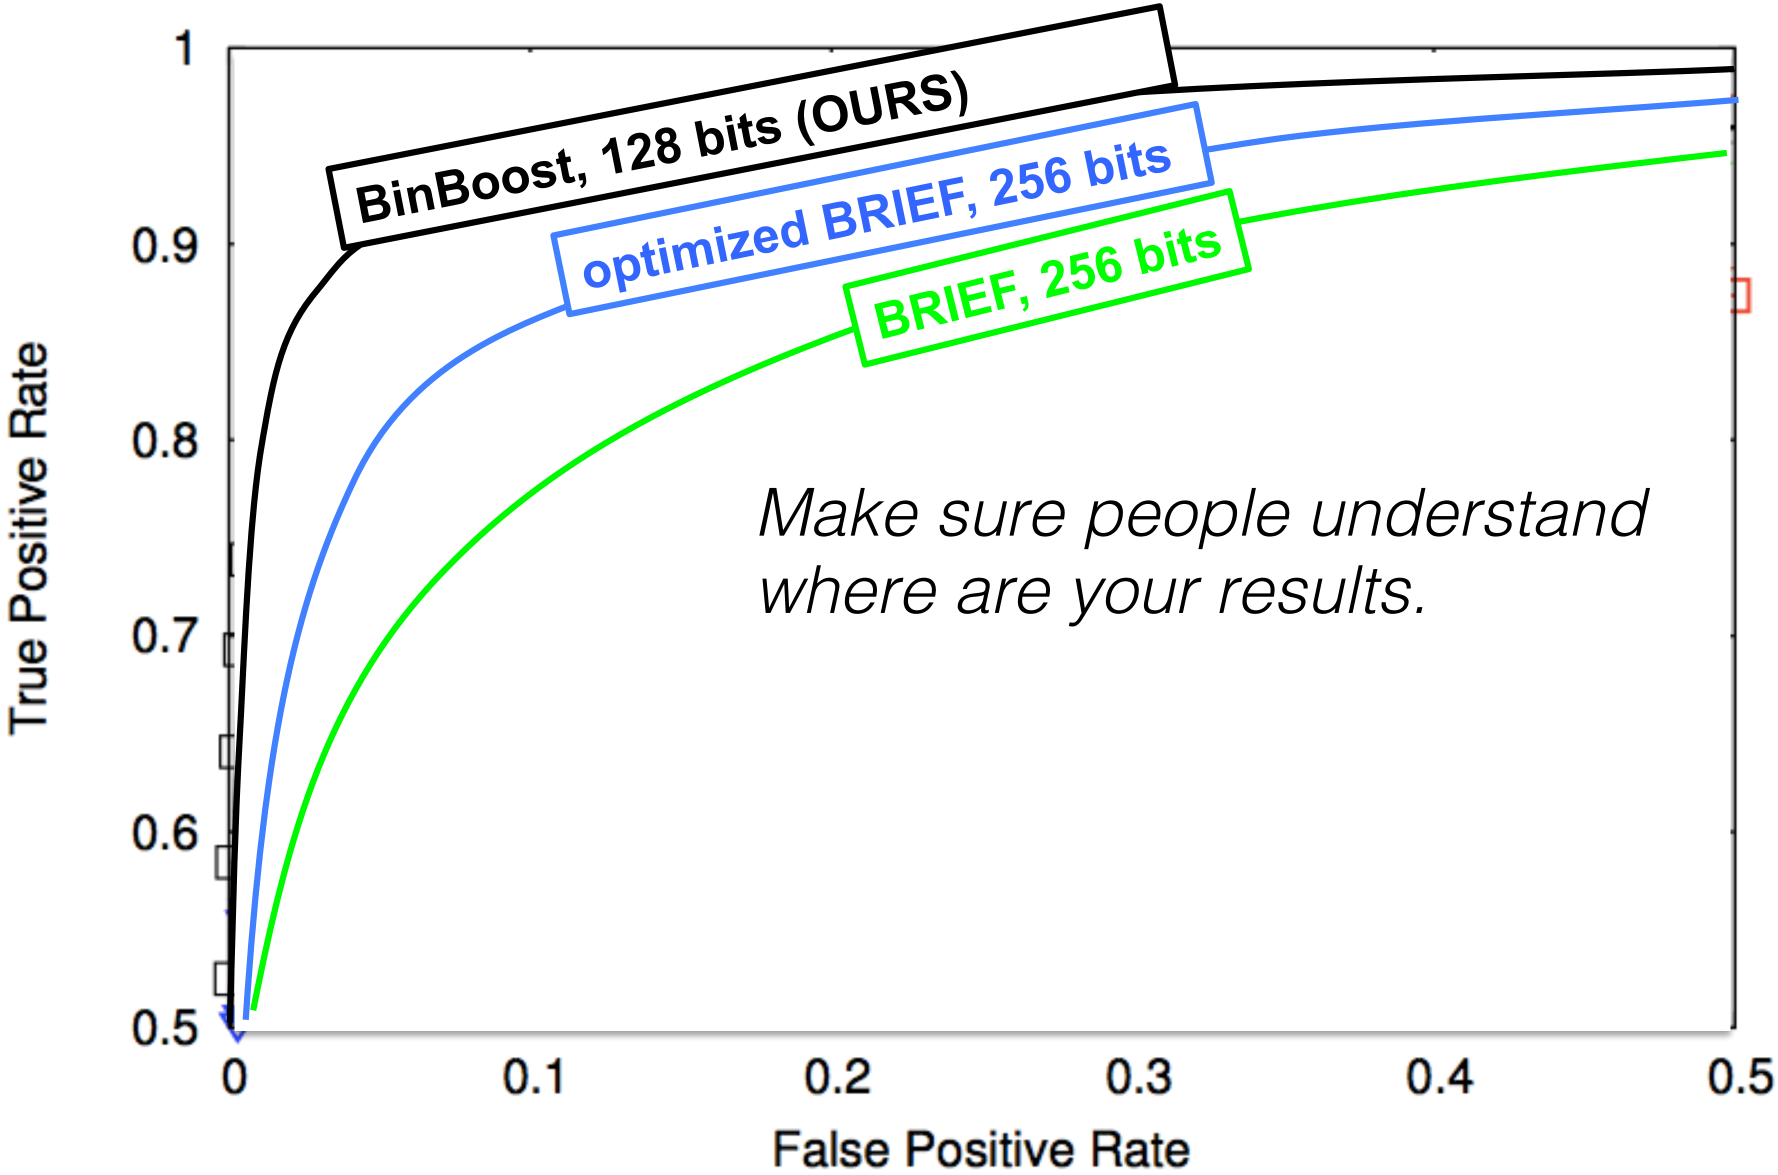

Make sure people understand where are your results.

Use PowerPoint/Keynote (except if you really have a lot of maths):

- you can draw figures easily;
- you can do last minute changes easily;
- you can see your script during the presentation;

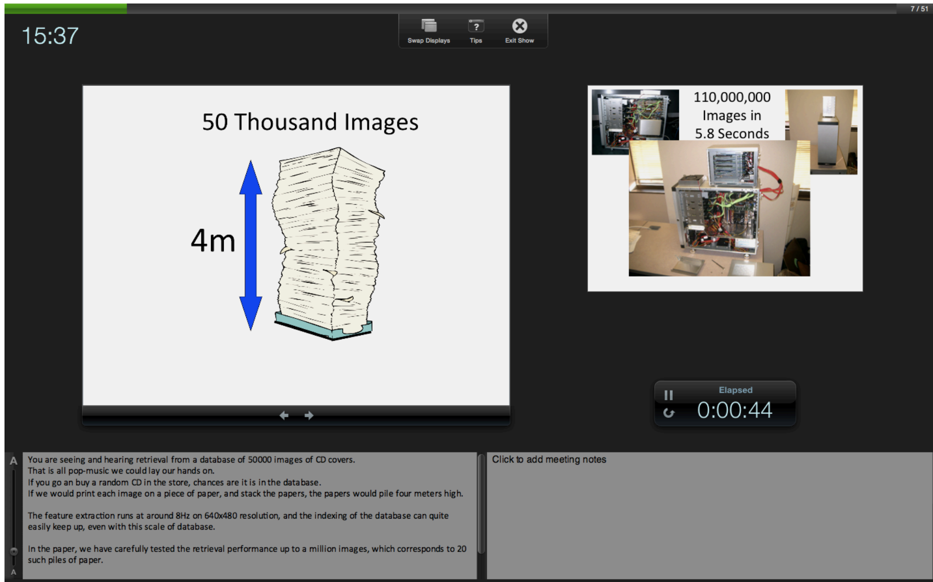

But be careful!

université de BORDEAUX

Title

The format Powerpoint/ Keynote encourages makes the slides difficult to remember (and boring).

- First item;
- Second item;
  - First sub-item;

- This is boring;
- And difficult to remember.

université de BORDEAUX

Animations

Don't use fancy effects

But constructing the slide progressively can be interesting sometimes.

Centerline Detection as a Classification Problem

Input image

Centerline Detection as a Classification Problem

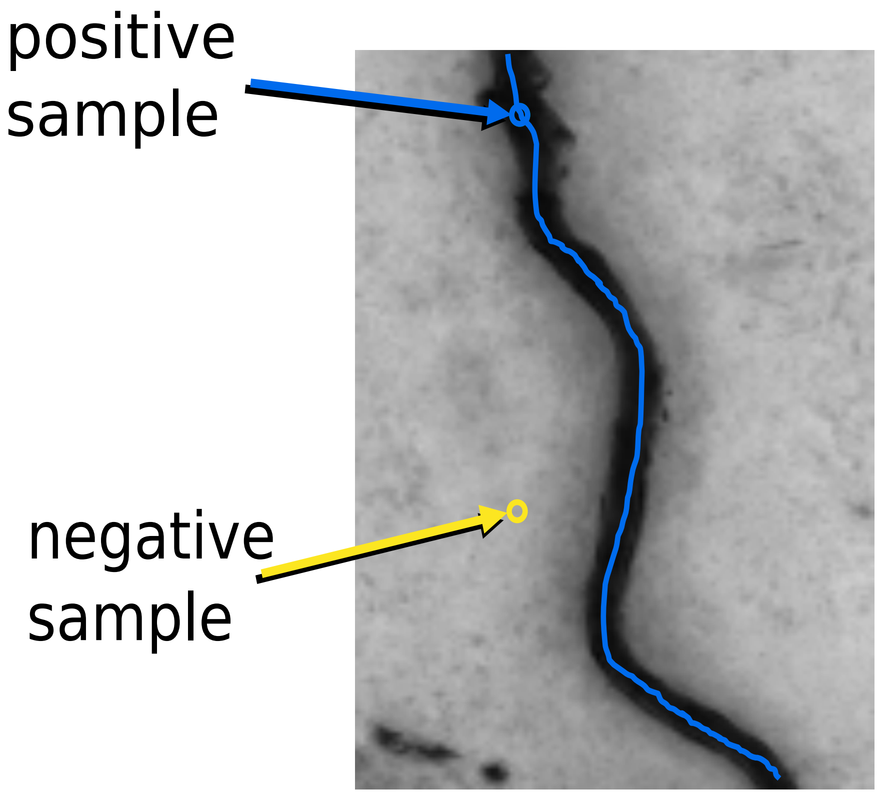

Input image

Centerline Detection as a Classification Problem

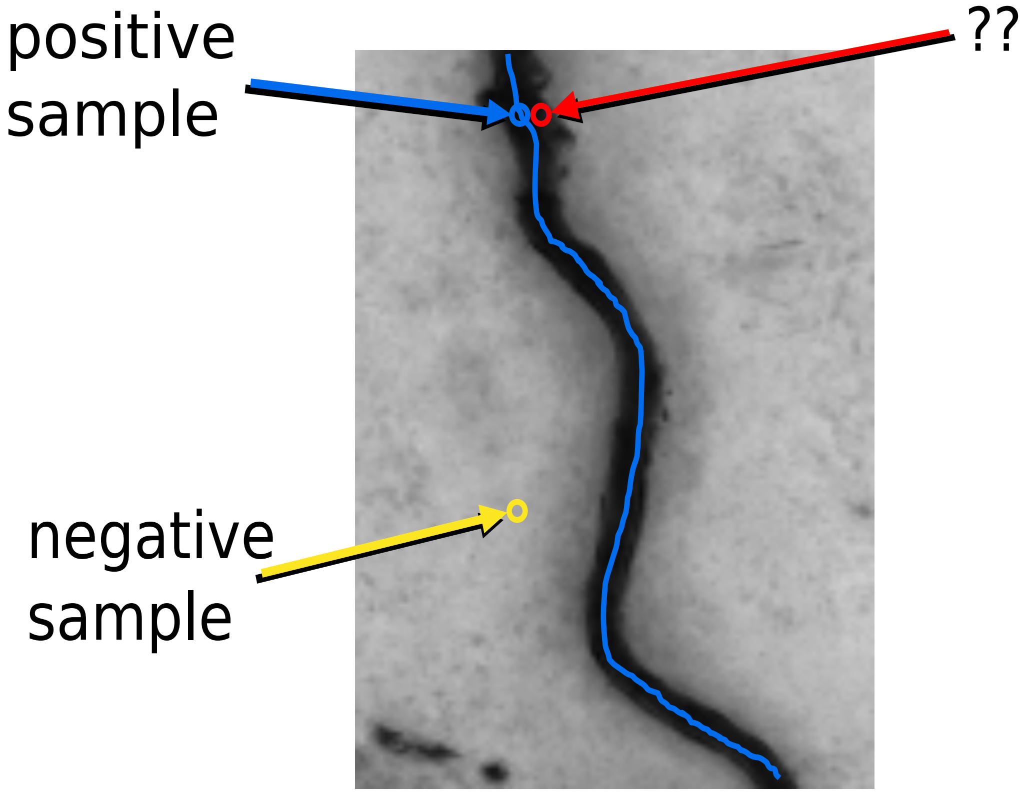

Input image

Presenting your Talk

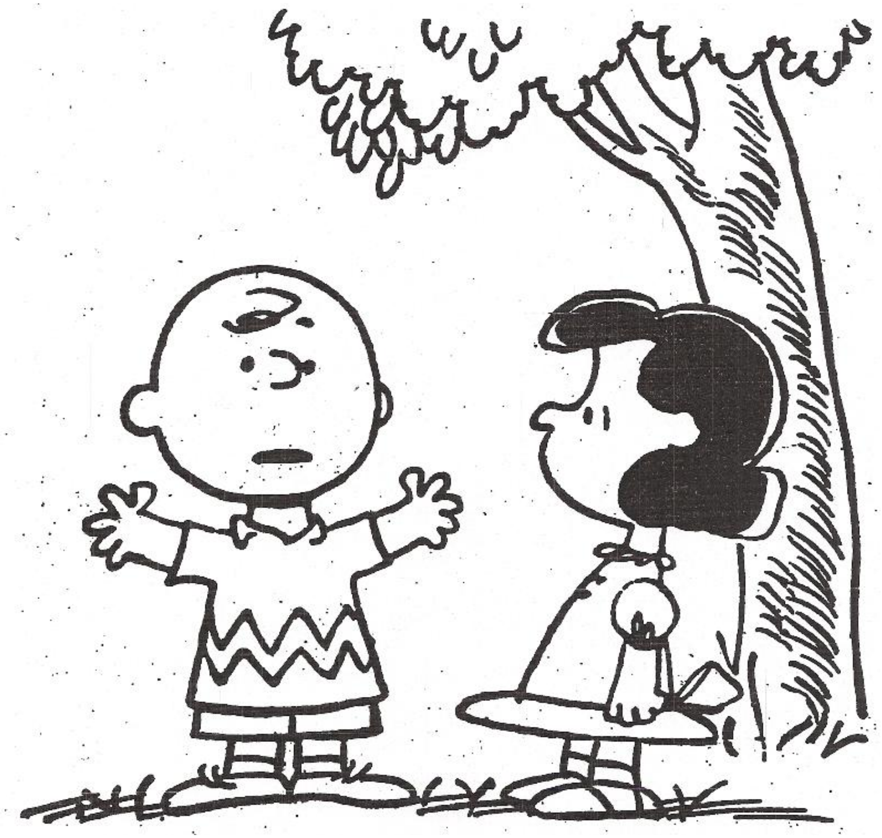

Rehearse (in front of people!)

Be enthusiastic!

Everybody is Nervous

- Deep breathing before;

- Script your first sentences precisely.

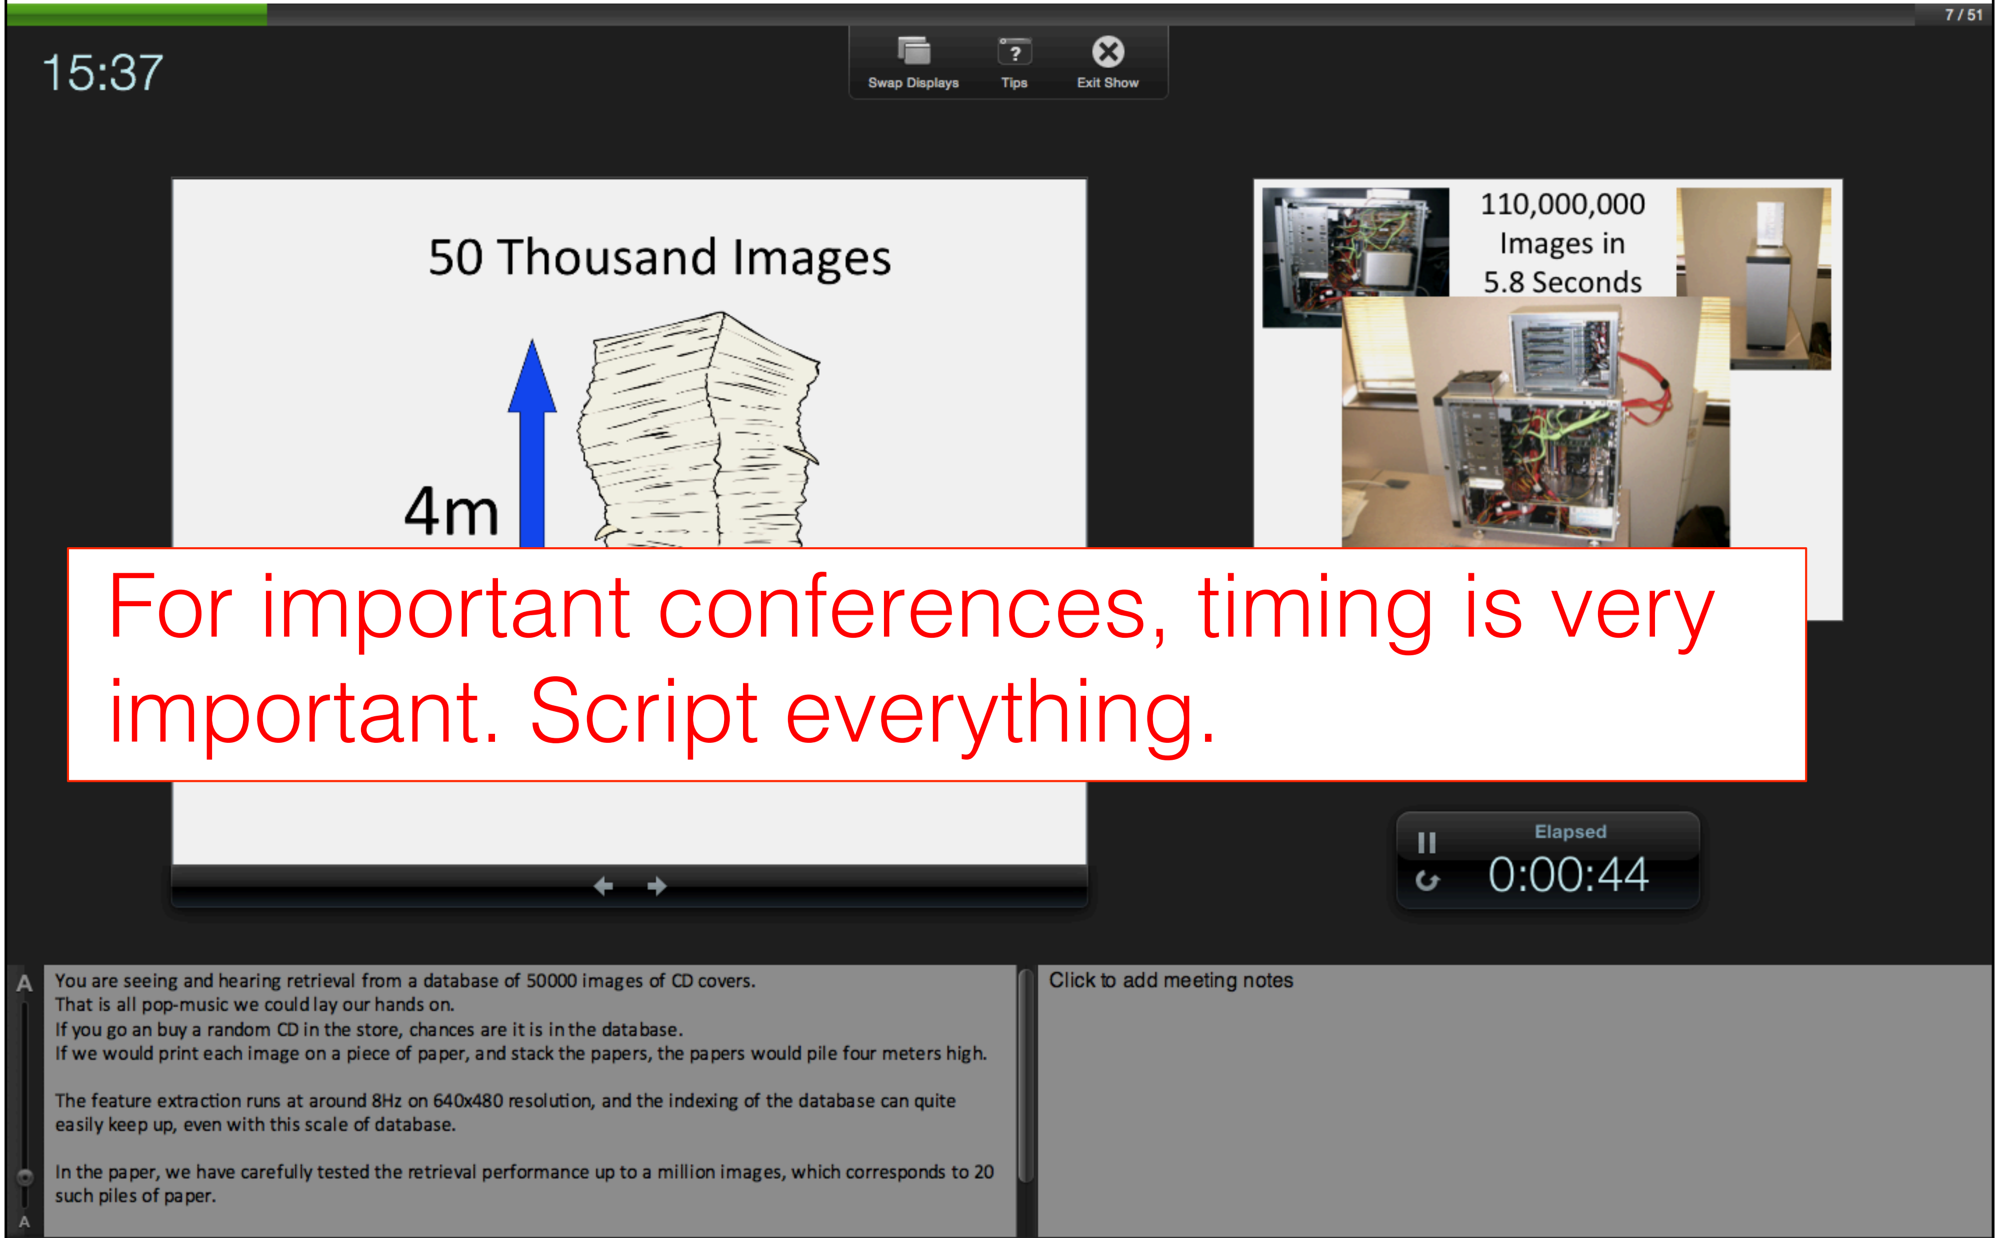

For important conferences, timing is very important. Script everything.

Watch Online Videos

Observe what the presenter did when you understand/are interested/are lost/are bored.

videolectures.net
exchange ideas & share knowledge

YouTube

TED
ted.com

A Good Reference

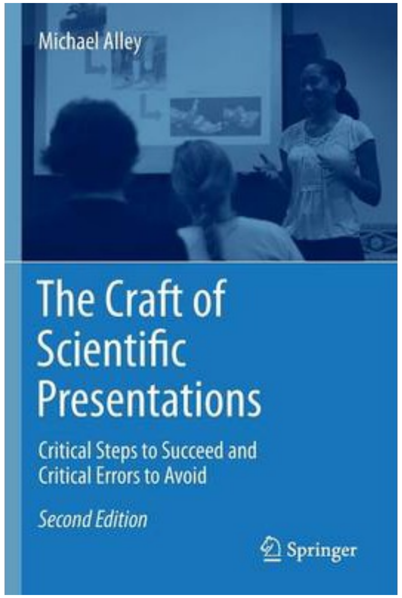

The Craft of Scientific Presentations (second edition)
Michael Alley

Springer

Make Sure the Audience Understands when You Are Done

Thanks!

Questions?

université de BORDEAUX