Writing a Good Research Paper

Vincent Lepetit

École des Ponts
ParisTech

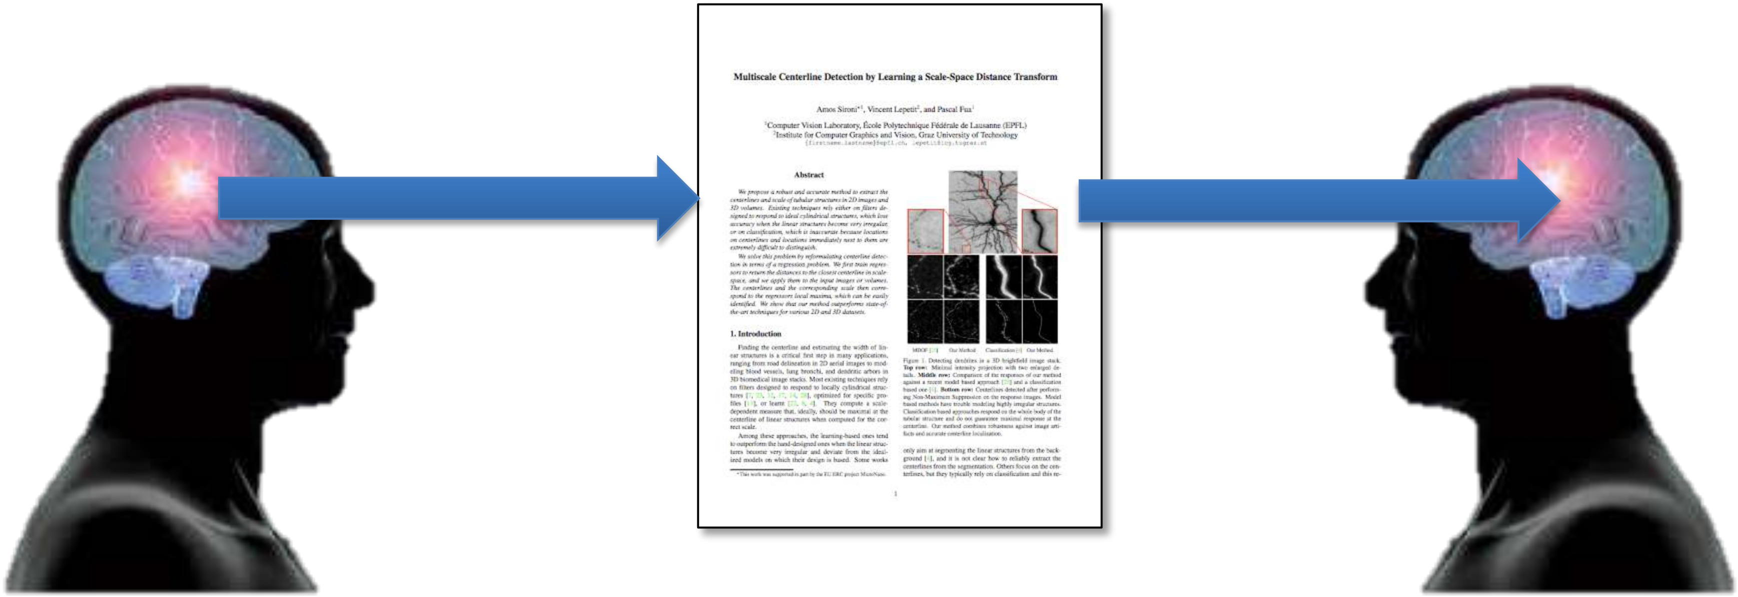

Your paper is the only thing the reviewers (and the readers) see of your work

They do not care about the quality of your code, the technical problems you encountered, …

→ your paper should be as good as possible

École des Ponts ParisTech

a good paper

Rule #1: Be as clear as possible

- a small misunderstanding can get your paper rejected;
- you need to guide the reviewer into your work, and anticipate what will be difficult to understand or could lead to a misunderstanding;
- and don't obfuscate.

Rule #2: Convince the reader your work is worth publishing

École des Ponts ParisTech

writing your paper

- title
- abstract
- introduction
- related work
- method
- experiments
- conclusion

done

this is not how it is going to be...

École des Ponts
ParisTech

writing your paper

add some text

get corrections from your advisor

adding a nice figure to explain your method

get feedback, take it into account

realize you need more experiments

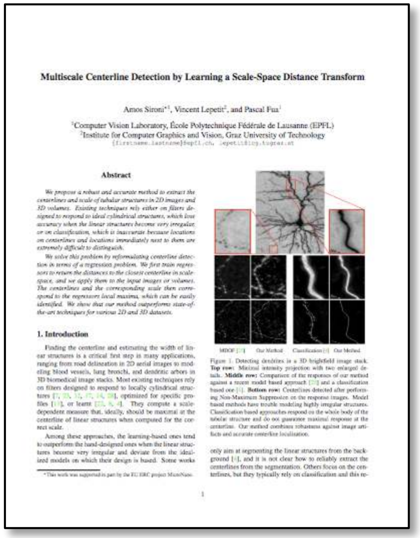

Multiscale Centerline Detection by Learning a Scale-Space Distance Transform

Amos Sironi*[1], Vincent Lepetit[2], and Pascal Fua[1]

[1]Computer Vision Laboratory, École Polytechnique Fédérale de Lausanne (EPFL)
[2]Institute for Computer Graphics and Vision, Graz University of Technology
{firstname.lastname}@epfl.ch, lepetit@icg.tugraz.at

Abstract

We propose a robust and accurate method to extract the centerlines and scale of tubular structures in 2D images and 3D volumes. Existing techniques rely either on filters designed to respond to ideal cylindrical structures, which lose accuracy when the linear structures become very irregular, or on classification, which is inaccurate because locations on centerlines and locations immediately next to them are extremely difficult to distinguish.

We solve this problem by reformulating centerline detection in terms of a regression problem. We first train regressors to return the distances to the closest centerline in scale-space, and we apply them to the input images or volumes. The centerlines and the corresponding scale then correspond to the regressors local maxima, which can be easily identified. We show that our method outperforms state-of-the-art techniques for various 2D and 3D datasets.

1. Introduction

Finding the centerline and estimating the width of linear structures is a critical first step in many applications, ranging from road delineation in 2D aerial images to modeling blood vessels, lung bronchi, and dendritic arbors in 3D biomedical image stacks. Most existing techniques rely on filters designed to respond to locally cylindrical structures [7, 23, 12, 17, 14, 28], optimized for specific profiles [11], or learnt [22, 8, 4]. They compute a scale-dependent measure that, ideally, should be maximal at the centerline of linear structures when computed for the correct scale.

Among these approaches, the learning-based ones tend to outperform the hand-designed ones when the linear structures become very irregular and deviate from the idealized models on which their design is based. Some works only aim at segmenting the linear structures from the background [4], and it is not clear how to reliably extract the centerlines from the segmentation. Others focus on the centerlines, but they typically rely on classification and this re-

MDOF [22] Our Method Classification [4] Our Method

Figure 1. Detecting dendrites in a 3D brightfield image stack. **Top row:** Minimal intensity projection with two enlarged details. **Middle row:** Comparison of the responses of our method against a recent model based approach [22] and a classification based one [4]. **Bottom row:** Centerlines detected after performing Non-Maximum Suppression on the response images. Model based methods have trouble modeling highly irregular structures. Classification based approaches respond on the whole body of the tubular structure and do not guarantee maximal response at the centerline. Our method combines robustness against image artifacts and accurate centerline localization.

*This work was supported in part by the EU ERC project MicroNano.

1

École des Ponts
ParisTech

the importance of feedback

- a small misunderstanding can get your paper rejected;
- it is difficult to read your own paper with the same state-of-mind as a reviewer;

→ getting the feedback of somebody who does not know your work is really important even if it can hurt.

- Each reader can only read your paper for the first time once! Use them carefully.

- The reviewer is always right. If they did not understand something, it is because you did not explain it clearly enough.

paper writing as an optimization process

$$-\log P_{\text{accepted}} = \mathcal{L}_{\text{clarity}} + \mathcal{L}_{\text{convincingness}}$$

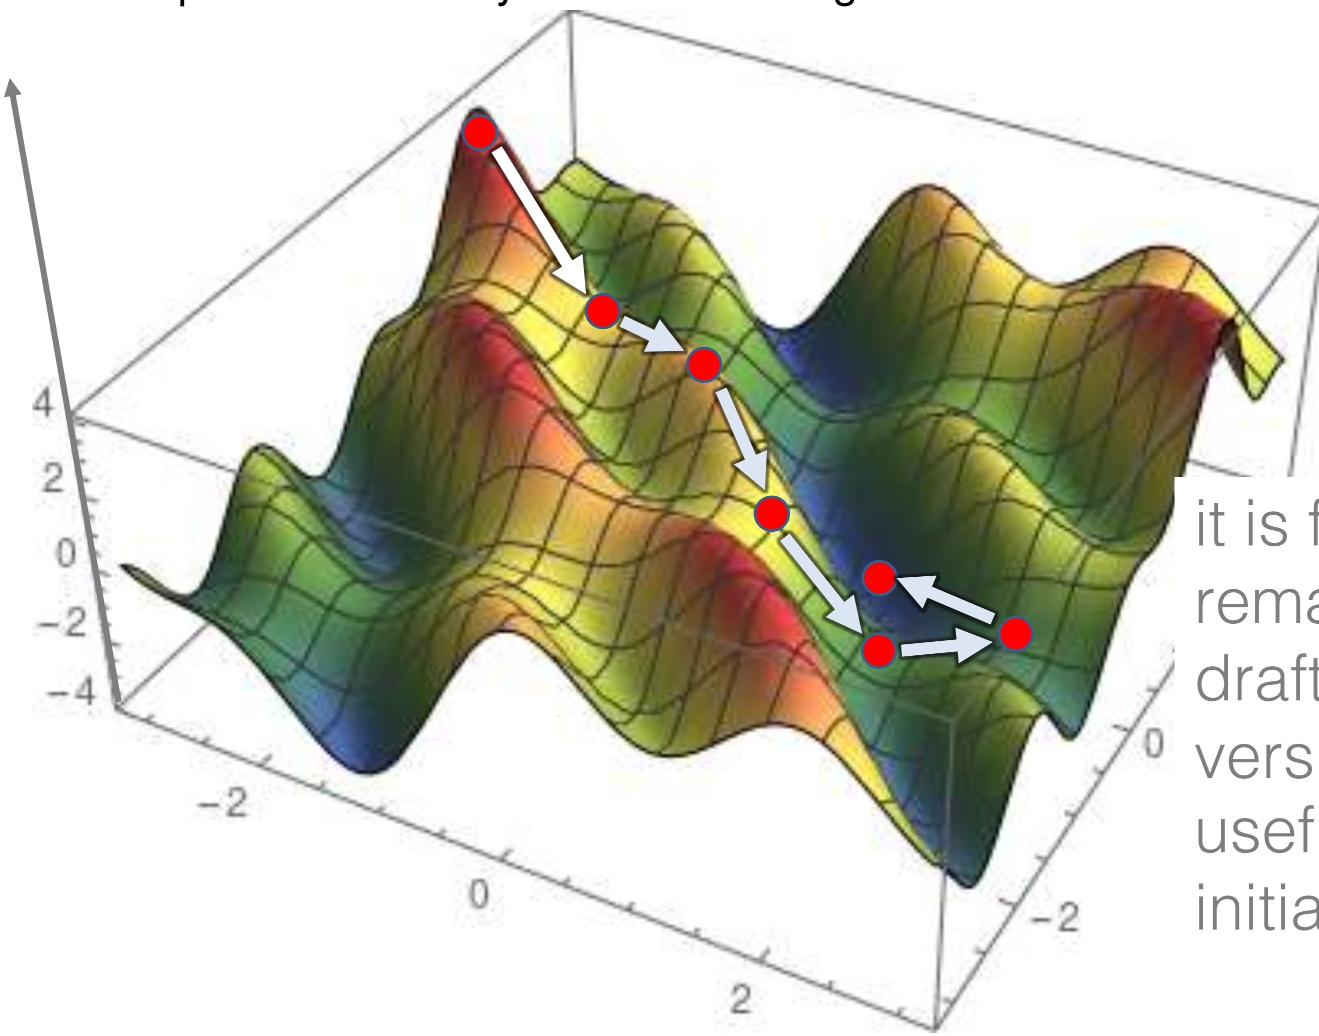

it is fine if nothing remains of your draft in the final version, it was useful as an initialization

saving some precious iterations

First write short sentences, you can make the text sound nicer latter.

Keep the paragraphs short as well.

Use sections / subsections / paragraphs to structure your text.

keep a clear structure

Doc. 3

KINETIC THEORY OF THERMAL EQUILIBRIUM AND OF THE SECOND LAW OF THERMODYNAMICS

by A. Einstein

[*Annalen der Physik* 9 (1902): 417-433]

Great as the achievements of the kinetic theory of heat have been in the domain of gas theory, the science of mechanics has not yet been able to produce an adequate foundation for the general theory of heat, for one has not yet succeeded in deriving the laws of thermal equilibrium and the second law of thermodynamics using only the equations of mechanics and the probability calculus, though Maxwell's and Boltzmann's theories came close to this goal.
[1] The purpose of the following considerations is to close this gap. At the same time, they will yield an extension of the second law that is of importance for
[2] the application of thermodynamics. They will also yield the mathematical expression for entropy from the standpoint of mechanics.

§1. *Mechanical model for a physical system*

Let us imagine an arbitrary physical system that can be represented by a mechanical system whose state is uniquely determined by a very large number of
[3] coordinates $p_1 \dots p_n$ and the corresponding velocities

$$\frac{dp_1}{dt}, \dots \frac{dp_n}{dt}.$$

Let their energy E consist of two additive terms, the potential energy V and the kinetic energy L. The former shall be a function of the coordinates alone, and the latter shall be a quadratic function of

$$\frac{dp_\nu}{dt} = p'_\nu ,$$

whose coefficients are arbitrary functions of the p's. Two kinds of external forces shall act upon the masses of the system. One kind of force shall be derivable from a potential V_a and shall represent external conditions (gravity, effect of rigid walls without thermal effects, etc.); their potential may contain time explicitly, but its derivative with respect to time should be very small. The other forces shall not be derivable from a potential and shall vary rapidly. They have to be conceived as the forces that produce the influx of heat. If such forces do not act, but V_a depends explicitly on time, then we are dealing with an adiabatic process.

Also, instead of velocities we will introduce linear functions of them, the momenta $q_1, \dots, q_n$, as the system's state variables, which are defined by n equations of the form

$$q_\nu = \frac{\partial L}{\partial p'_\nu},$$

where L should be conceived as a function of the $p_1, \dots, p_n$ and $p'_1, \dots, p'_n$.

§2. *On the distribution of possible states between N identical adiabatic stationary systems, when the energy contents are almost identical.*

Imagine infinitely many (N) systems of the same kind whose energy content is continuously distributed between definite, very slightly differing
values $\bar{E}$ and $\bar{E} + \delta E$. External forces that cannot be derived from a poten- [4]
tial shall not be present, and V_a shall not contain the time explicitly, so that the system will be a conservative one. We examine the distribution of states, which we assume to be stationary.

We make the assumption that except for the energy $E = L + V_a + V_i$, or a [5]
function of this quantity, for the individual system, there does not exist any
function of the state variables p and q which remains constant in time; we [6]
shall henceforth consider only systems that satisfy this condition. Our assumption is equivalent to the assumption that the distribution of states of our systems is determined by the value of E and is spontaneously established from any arbitrary initial values of the state variables that satisfy our condition regarding the value of energy. I.e., if there would exist for the

[https://einsteinpapers.press.princeton.edu/]

use the active form

The passive form can be boring and ambiguous:

“the method is called..”
→ “we call our method ..”

“it can be seen in Table 1..”
→ “Table 1 shows that..”

don’t be afraid of repetitions

Avoid pronouns such as “it”, “they”:

“..to reach a “high-level understanding” of the scene’s geometry, by simplifying it drastically.”
→ “..to reach a “high-level understanding” of the scene’s geometry, by simplifying this geometry drastically.”

In general, do not be afraid to repeat key words to avoid ambiguities.

the/a/<nothing> in English

“the COCO dataset” or simply “COCO”

NOT

“COCO dataset” or “the COCO”

“the object” refers to a specific object;
“an object” refers to any object.

equations

An equation should be part of a sentence. The reader should have a decent understanding of the method by replacing the equation by “blah” for the first read.

Notice the punctuation (the comma after the equation):

Given training samples $\{(f_i, y_i)\}_i$, where $f_i = f(\mathbf{x}_i, I_i) \in \mathbb{R}^M$ is the feature vector corresponding to a point $\mathbf{x}_i$ in image I_i and $y_i = d(\mathbf{x}_i)$, GradientBoost approximates $y(\cdot)$ by a function of the form

$$\varphi(f(\mathbf{x}, I)) = \sum_{k=1}^{K} \alpha_k h_k(f(\mathbf{x}, I)) \,, \tag{4}$$

where $h_k : \mathbb{R}^M \to \mathbb{R}$ are weak learners and $\alpha_k \in \mathbb{R}$ are weights. Function φ is built iteratively, selecting one weak

nice looking equations

$$-\log P_\text{accepted} = \mathcal{L}_\text{clarity} + \mathcal{L}_\text{convincingness}$$

```
-\log P_\text{accepted} = \calL_\text{clarity} +
\calL_\text{convincingness}
```

NOT:

$$-\log P_{accepted} = \mathcal{L}_{clarity} + \mathcal{L}_{convincingness}$$

```
-\log P_{accepted} = \calL_{clarity} + \calL_{convincingness}
```

École des Ponts
ParisTech

default tables in LaTeX are ugly

signal processing concept	algebraic concept (coordinate free)	in coordinates
fi lter	$h \in \mathcal{A}$ (algebra)	$\phi(h) \in \mathbb{C}^{I \times I}$
signal	$s = \sum s_i b_i \in \mathcal{M}$ ($\mathcal{A}$-module)	$\mathbf{s} = (s_i)_{i \in I} \in \mathbb{C}^I$
fi ltering	$h \cdot s$	$\phi(h) \cdot \mathbf{s}$
impulse	base vector $b_i \in \mathcal{M}$	$\mathbf{b}_i = (\dots, 0, 1, 0 \dots)^T \in \mathbb{C}^I$
impulse response of $h \in \mathcal{A}$	$h \cdot b_i \in \mathcal{M}$	$\phi(h) \cdot \mathbf{b}_i = (\dots, h_{-1}, h_0, h_1, \dots)^T \in \mathbb{C}^I$
Fourier transform	$\Delta: \mathcal{M} \to \bigoplus_{\omega \in W} \mathcal{M}_\omega$	$\mathcal{F}: \mathbb{C}^I \to \bigoplus_{\omega \in W} \mathbb{C}^{d_\omega}$
		$\Leftrightarrow \phi \to \bigoplus_{\omega \in W} \phi_\omega$
spectrum of signal	$\Delta(s) = (s_\omega)_{\omega \in W} = \omega \mapsto s_\omega$	$\mathcal{F}(\mathbf{s}) = (\mathbf{s}_\omega)_{\omega \in W} = \omega \mapsto \mathbf{s}_\omega$
frequency response of $h \in \mathcal{A}$		$(\phi_\omega(h))_{\omega \in W} = \omega \mapsto \phi_\omega(h)$

signal processing concept	**algebraic concept (coordinate free)**	**in coordinates**
filter	$h \in \mathcal{A}$ (algebra)	$\phi(h) \in \mathbb{C}^{I \times I}$
signal	$s = \sum s_i b_i \in \mathcal{M}$ ($\mathcal{A}$-module)	$\mathbf{s} = (s_i)_{i \in I} \in \mathbb{C}^I$
filtering	$h \cdot s$	$\phi(h) \cdot \mathbf{s}$
impulse	base vector $b_i \in \mathcal{M}$	$\mathbf{b}_i = (\dots, 0, 1, 0, \dots)^T \in \mathbb{C}^I$
impulse response of $h \in \mathcal{A}$	$h \cdot b_i \in \mathcal{M}$	$\phi(h) \cdot \mathbf{b}_i = (\dots, h_{-1}, h_0, h_1, \dots)^T \in \mathbb{C}^I$
Fourier transform	$\Delta: \mathcal{M} \to \bigoplus_{\omega \in W} \mathcal{M}_\omega$	$\mathcal{F}: \mathbb{C}^I \to \bigoplus_{\omega \in W} \mathbb{C}^{d_\omega} \Leftrightarrow \phi \to \bigoplus_{\omega \in W} \phi_\omega$
spectrum of signal	$\Delta(s) = (s_\omega)_{\omega \in W} = \omega \mapsto s_\omega$	$\mathcal{F}(\mathbf{s}) = (\mathbf{s}_\omega)_{\omega \in W} = \omega \mapsto \mathbf{s}_\omega$
frequency response of $h \in \mathcal{A}$	n.a.	$(\phi_\omega(h))_{\omega \in W} = \omega \mapsto \phi_\omega(h)$

http://www.inf.ethz.ch/personal/markusp/teaching/guides/guide-tables.pdf

Google “making nice tables latex ethz”

a trick to make tables (and other stuff) fit

```
\begin{table}
  \caption{\label{tab:auc_shapenet}{\bf title}, blah.
  \centering
  \scalebox{0.96}{
    \begin{tabular}{@{}lcccccccc@{}}
    ...
    \end{tabular}
  }
\end{table}
```

École des Ponts
ParisTech

figures

graphs and images should be large enough, including the labels.

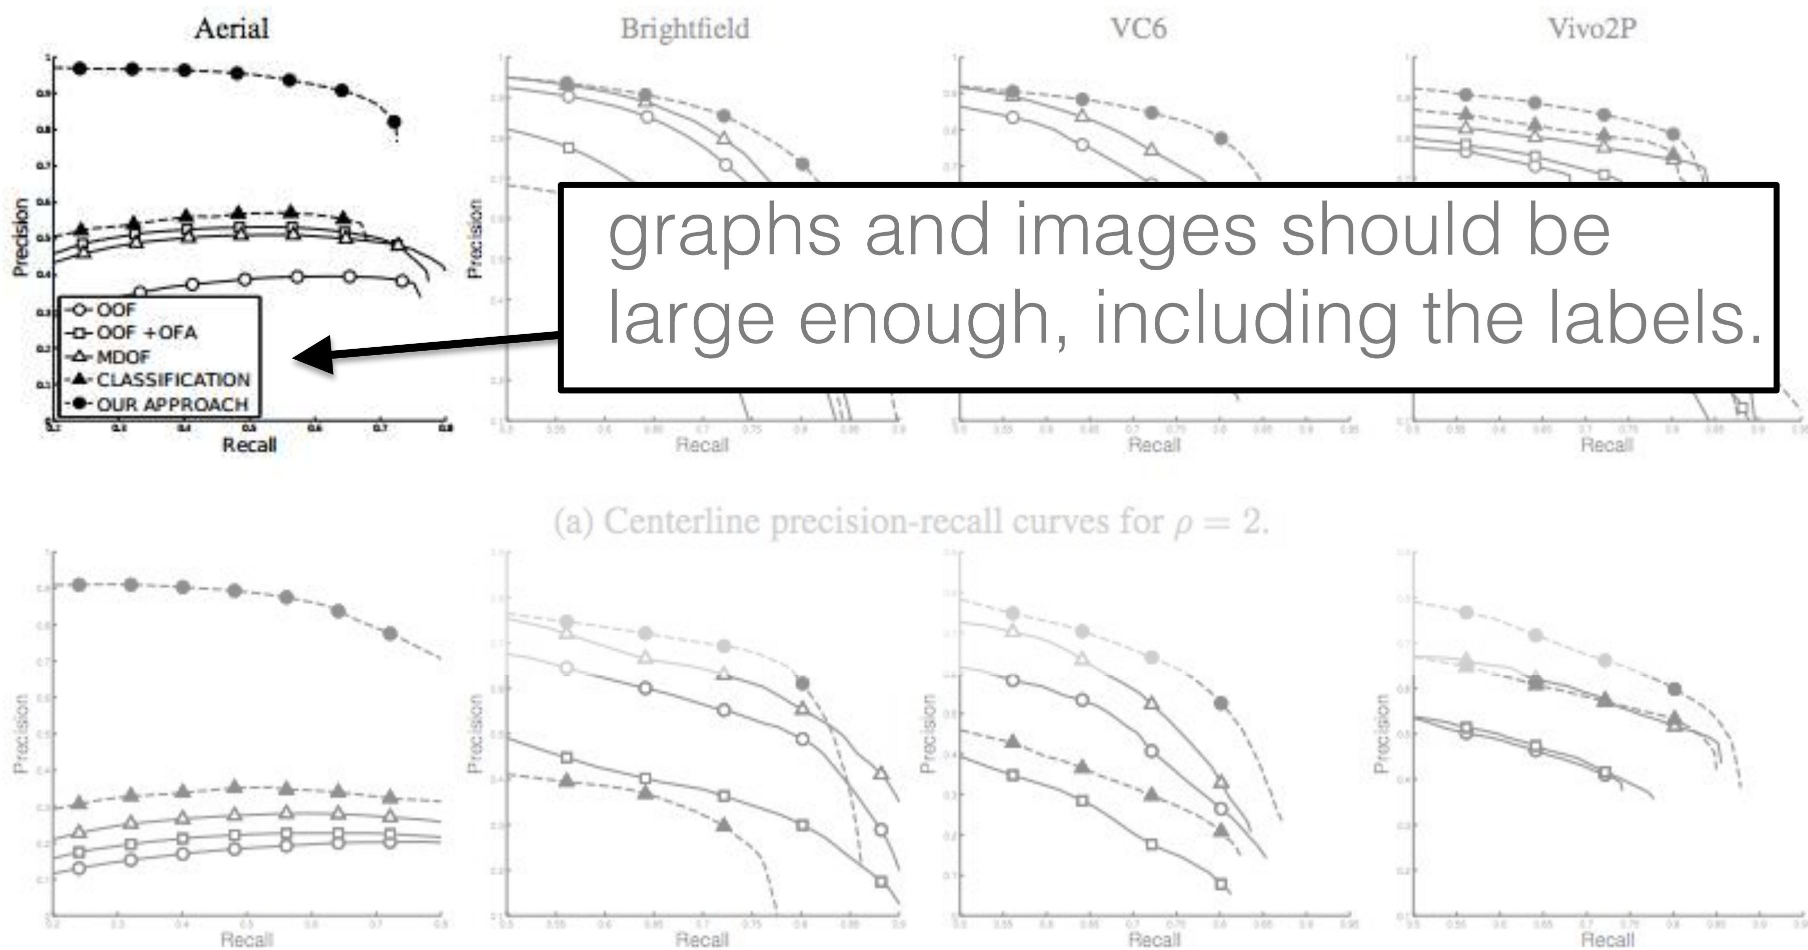

(a) Centerline precision-recall curves for $\rho = 2$.

(b) Segmentation precision-recall curves for $\delta = 0.4$.

Figure 8. Precision Recall curves. Our method outperforms the others on all the datasets we considered, both for centerline detection and joint centerline and radius estimation.

captions

name description

Figure 8. **Precision Recall curves.** Our method outperforms the others on all the datasets we considered, both for centerline detection and joint centerline and radius estimation.

Caption should start with the **name** of the figure.

Then give a **description**: Tell the reviewer what they should look at in the figure.

The figure and its caption should be self-contained.

referencing figures

Then, we can rely on simple non-maximum suppression to localize the centerlines. We will show in the next section that this solution is significantly more robust than both classification-based and filter-based methods (see Fig. 3).

no

but:

Then, as shown in Figure 3, we can rely on simple non-maximum suppression to localize the centerlines. We will show in the next section that this solution is significantly more robust than both classification-based and filter-based methods.

yes

École des Ponts ParisTech

use clear phrases

Don’t use jargon, don’t use smart-sounding phrases:

“Don't say 'reflective acoustic wave.' Say 'echo.’” [Richard Feynman]

latex code is code

it should be clear as well:

```
\subsection{Maximizing Surface Coverage Gain on a Binary Occupancy Map}

Here, we consider a binary occupancy map $\occfield:\IR^3 \rightarrow
\{0,1\}$ representing the volume of the target object or scene. We will relax
our derivations to a probabilistic occupancy map when looking for the next
best view in the next subsections. From the binary map $\occfield$, we can
define the set $\chi$ of occupied points, \ie, the set of points $x$
verifying $\occfield(x) = 1$, its surface as the boundary $\partial\chi$, and
the surface coverage $C(c)$ achieved by a camera pose $c=(c_\pos,
c_\rot)\in\calC:=\IR^3\times SO(3)$ as the following surface integral:
%
\begin{equation}
    \label{eqn:surface_coverage_definition}
    C(c) = \frac{1}{|\partial\chi|_S} \int_{\partial\chi}
\mathbb{1}_{\chi_c}(x) \cdot \absvis_c(x)\,\text{d}x \> ,
\end{equation}
%
where $|\partial \chi|_S := \int_{\partial\chi}\text{d}x$ is the area of
surface $\partial\chi$. $\chi_c \subset \chi$ is the subset of occupied
points contained in the field of view of camera $c$, and $\absvis_c(x)$ is
the visibility of point $x$, \ie, $\absvis_c(x) =
\mathbb{1}\left(\occfield\left(\{(1-\lambda) c_\pos + \lambda x \text{ such
that } \lambda\in[0,1)\}\right)=\{0\}\right)$.
```

École des Ponts ParisTech

latex code is code

you can introduce variables:

```
\newcommand{\calL}{\mathcal{L}}
\newcommand{\obs}{\text{obs}}
\newcommand{\reg}{\text{reg}}
```

to write

```
\calL = \calL_\obs + \calL_\reg
```

instead of

```
\mathcal{L} = \mathcal{L}_\text{obs} +
\mathcal{L}_\text{obs}
```

École des Ponts
ParisTech

references

[1] Ricardo Cabral and Yasutaka Furukawa. Piecewise Planar and Compact Floorplan Reconstruction from Images. In Conference on Computer Vision and Pattern Recognition, 2014.

[2] Yu-Wei Chao, Wongun Choi, Caroline Pantofaru, and Silvio Savarese. Layout Estimation of Highly Cluttered Indoor Scenes Using Geometric and Semantic Cues. In Image Analysis and Processing, 2013. 2

[3] Jiacheng Chen, Chen Liu, Jiaye Wu, and Yasutaka Furukawa. Floor-SP: Inverse CAD for Floorplans by Sequential Room-Wise Shortest Path. In Conference on Computer Vision and Pattern Recognition, 2019.

Be consistent, it looks more professional:

- always full first names;
- always the same way to reference a conference (eg dont mix CVPR with Conference on Computer Vision..),

etc.

For me, the name of the conference and the year are enough.

École des Ponts ParisTech

the `\cite` command

the ~ is a non-breaking space

```
Earlier methods~\cite{Oikonomidis2011Tracking,
Kyriazis2014Scalable3T, Oikonomidis2012TrackingTA}
for generative hand pose estimation during
interaction used complex optimization methods to fit
a parametric hand model to RGBD data.
\cite{hampali2020honnotate} proposed multi-frame
op[illegible]on to fit hand and object models to RGBD
data from multiple RGBD cameras.
```

no ~ here

the introduction

École des Ponts
ParisTech

First impressions are very important.

After reading the introduction, the reviewer is likely to have already decided (maybe only unconsciously) if they will accept your paper or not

(if nothing is technically wrong in the method section and the results are convincing)

introduction structure at CVPR/ICCV/ECCV/..

1st paragraph: Why is the problem important?

2nd paragraph: Why is the problem difficult?

3rd paragraph: What is your key idea? You can point to the teaser figure.

4th paragraph: Give the important technical details.

5th paragraph: How do you evaluate your method? What are the most impressive results?

6th paragraph: A summary of your contribution (possibly as a bullet list).

7th paragraph: Maybe a short overview of the paper.

introduction structure at CVPR/ICCV/ECCV/..

1st paragraph: Why is the problem important?

don't discuss related work too much here.

2nd paragraph: Why is the problem difficult?

Give the intuition of your idea, get the reader hooked

3rd paragraph: What is your key idea? You can point to the teaser figure.

4th paragraph: Give the important technic

I often have to move aspects mentioned in the Method section here

5th paragraph: How do you evaluate your method? What are the most impressive results?

you can already explain here why you did not evaluate on some dataset the reviewer could think of.

6th paragra

7th paragra

you can also add a paragraph after the 5th paragraph to explain what your method is *not* if you think the reviewer can see relations with unrelated previous work

in general, try to anticipate what the reviewer will need to understand correctly your method

École des Ponts ParisTech

teaser

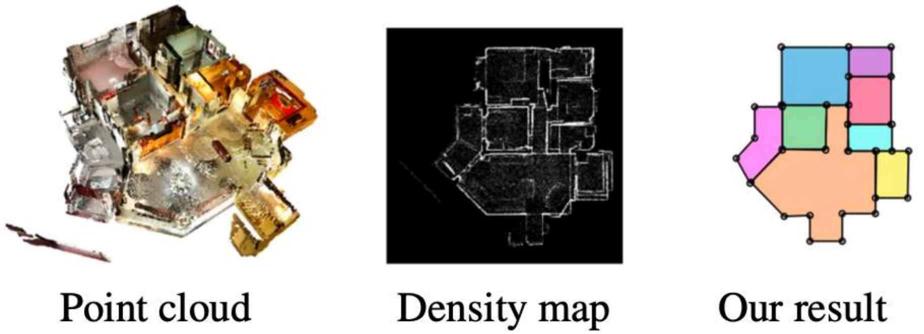

Figure 1: Given a density map, *i.e.*, the top view of the 3D point cloud of a floor, we retrieve an accurate floor map that successfully recovers variety of room shapes.

the reviewer should be able to understand the contribution of the paper from the teaser only

École des Ponts ParisTech

iterating on text [example from Bill Freeman]

An important part of being clear is being **concise**.

Original:

The underlying assumption of this work is that the estimate of a given node will only depend on nodes within a patch: this is a locality assumption imposed at the patch-level. This assumption can be justified in case of skin images since a pixel in one corner of the image is likely to have small effect on a different pixel far away from itself.

Revised:

We assume local influence--that nodes only depend on other nodes within a patch. This condition often holds for skin images, which have few long edges or structures.

the method section

École des Ponts
ParisTech

method section

Do NOT describe your method step by step!

You would quickly lose your reader in technical details.

Instead:

- Start with an overview of the section;
- then, give a general description of the method;
- end with the technical details.

→ Always from the more general to the more detailed explanation

for example

3.1 Learning a Regressor for Fixed Radius Structures

Let us momentarily assume that the linear struct
Let C be the set of centerline points and
Euclidean distance transform, that is, $\calD_C(\b
location $\bx$ to the closest location in C.

More general,
simpler problem

…

Second, a regressor trained to associate to a feature vector $f …e value of $d(\bx)$ can only do so approximately. As a result, th… refore no guarantee that its maximum is exactly on the centerline. To… obustness to noise, we have therefore found it effective to train our reg… reproduce a distance function whose extremum is better defined… ur actual implementation, we take it to be

…

More detailed

3.2 Handling Structures of Arbitrary Radius

In the previous section, we focused on structures of kn… neral, however, structures of many different radii are present. … lize our approach to this multi-scale situation, …

École des Ponts
ParisTech

the overview of the method section can be pretty long

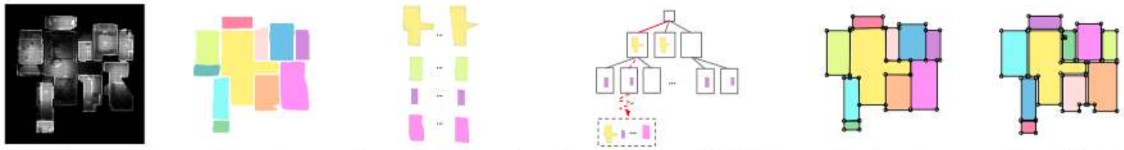

Figure 2: **Overview of our MonteFloor method.** Given a 3D point cloud, we first create a density map of a floor. We then detect room segments using Mask-RCNN as in Floor-SP [7]. Note the false positive at the bottom of the green segment on the left hand side. We polygonize each segment in different ways and obtain multiple room proposals from each room segment. We rely on MCTS and our objective function to select the correct room proposals, and our refinement step to adjust jointly the shapes of the room proposals to the input density map.

3. Method

Figure 2 gives an overview of our MonteFloor method: Given a 3D point cloud of a scene, we first create a top-view density map of this point cloud. We use Mask R-CNN [14] trained to detect rooms in such density maps and we polygonalize the detections to obtain a set of room proposals. Some proposals will correspond, at least coarsely, to actual rooms but others are only false positives. We use MCTS to find which room proposals make together the best fit to the input density map. The MCTS search is guided by a score predicted using a 'metric network' trained to predict the Intersection-over-Union between the selected room proposals and the floor map ground truth. Because the shapes of the correct room proposals from Mask R-CNN correspond only coarsely to the real rooms, we optimise their shapes while performing the search in MCTS. This is done by introducing a differentiable method for rendering polygonal shapes.

In the following, we detail:

- How exactly we obtain the room proposals;
- How we use MCTS to select the room proposals;
- Our objective function, involving our metric network and regularization terms;
- How we refine the room proposals' locations and shapes within MCTS;
- How exactly we compute the density map given a 3D point cloud.

École des Ponts ParisTech

notations

- Don't start the description of the method with a list of notations!
- Introduce the notations only when needed:

Given training samples $\{(f_i, y_i)\}_i$, where $f_i = f(\mathbf{x}_i, I_i) \in \mathbb{R}^M$ is the feature vector corresponding to a point $\mathbf{x}_i$ in image I_i and $y_i = d(\mathbf{x}_i)$, GradientBoost approximates $y(\cdot)$ by a function of the form

$$\varphi(f(\mathbf{x}, I)) = \sum_{k=1}^{K} \alpha_k h_k(f(\mathbf{x}, I)) \, , \tag{4}$$

where $h_k : \mathbb{R}^M \to \mathbb{R}$ are weak learners and $\alpha_k \in \mathbb{R}$ are weights. Function φ is built iteratively, selecting one weak

École des Ponts ParisTech

results Section

- As for other sections, start with an overview;
- Presents experiments that will show your approach is correct;
- You need comparisons with previous methods. If you are the first doing something, this is great but you still need a baseline.
- Don’t just give final quantitative results, try to give insights on your methods with additional experiments.

related work

Not a mere description of the state-of-the art.

Serves two purposes:

- show you know the state-of-the-art;
- show your method solves
  - aspects of the problem that were not solved before, or
  - a new problem.

École des Ponts ParisTech

2. Related Work

short introduction describing the structure of the section

Centerline detection methods can be classified into two main categories, those that use hand-designed filters and those that learn them from training data. We briefly review both kinds below.

Every section should start with a short overview of the section.

Hand-Designed Filters […]

The best is to write this section first and then its introduction.

École des Ponts ParisTech

2. Related Work

short introduction describing the structure of the section

Centerline detection methods can be classified into two main categories, those that use hand-designed filters and those that learn them from training data. We briefly review both kinds below.

short description of a family of methods

Hand-Designed Filters Such filters also fall into two main categories. The first is made of Hessian-based approaches [1, 2, 3] that combine the eigenvalues of the Hessian to estimate the probability that a pixel or voxel lies on a centerline. The main drawback of these approaches is that the required amount of Gaussian blur to compute the Hessian may result in confusion between adjacent structures, especially when they are thick.

[…]

- For each method, explain why or when they are not as good as your method, but be fair
- Be accurate, the authors are likely to be your reviewers!

École des Ponts ParisTech

conclusion

- avoid giving a summary of your paper;
- avoid proposing future work.

5. Conclusion

We have introduced an efficient regression-based approach to centerline detection, which we showed to outperform both methods based on hand-designed filters and classification-based approaches.

We believe our approach to be very general and applicable to other linear structure detection tasks when training data is available. For example, given a

if you think
- your approach can be applied to other problems, or
- points to new research directions, or
- is a milestone for a longer term direction

mention it and explain why.

École des Ponts ParisTech

abstract

- Write it last, you will save time and energy.

- Should be concise, but still have all the points to convey the idea:
  - Here is a problem
    - It is an interesting problem
    - It is an unsolved problem
  - Here is my idea
  - My idea works (details, data)
  - Here is how my idea compares to other people's approaches

École des Ponts ParisTech

We propose a robust and accurate method to extract the centerlines and scale of tubular structures in 2D images and 3D volumes. Existing techniques rely either on filters designed to respond to ideal cylindrical structures, which lose accuracy when the linear structures become very irregular, or on classification, which is inaccurate because locations on centerlines and locations immediately next to them are extremely difficult to distinguish.

We solve this problem by reformulating centerline detection in terms of a *regression* problem. We first train regressors to return the distances to the closest centerline in scale-space, and we apply them to the input images or volumes. The centerlines and the corresponding scale then correspond to the regressors local maxima, which can be easily identified. We show that our method outperforms state-of-the-art techniques for various 2D and 3D datasets.

École des Ponts
ParisTech

the input and the output
should be explicitly stated

don't lose time with generalities

We propose a robust and accurate method to extract the centerlines and scale of tubular structures in 2D images and 3D volumes. Existing techniques rely either on filters designed to respond to ideal cylindrical structures, which lose accuracy when the linear structures become very irregular, or on classification, which is inaccurate because locations on centerlines and locations immediately next to them are extremely difficult to distinguish.

important, yet unsolved problem

We solve this problem by reformulating centerline detection in terms of a *regression* problem. We first train regressors to return the distances to the closest centerline in scale-space, and we apply them to the input images or volumes. The centerlines and the corresponding scale then correspond to the regressors local maxima, which can be easily identified. We show that our method outperforms state-of-the-art techniques for various 2D and 3D datasets.

description of the contribution,
give the intuition but don't be vague

the proposed method
outperforms the state-of-the-art

École des Ponts
ParisTech

avoiding the writer's block

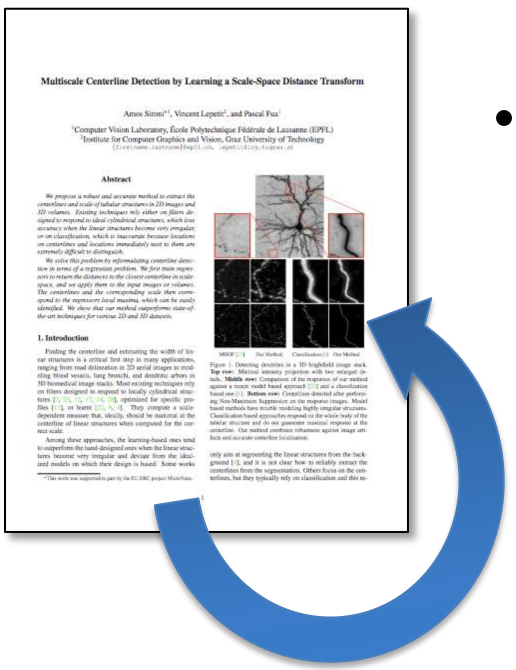

Multiscale Centerline Detection by Learning a Scale-Space Distance Transform

Amos Sironi*[1], Vincent Lepetit[2], and Pascal Fua[1]

[1]Computer Vision Laboratory, École Polytechnique Fédérale de Lausanne (EPFL)
[2]Institute for Computer Graphics and Vision, Graz University of Technology

Abstract

We propose a robust and accurate method to extract the centerlines and scale of tubular structures in 2D images and 3D volumes. Existing techniques rely either on filters designed to respond to ideal cylindrical structures, which lose accuracy when the linear structures become very irregular, or on classification, which is inaccurate because locations on centerlines and locations immediately next to them are extremely difficult to distinguish.

- start writing without thinking too much about the quality of your text;
- then iterate on your text, making it clearer and more convincing at each iteration:
  - Write the sections' overviews at the beginning of each section,
  - Make sure you followed the points from the first part of this talk;
  - make sure your paragraphs are short,
  - add figures,
  - etc.

École des Ponts
ParisTech

writing a paper with co-authors

working in parallel

PhD student

PhD advisor

start writing early (3 weeks before the deadline)

write as much as possible:
- you can use placeholders for the figures, but describe [caption and/or text] the figures anyway.
- you can have empty tables, but describe them anyway based on the results you expect.

make sure you followed the points from the first part of the talk and tell your advisor they can look at your text

you can do more experiments in parallel

does their job

keeping tracks of the changes and leaving remarks

We use commands of the form:

```
\newcommand{\vincent}[1]{{\color{ForestGreen} #1}}
\newcommand{\vincentrmk}[1]{{\color{ForestGreen} {\bf VL: #1}}}
```

to get:

As shown in Fig. 3, our architecture first detects keypoints that are likely to correspond to the 2D locations of hand joints and encodes them as input to the keypoint-joint association stage. The keypoints are encoded with their spatial locations and the image features at these locations. The self-attention layers in the Keypoint Transformer disambiguate the keypoints and associates them with different joint types and a background class. **v: I did not get this point:** The (single) cross-attention layer then selects these “identity-aware keypoints” to predict root-joint-relative pose parameters of both hands, plus additional parameters such as the translation between the hands and hand shape parameters.

École des Ponts
ParisTech

organizing Overleaf

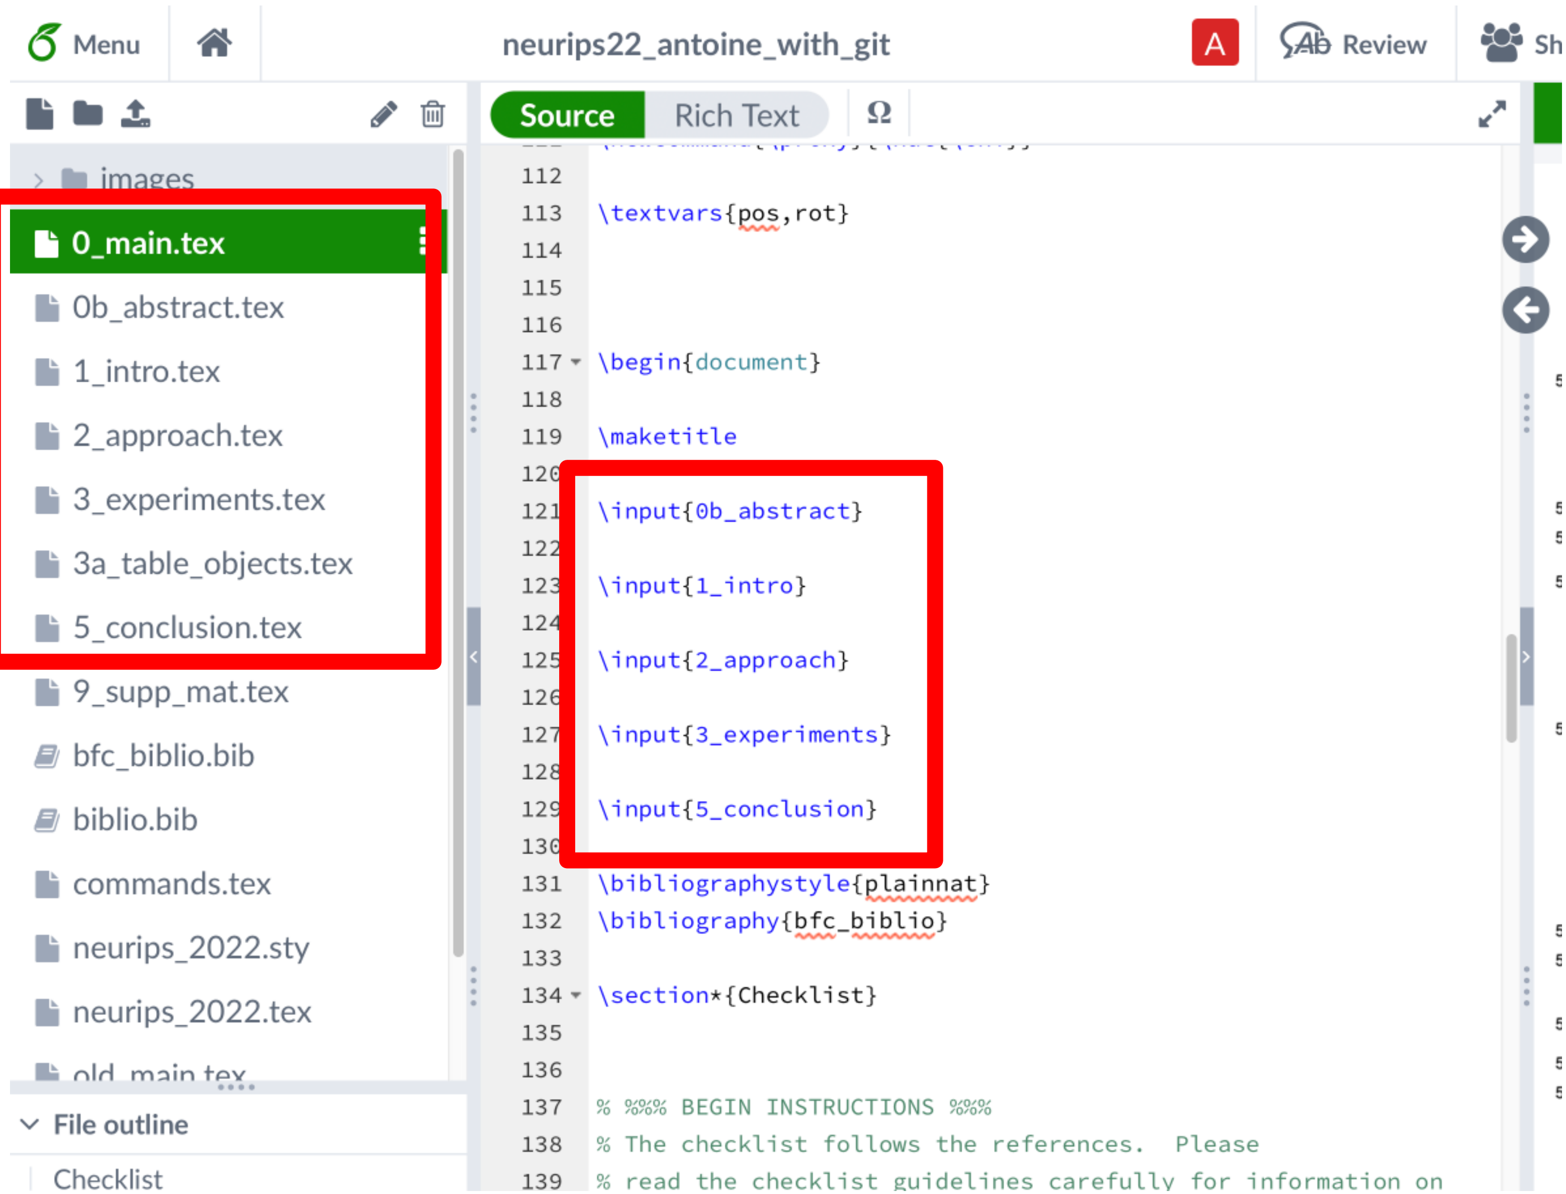

my LaTeX commands

you can get the file at: https://vincentlepetit.github.io/files/vincents_commands.tex

```
% \addeditor{vincent}{VL}{0.0, 0.5, 0.0}
%
% adds the following commands:
%
% \vincent{text},
% \vincentrmk{remark}, and
% \vincentrpl{newtext}{oldtext}
%
% Use \showeditstrue to show the edits, and
%     \showeditsfalse for a clean version.
%
%
% \textvars{pose,rot}
% adds the following commands:
% \pose, which is replaced by \text{pose} and
% \rot, which is replaced by \text{rot}
%
%
% \calA for \mathcal{A}, etc.
% \bA for \textbf{A}, etc.
% \ba for \textbf{a}, etc.
% \IR for \mathds{R}, etc.
```

École
ParisTech

conclusion

Should we start writing as soon as we have the first results?

Remember to write as much text as possible, even if you are still waiting for the final results (use placeholders);

Remember the advices for saving iterations.

These are not rigid rules, it takes a lifetime to learn how to write well.

École des Ponts
ParisTech