
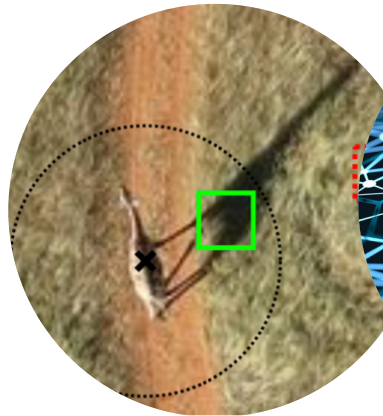


Geospatial computer vision

Let's talk about deep learning for remote sensing

AMID

Devis Tuia, Wageningen University, Jan 27





Let's talk about something new!

Let's dialogue with images!

AMID

Devis Tuia, Wageningen University, Jan 27



Dialoguing with images... why?

- We want to make remote sensing more accessible to **non-experts**
- Non-experts are ... **non technical** experts
- Non-experts want answers to **specific questions**

We are good at solving single tasks




We are good at solving single tasks



How many cars?

Deep Neural Network



Update parameters

11

Loss function

Ground Truth



12

[Lobry and Tuia, JURSE 2019; Lang et al., LPS 2019]

We are good at solving single tasks



Are there trees?



[Audebert et al., Maggiori et al.; Volpi and Tuia; ...]

Dialoguing with images... why?

- We want to make remote sensing more accessible to **non-experts**
- Non-experts are ... **non technical** experts
- Non-experts want answers to **specific questions**
you don't know in advance what's coming

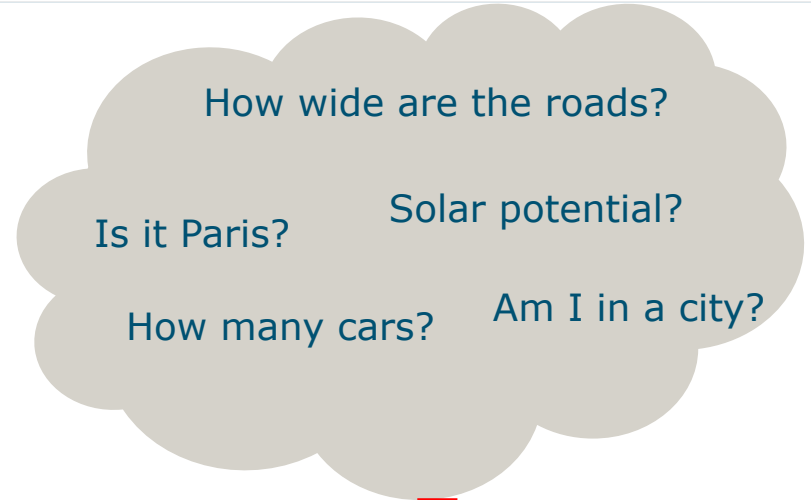
Dialoguing with images... why?

- We want to make remote sensing more accessible to **non-experts**
- Non-experts are ... **non technical** experts

- Non-experts want answers to **specific questions**
you don't know in advance what's coming

- Non-experts want to formulate questions as **sentences**
in english, not in convolutions

We are not very good at reacting to unforeseen questions



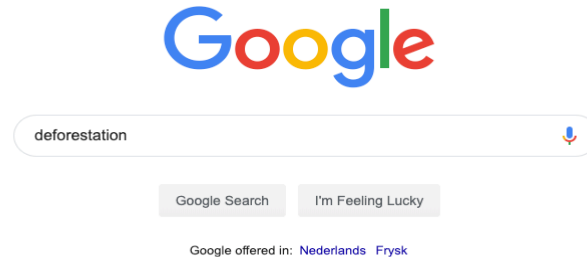
ANSWER



What do we need?

- For web-search it works a bit like that.

[Gmail](#) [Images](#)  [Sign in](#)



What do we need?

- For web-search it works a bit like that.

The image shows a Google search interface for the term "deforestation". The search bar contains the text "deforestation" and a search icon. Below the search bar, there are navigation tabs for "All", "Images", "News", "Videos", "Maps", and "More". The search results show "About 23,700,000 results (0,51 seconds)".


The first result is a featured snippet from Live Science, dated April 3, 2018. It includes a large image of a deforested landscape with tree stumps. The text reads: "Deforestation is the permanent destruction of forests in order to make the land available for other uses. An estimated 18 million acres (7.3 million hectares) of forest, which is roughly the size of the country of Panama, are lost each year, according to the United Nations' Food and Agriculture Organization (FAO)." Below the text is a link to the article: "Deforestation: Facts, Causes & Effects - Live Science" with the URL "https://www.livescience.com/27692-deforestation.html".

To the right of the featured snippet is a grid of six smaller images related to deforestation, with a "More images" button. Below this grid is a "Deforestation" card with a share icon and a brief definition: "Deforestation, clearance, clearcutting or clearing is the removal of a forest or stand of trees from land which is then converted to a non-forest use. Deforestation can involve conversion of forest land to farms, ranches, or urban use. The most concentrated deforestation occurs in tropical rainforests. Wikipedia".

Below the featured snippet is a "Top stories" section with three video thumbnails. The first is titled "Brazil registers huge" and shows a dirt road cutting through a forest. The second is titled "Brazil: huge rise in" and shows a deforested area with a "The Guardian" logo. The third is titled "Sham rise in Amazon" and shows a forest with a play button icon.

What do we need?


- For web-search it works a bit like that.
- With satellite images it just doesn't work
(it's normal. It wasn't built for that)



But what if you could... ask questions to remote sensing images?

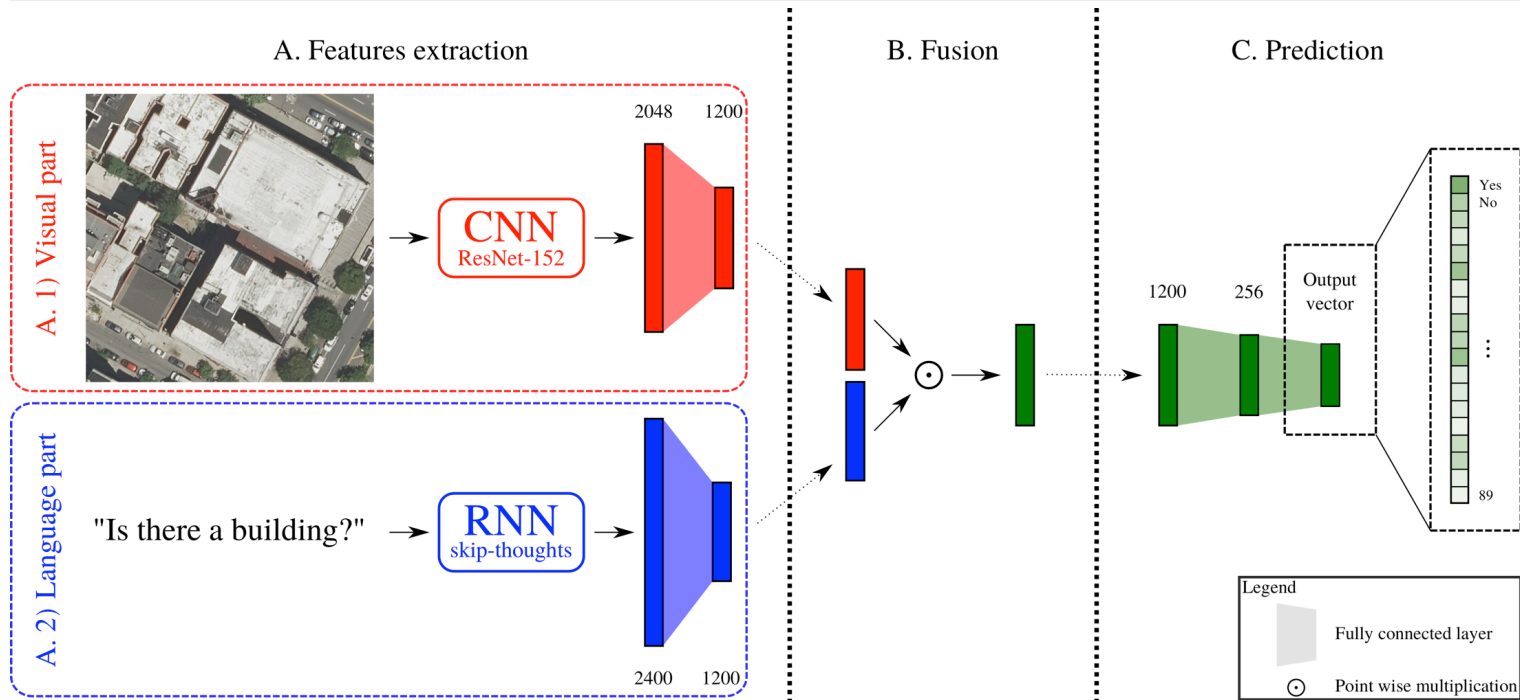


Remote sensing visual question answering (RSVQA)



[Lobry, Marcos, Murray, Tuia, IEEE TGRS]


Remote sensing visual question answering (RSVQA)



How do we train this monster?

We created a dataset of

- Sentinel-2 images (RGB)
 - 9 scenes
 - 772 tiles (256 x 256)
- OpenStreetMap layers
- Covers the whole Netherlands



How do we train this monster?



We generated **77'232** {image, **question**, answer} triplets

Check on OSM


Elements catalog

Road
Water area
Commercial building
Industrial building


Residential area
Retail
⋮
Religious area

} Land usages

Attributes catalog




Relations catalog



Random generator

Questions catalog




How do we train this monster?




We generated **77'232** {image, **question**, answer} triplets


"How many roads are present in the image?"



"Is there a small retail place?"



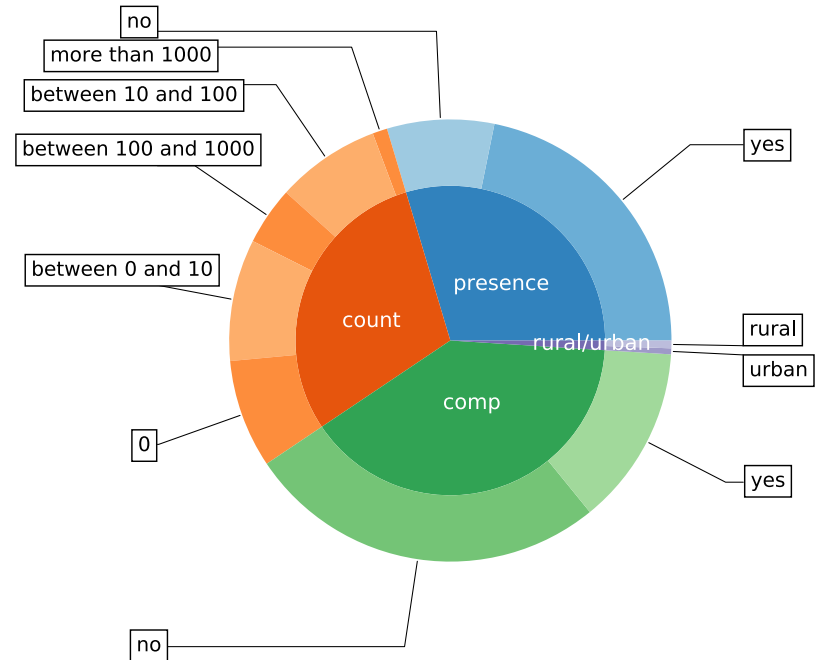
"Is there more buildings at the top of a circular religious place than roads in the image?"



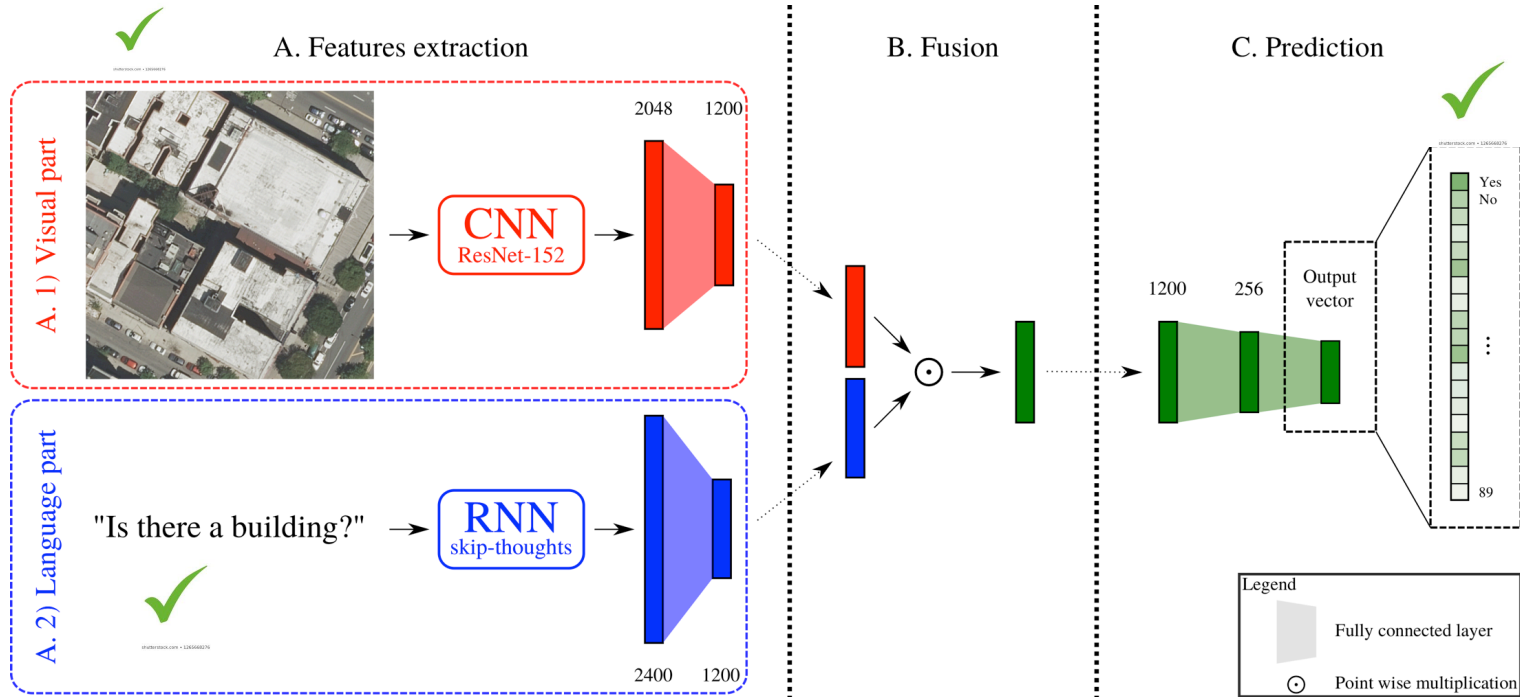
How do we train this monster?



We generated **77'232** {image, question, **answer**} triplets



Remote sensing visual question answering (RSVQA)



Results – Sentinel 2

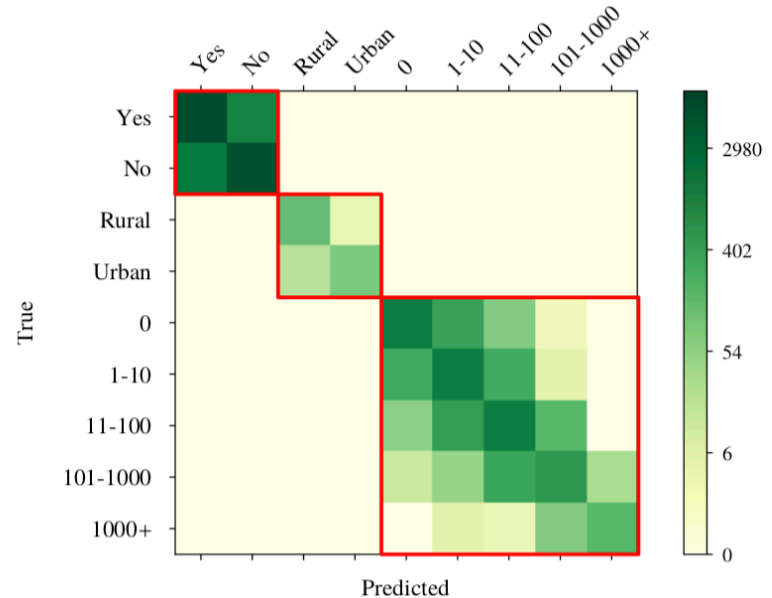
79% overall accuracy!

73% if randomizing the image part

Count questions less accurate

Type	Accuracy
Count	67.01% (0.59%)
Presence	87.46% (0.06%)
Comparison	81.50% (0.03%)
Rural/Urban	90.00% (1.41%)
AA	81.49% (0.49%)
OA	79.08% (0.20%)

The model can make a good distinction between types of questions



Results – Sentinel 2



Is it a rural or an urban area?

Ground truth	Prediction
Rural	Rural



Are there less buildings than water areas?

Ground truth	Prediction
No	No

Results – Sentinel 2



Is it a rural or an urban area?

Ground truth

Urban

Prediction

Urban



Are there more water areas than commercial buildings?


Ground truth

Yes

Prediction

No

Results – Sentinel 2



Are there more water areas than commercial buildings?

Ground truth

Yes

Prediction

No



Are there more water areas than commercial buildings?


Ground truth

Yes

Prediction

No

Results – Sentinel 2



Are there more water areas than commercial buildings?

Ground truth

Yes

Prediction

No



Are there more water areas than commercial buildings?

Ground truth

Yes

Prediction

No

Results – Sentinel 2



Are there more water areas than commercial buildings?	
Ground truth	Prediction
Yes	No



Are there more water areas than commercial buildings?	
Ground truth	Prediction
Yes	No




Where is this going?

2020





Where is it?





Where is this going?

2020



Where is it?

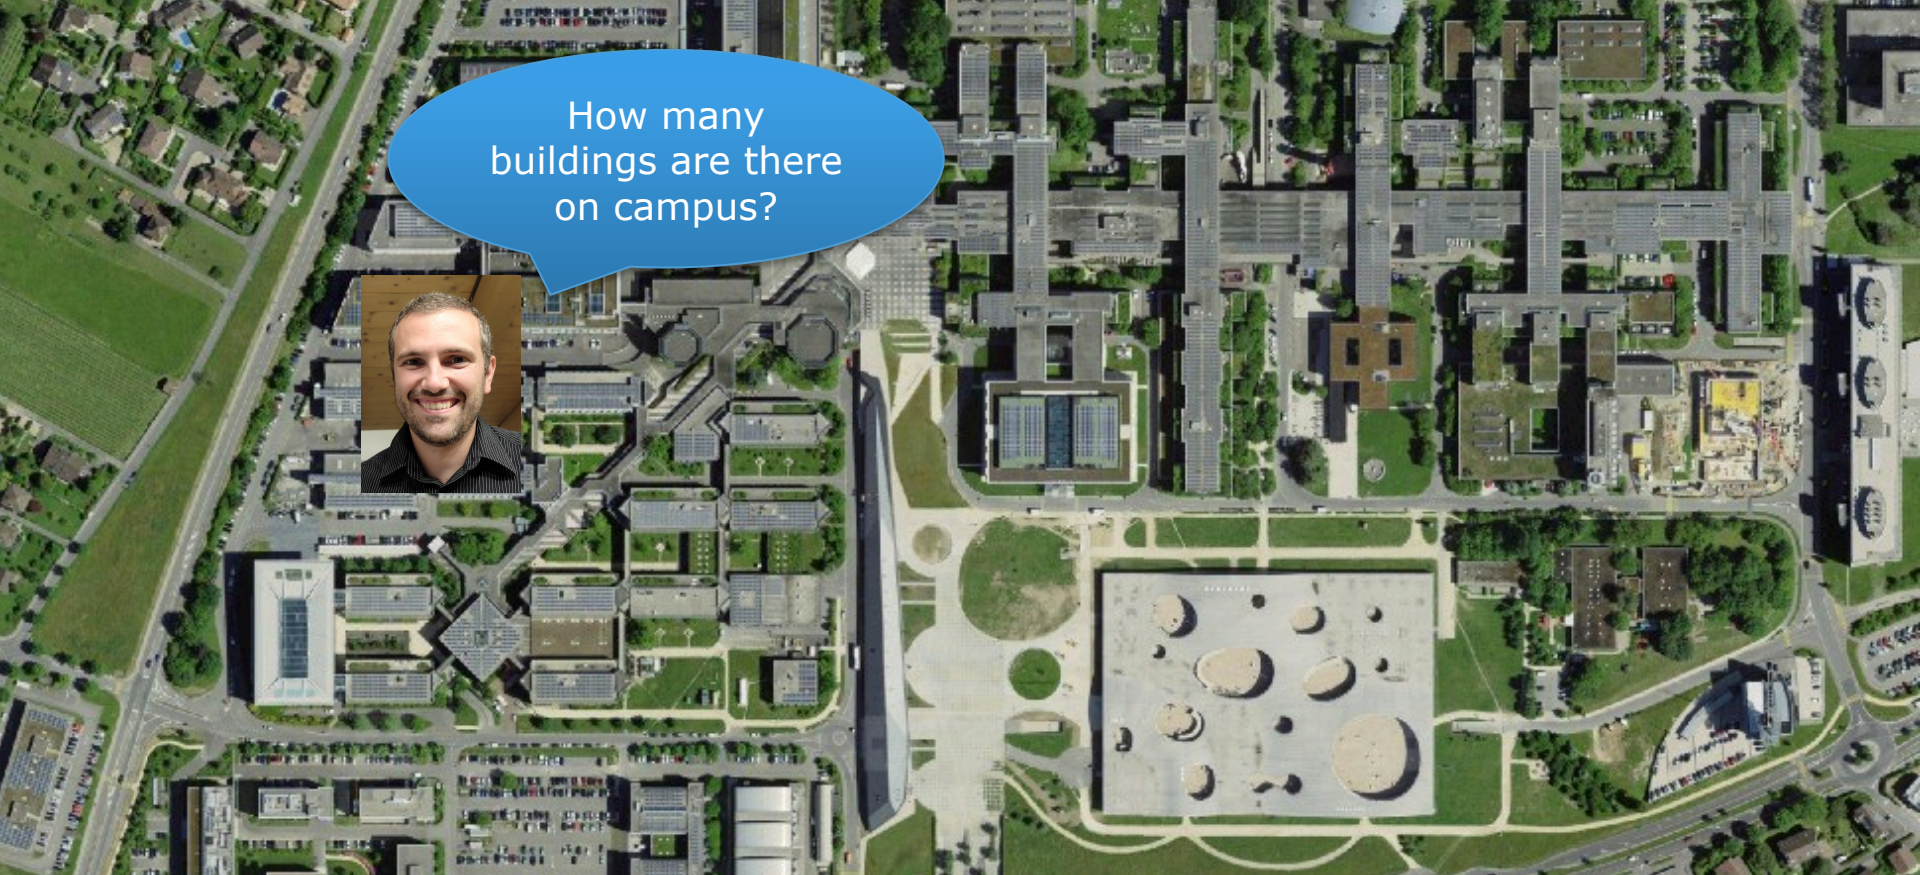
It is EPFL,
Lausanne




Where is this going?

2020



An aerial photograph of a university campus. A blue speech bubble is overlaid on the top left, containing the text 'How many buildings are there on campus?'. Below the speech bubble is a small portrait of a smiling man with short dark hair and a beard, wearing a dark shirt. The campus features numerous buildings, green spaces, and a large central area with several circular structures.

How many buildings are there on campus?





Where is this going?

2020



How many buildings are there on campus?

57





Where is this going?

2020



Which ones are new?

mmm... let me check...




Where is this going?

2020



Where is this going?

2007







Where is this going?

2020





These.





Where is this going?

2020



These.


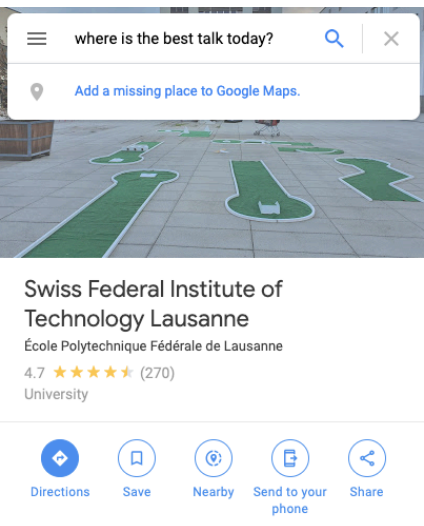


Where is this going?


2007

RSVQA - Summary

- Joins image recognition and natural language processing deep models
- Opens use of EO image data to the laymen
- Towards an EO search engine



Thanks!



devis.tuia.googlepages.com (with links to codes!)

devis.tuia@wur.nl

Coming soon:

The book

“Deep learning for the Earth Sciences”

A comprehensive approach to remote sensing, climate science and geosciences

Editors: Gustau Camps-Valls, Xiao Xiang Zhu, Devis Tuia, Markus Reichstein

Publisher: Wiley & Sons, inc., 2020

e l l i s

European Laboratory for Learning and Intelligent Systems

Program “ML for Earth and Climate Science”