# Mining High Resolution Earth Observation Data Cubes


Andreas Züfle
George Mason University
USA
azufle@gmu.edu

Konrad Wessels
George Mason University
USA
kwessel4@gmu.edu

Dieter Pfoser
George Mason University
USA
dpfoser@gmu.edu



## ABSTRACT

Earth observation data is collected by ever-expanding fleets of satellites including Landsat1-8, Sentinel1 & Sentinel2, SPOT1-7 and WorldView1-3. These satellites generate at spatial resolutions (pixel size) from $30m$ to $31cm$ and provide revisit rates of as frequent as every 5 days. This allows us not only to look at high-resolution images of every corner of the Earth, but also to track events and observe change over time. During the past 5 years, medium spatial resolution satellite data ($30 - 10m$ pixels) have developed very high temporal revisit frequencies of 5-16 days and spatial-temporal structures have been developed to manage these vast data sets. However, high resolution satellite images and rapidly increasing revisit rates create major data management and mining challenges. This work discusses six challenges of integrating observations at different times, from different sensors, at different spatial resolutions and different temporal frequencies into a unified Earth Observation Data Cube, that is, a tensor of location, time, and spectral bands. Challenges include creating a unified data cube from heterogeneous sensors, scaling geo-registration (mapping pixel between images), accounting for uncertainty across observations, imputing missing observations, broad area event detection, and ultimately, predicting the future state of our planet. With such a unified Earth Observation Data Cube in place, we describe potential application areas such as detecting anthropogenic land cover change, early warning of natural hazards, tracing movement of animals, finding missing airplanes, and rapid detection of forest fires.


## CCS CONCEPTS

• **Information systems** → **Spatial-temporal systems**; **Location based services**.


**ACM Reference Format:**

Andreas Züfle, Konrad Wessels, and Dieter Pfoser. 2021. Mining High Resolution Earth Observation Data Cubes. In *17th International Symposium on Spatial and Temporal Databases (SSTD '21), August 23–25, 2021, virtual, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3469830.3470917


## 1 INTRODUCTION

Due to advances in imaging and satellite technology, an overwhelming amount of Earth Observation (EO) data is collected and made publicly available at an unprecedented spatial and temporal resolution. For example, the Worldview-3 satellite observes the world at a *resolution of* $31cm$ *per pixel* [1], which translates into 10.4 million pixels per $km^2$, and covers $680,000km^2$ a day resulting in more than 7 trillion pixels per day. Data has been collected since 2014 and is openly available to paying customers [7]. Other programs such as Landsat [27] and Sentinel-2 [8] now coordinate up to four satellites to produce a Harmonized Landsat Sentinel2 product, which will reduce the traditional 16-day revisit frequency of a single Landsat down to 3 days and creating a flood (or rather a Tsunami) of open and publicly available data in the process.

Besides frequent satellite imagery, other spatio-temporal environmental data have become publicly available. For example, MERRA-2 Modern-Era Retrospective analysis for Research and Applications Dataset [10, 24] provides environmental parameters such as temperature, humidity, and precipitation every hour since 1979, but with a coarser spatial resolution of $50km$. MERRA-2 datacubes have been used, for example, to predict the rapid intensification of tropical storms [17].

Together, satellite observations of land, oceans and atmosphere along with models, can capture and monitor the entire earth systems, up to the scale of the entire Earth. Abstractly speaking, Earth observation data maps location and time to a series of observations, such as spectral features (color intensities), temperature, or precipitation.

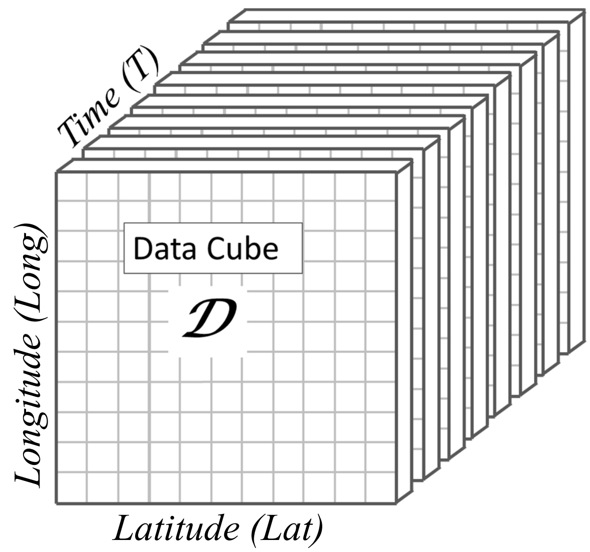


**Figure 1: Schematic of an Earth observation data cube**

[1] The panchromatic (greyscale) sensor resolution of Worldview-3 is at $31cm$ whereas the multi-spectral (color) sensor has a resolution of $1.24m$.



*SSTD '21, August 23–25, 2021, virtual, USA*



ACM ISBN 978-1-4503-8425-4/21/08…$15.00
https://doi.org/10.1145/3469830.3470917

*Definition 1.1 (Observation).* Let $\mathcal{V} = \{var_1, ..., var_N\}$ denote a set of observable variables. We let $O = var_1 \times ... \times var_N$ denote the observation variable space and we let an observation $obs \in O$ map each variable to a value.

Earth observation data consists of a geolocation, a time (often aggregated to days), and an observation. Earth observation datasets can be captured by so-called data cubes [20].

*Definition 1.2 (Earth Observation Datacube).* Let $G = Lat \times Long$ denote a spatial grid, let $Time$ denote a time grid, and let $O = var_1 \times ... \times var_N$ denote an observation variable space. An Earth Observation Datacube $\mathcal{D} := O^{\mathcal{D} \in lat \times long \times Time}$ is a three-mode tensor such that $\mathcal{D}_{lat \in Lat, long \in Long, t \in Time} \in O$ is an observation made at location $(lat, long)$ at time $t$.

A visual representation of an Earth observation data cube is depicted in Figure 1. We note that an Earth observation data cube may be extremely large. Using $31cm$ resolution observations as measured by Worldview-3, the planet spans 2.05 quadrillion pixels, yielding approximately a $|Lat| = 45,000,000$ by $|Long| = 45,000,000$ grid measured for approximately $|Time| = 3000$ days since the start of Worldview-3 in 2013. Thus, for Worldview-3, the corresponding Datacube $\mathcal{D}$ would have more than six quintillion ($6 \cdot 10^{18}$) cells.

## 2 TOWARDS A UNIFIED EARTH OBSERVATION DATA CUBE

Although significant progress has been made to create datacube standards and operational implementations (e.g., Open DataCube) [6, 11, 12], a number of challenges have to be solved to create a unified Earth observation data cube that captures data from different satellites, at different spatio-temporal resolutions, and with different levels of uncertainty. Prior work [20] defines Earth observation as a collection of data cubes, each pertaining to an individual sensor. While such an approach is sufficient to analyze individual scenes and to compare different sensors, it does not allow linking data across data cubes. To obtain a greater picture (technically, a greater video) of Earth, it is paramount to unify or link data cubes from different sensors to provide unified access to observations from a multitude of sensors.

Challenge 1 (A Unified Earth Observation Data Cube). *Given a collection of Earth observation data cubes. The challenge is to unify the information across all the data cubes into a single data cube $\mathcal{D}$.*

Such sensors may not only include large Earth observation projects using state-of-the-art satellite technology, but also include data observed and volunteered by individuals using drones to observe small areas that may span only a handful of square kilometers. We call such data *Volunteered Earth Observation Data* (VEOD). The greater challenge will be to integrate such VEOD in a single Earth observation data cube. Such integration may enabling an era of volunteered Earth observation data where individual users can upload Earth observation data captured by private drones in a common and public Earth observation data cube.

Towards Challenge 1, Section 2.1 describes the challenge of co-registering large-scale Earth observation data cubes, that is, mapping the precise locations observed in one observation to the same location in another observation. Section 2.2 discusses the challenge of estimating the uncertainty of observations of individual sensors and the challenge of minimizing uncertainty across sensors in a unified data cube. Section 2.3 postulates the challenge of handling the sparsity of an Earth observation data cube by imputing missing observations.

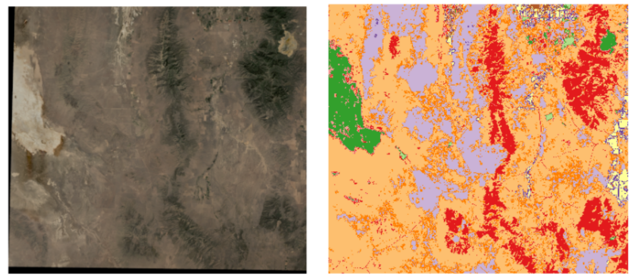

(a) Earth Observation Data (b) Co-Registered Land Cover Data

**Figure 2: Example of Image Co-Registration**

### 2.1 Earth Observation Data Cube Fusion

In remote sensing, multi-sensor image fusion or (co-)registration is the process of combining relevant information from two or more images into a single image [4]. The remote sensing community has already developed algorithms for this purpose [19]. For example, Figure 2(a) shows Harmonized Landsat Sentinel-2 (HLS) data obtained via NASA's Land Processes Distributed Active Archive Center (LP DAAC) [1]. To obtain a ground truth of labeled land cover information (such as forest cover, wetlands, urban, etc.), Figure 2(b) shows co-registered land cover information obtained from the United States Geological Survey (USGS) National Land Cover Database [14]. In a nutshell, co-registration joins two observations using a matching spatial location as the join predicate. Existing solutions for co-registration however assume that observed images capture the same, known area of limited size.

Challenge 2 (Scalable Earth Observation Data Cube Fusion). *An open challenge is to scale co-registration algorithms to a global scale to allow the fusion of whole data cubes, which may each pertain to quadrillions of observations across space and time, at different resolutions and levels of data quality.*

### 2.2 Uncertainty Management

Different earth observation sensors capture the planet at different resolutions and different levels of quality. Drones may capture a small area of the planet with extremely high centimeter resolution, some satellites may observe data at $10m$ resolution while others may only provide $30m$ resolution data. How can we integrate all such observations across data cubes without loss of information? Simply taking the highest resolution observation will discard information from lower resolutions which could be used to correct errors. As observations may be made at different times, it becomes a challenge when to switch to a more recent but lower resolution observation. Some observations may also be covered, partially or completely, by clouds or shadows, thus blurring or masking the true observations on the ground. Another challenge is that at a given time, some observations may be obsolete and require extrapolation from the

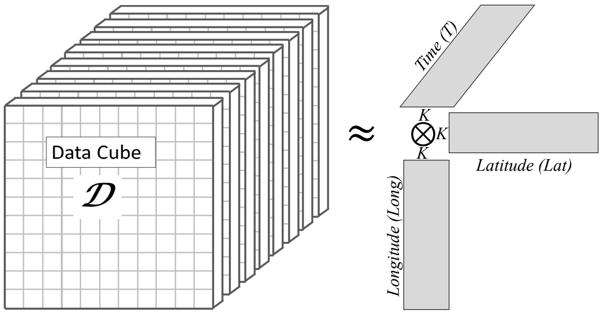


**Figure 3: Schematic of an Earth observation data cube**

most recent observations or interpolation using observations made at later times.

All these challenges have in common that for a location on Earth, our information may be uncertain, obsolete, inconsistent, or contradicting. How can we capture the underlying uncertainty and leverage uncertainty for more reliable information retrieval and data mining?

Challenge 3 (Managing Uncertainty in Earth Observation Data Cubes). *Existing Earth Observations Data Cubes store deterministic values. Due to uncertain, obsolete, inconsistent, or contradicting observations, the true data values are unknown. How can we capture the inherent uncertainty of observations to improve decision making?*

To tackle this challenge, researchers may be able to leverage advances in uncertain spatio-temporal data management as surveyed in recent tutorials [30–32]. Uncertain spatio-temporal databases treat spatial information as a random variable whose probability distribution is estimated using observed data. While the size of an earth observation data cube may make detailed probability distributions infeasible, parametric models may be useful to assess the variance of each observation as a measure of reliability.

## 2.3 Imputation of Missing Values

While an Earth Observation Data Cube as defined in Definition 1.2 is extremely large, it is also very sparse. For most $(lat, long, time)$ triples there are typically no available observations in the corresponding location and on the corresponding day. While satellites like Worldview-3 collect data from vast areas of the Earth each day (680,000$km^2$ for the case of Worldview-3) this is less than 0.2% of the total area of 512,072,000$km^2$ of Earth.

Challenge 4 (Earth Observation Imputation). *Given a sparse Earth observation data cube, we formulate the challenge of estimating what an observation, in the past, would have looked like if a sensor would have observed it.*

A classic approach for tensor imputation uses tensor factorization to represent each mode of a tensor by a small set of $K$ latent features [29] as sketched in Figure 3. Following an encoder-decoder paradigm, this encoded representation of the tensor is then decoded (via tensor multiplication) to obtain a data cube having missing values estimated. While such an approach works well to impute missing values in non-spatial applications such as recommender systems (where a tensor may hold user ratings and modes may correspond to users, products, and context features) (cf. , e.g., [22]), it is not clear whether such an approach is promising for spatial data. The problem is that the factorized representation of $K$-latitude features and $K$-longitude features (as shown on the right of Figure 3) treats latitude and longitude as independent modes, without any notion of spatial proximity. A spatially-aware Earth observation data cube factorization that employs a convolution of space may be a promising direction to improve the imputation of missing Earth observations.

# 3 MINING A UNIFIED EARTH OBSERVATION DATA CUBE

Having a unified Earth observation data cube at hand will enable impactful new applications and research directions with major impact on our lives. Speaking abstractly, an Earth observation data cube represents a (very large) video of Earth. A first vision of broad area search is to detect and retrieve events of interest, such as the location (in space and time) of a missing plane, the location and movement of protected animals, or an emerging forest fire. A second vision asks how our planet evolves into the future [2]. How is climate changing? Which places will suffer from pollution? How will land cover change in the next decades? This section describes these two visions in details.

## 3.1 Broad Spatio-Temporal Area Search

Our ability to extra information and recognize objects on images and videos has improved tremendously in the last decade, allowing us to detect objects as videos are streamed [23]. Such solutions can be used to monitor a region to detect flooding, or forest fires [18] using high resolution imagery and videos collected by UAVs. A limiting factor for such approaches is that one first has to know where to look to find objects of interest. What if a forest fire breaks out in an area that is not monitored?

To address this problem, we can utilize EOD and monitor the whole planet at a lower spatial and temporal resolution. Existing object detection algorithms [23, 26] first scan an image or video for candidate objects and then use a classifier to identify the type of object. Earth-wide data cubes contain many objects and the vast majority of them represents negative examples. A novel challenge is to develop object detection algorithms capable of processing data on the scale of the whole planet.

Challenge 5 (Earth-Wide Object or Event Detection). *Given an earth observation data cube, a challenge is to detect objects or events of interest such as forest fires, flooding, and other natural or anthropogenic events automatically, and on a planet-wide scale without a priori knowledge of the location of these events.*

Speaking figuratively, the challenge is to find needles in a haystack, or "Waldo" in a crowd of people [13]. We have capabilities to classify needles and hay straws and if the image of a person depicts Waldo. However, efficiently searching through an Earth-wide multi-year data cube discarding potentially billions of true negatives remains an open challenge. This challenge may require spatial index structures and efficient algorithms to quickly identify true negatives with high confidence, and to leverage parallel and distributed computing solutions for spatial data [9, 28] to search concurrently in many

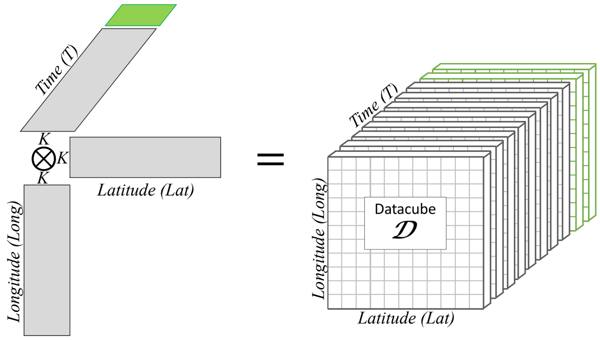


**Figure 4: Predicting future states of an Earth observation data cube. Left: Factorized representation to an EO data cube with predicted future temporal features. Right: Corresponding decoded datacube with future slices extrapolated via matrix multiplication.**

places at once. Currently, operational systems are available that detect, e.g., deforestation in the tropics using time series of Landsat imagery, or active fires using the thermal channels of MODIS satellite sensor. However, these individual systems are tailored to identifying only one specific phenomenon, or a finite study area using one or maybe two satellite sensors. The vision of the unified Data Cube would enable a "God's Eye view", which is able to automatically detect multiple phenomena using a large number of diverse satellite sensors. Artificial Intelligence will play a central role in deriving situational awareness and geo-spatial insight from all the earth observation data.

### 3.2 Predicting Future Earth

We have gained great capabilities in predicting future states of complex spatio-temporal systems such as traffic [5, 16], crime [25], and events [21]. A novel challenge is to leverage these predictive capabilities to predict future states of the Earth. How will land cover change in the next years, given specific drivers? How will urban areas grow under various economic and policy scenarios (and pandemics)? Will land degradation such as desertification continue? By forecasting future states of an unified Earth observation data cube, we may be able to answer such questions.

Challenge 6 (Earth Observation Prediction). *Given an Earth observation data cube* $\mathcal{D} := O^{\mathcal{D}\in lat\times long\times Time}$*, capturing Earth observation at past and current times* $Time$ *a challenge is to predict future states* $\mathcal{D} := O^{\mathcal{D}\in lat\times long\times Time'}$ *where* $Time'$ *captures future times that have not yet been observed.*

One approach to estimate future Earth observations may leverage a encoder-decoder framework using non-negative matrix factorization as used in recent work to predict states in a traffic network [3]. Such an approach may predict latent features of time in the encoded latent feature space and decodes these features using matrix multiplication to predict future observations. Other approaches may leverage GeoAI to find representations of space and time to predict future observation of the data cube [15].

## 4 CONCLUSIONS AND FUTURE WORK

Data cubes are important data management and analysis constructs when it comes to extracting knowledge from the massive Earth Observation Data that is being collected by a steeply increasing number of satellites. This work outlines some of the challenges when in comes to integrating data at varying spatial an temporal granularities and coming from different sensors. While there are no prescribed solutions, integrating such data in one large data cube or linking across numbers of smaller data cubes, it is paramount to provide unified access to such data resources. We mention six specific challenges that relate implementation, i.e., integration, scaling, uncertainty management, and missing values, as well application-type problems, i.e., "needle-in-haystack" broad area search, and prediction of future states (of the Earth). These challenges should provide a blueprint for geospatial data science research as it relates to Earth Observation Data.

## REFERENCES

[1] U.S. Geological Survey (USGS) & National Aeronautics and Space Administration (NASA). Accessed 07/15/2021. Land Processes Distributed Active Archive Center (LP DAAC). https://lpdaac.usgs.gov. (Accessed 07/15/2021).
[2] S Aigner and M Körner. 2019. Futuregan: Anticipating the Future Frames of Video Sequences Using Spatio-Temporal 3d Convolutions in Progressively Growing Gans. *ISPRS Photogrammetry, Remote Sensing and Spatial Information Sciences* 4216 (2019), 3–11.
[3] Abdelkader Baggag, Sofiane Abbar, Ankit Sharma, Tahar Zanouda, Abdulaziz Al-Homaid, Abhiraj Mohan, and Jaideep Srivasatava. 2019. Learning Spatiotemporal Latent Factors of Traffic via Regularized Tensor Factorization: Imputing Missing Values and Forecasting. *IEEE Transactions on Knowledge and Data Engineering* (2019).
[4] Chen Chen, Yeqing Li, Wei Liu, and Junzhou Huang. 2015. SIRF: Simultaneous satellite image registration and fusion in a unified framework. *IEEE Trans. on Image Processing* 24, 11 (2015), 4213–4224.
[5] Fanglan Chen, Zhiqian Chen, Subhodip Biswas, Shuo Lei, Naren Ramakrishnan, and Chang-Tien Lu. 2020. Graph Convolutional Networks with Kalman Filtering for Traffic Prediction. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*. 135–138.
[6] Trevor Dhu, Gregory Giuliani, Jimena Juarez, Argyro Kavvada, Brian Killough, Paloma Merodio, Stuart Minchin, and Steven Ramage. 2019. National Open Data Cubes and Their Contribution to Country-Level Development Policies and Practices. *DATA* 4, 4 (DEC 2019). https://doi.org/{10.3390/data4040144}
[7] Digital Globe. 2013. WorldView-3 Data Sheet. http://content.satimagingcorp.com.s3.amazonaws.com/media/pdf/WorldView-3-PDF-Download.pdf. (2013).
[8] Matthias Drusch, Umberto Del Bello, Sébastien Carlier, Olivier Colin, Veronica Fernandez, Ferran Gascon, Bianca Hoersch, Claudia Isola, Paolo Laberinti, Philippe Martimort, et al. 2012. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote sensing of Environment* 120 (2012), 25–36.
[9] Ahmed Eldawy and Mohamed F Mokbel. 2015. Spatialhadoop: A mapreduce framework for spatial data. In *2015 IEEE 31st international conference on Data Engineering*. IEEE, 1352–1363.
[10] Ronald Gelaro, Will McCarty, Max J Suárez, Ricardo Todling, Andrea Molod, Lawrence Takacs, Cynthia A Randles, Anton Darmenov, Michael G Bosilovich, Rolf Reichle, et al. 2017. The modern-era retrospective analysis for research and applications, version 2 (MERRA-2). *Journal of climate* 30, 14 (2017), 5419–5454.
[11] Gregory Giuliani, Gilberto Camara, Brian Killough, and Stuart Minchin. 2019. Earth observation open science: Enhancing reproducible science using data cubes. (2019).
[12] Gregory Giuliani, Gilberto Camara, Brian Killough, and Stuart Minchin. 2019. Earth Observation Open Science: Enhancing Reproducible Science Using Data Cubes. *DATA* 4, 4 (DEC 2019). https://doi.org/{10.3390/data4040147}
[13] Martin Handford. 2019. *Where's Waldo?* Where's Waldo.
[14] Collin Homer, Jon Dewitz, Limin Yang, Suming Jin, Patrick Danielson, George Xian, John Coulston, Nathaniel Herold, James Wickham, and Kevin Megown. 2015. Completion of the 2011 National Land Cover Database for the conterminous United States–representing a decade of land cover change information. *Photogrammetric Engineering & Remote Sensing* 81, 5 (2015), 345–354.
[15] Yingjie Hu, Song Gao, Dalton Lunga, Wenwen Li, Shawn Newsam, and Budhendra Bhaduri. 2019. GeoAI at ACM SIGSPATIAL: progress, challenges, and future directions. *SIGSPATIAL Special* 11, 2 (2019), 5–15.
[16] Chao Huang, Chuxu Zhang, Jiashu Zhao, Xian Wu, Dawei Yin, and Nitesh Chawla. 2019. Mist: A multiview and multimodal spatial-temporal learning framework

for citywide abnormal event forecasting. In *The World Wide Web Conference*. 717–728.
[17] Nina Hubig, Philip Fengler, Andreas Züfle, Ruixin Yang, and Stephan Günnemann. 2017. Detection and Prediction of Natural Hazards Using Large-Scale Environmental Data. In *International Symposium on Spatial and Temporal Databases*. Springer, 300–316.
[18] Zhentian Jiao, Youmin Zhang, Jing Xin, Lingxia Mu, Yingmin Yi, Han Liu, and Ding Liu. 2019. A deep learning based forest fire detection approach using UAV and YOLOv3. In *2019 1st International Conference on Industrial Artificial Intelligence (IAI)*. IEEE, 1–5.
[19] Jacqueline Le Moigne, Nathan S Netanyahu, and Roger D Eastman. 2011. *Image registration for remote sensing*. Cambridge University Press.
[20] Miguel D Mahecha, Fabian Gans, Gunnar Brandt, Rune Christiansen, Sarah E Cornell, Normann Fomferra, Guido Kraemer, Jonas Peters, Paul Bodesheim, Gustau Camps-Valls, et al. 2020. Earth system data cubes unravel global multivariate dynamics. *Earth System Dynamics* 11, 1 (2020), 201–234.
[21] Maya Okawa, Tomoharu Iwata, Takeshi Kurashima, Yusuke Tanaka, Hiroyuki Toda, and Naonori Ueda. 2019. Deep mixture point processes: Spatio-temporal event prediction with rich contextual information. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 373–383.
[22] Manizheh Ranjbar, Parham Moradi, Mostafa Azami, and Mahdi Jalili. 2015. An imputation-based matrix factorization method for improving accuracy of collaborative filtering systems. *Engineering Applications of Artificial Intelligence* 46 (2015), 58–66.
[23] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 779–788.
[24] Michele M Rienecker, Max J Suarez, Ronald Gelaro, Ricardo Todling, Julio Bacmeister, Emily Liu, Michael G Bosilovich, Siegfried D Schubert, Lawrence Takacs, Gi-Kong Kim, et al. 2011. MERRA: NASA's modern-era retrospective analysis for research and applications. *Journal of climate* 24, 14 (2011), 3624–3648.
[25] Jiao Sun, Mingxuan Yue, Zongyu Lin, Xiaochen Yang, Luciano Nocera, Gabriel Kahn, and Cyrus Shahabi. 2021. CrimeForecaster: Crime Prediction by Exploiting the Geographical Neighborhoods' Spatiotemporal Dependencies. In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part V*. Springer International Publishing, 52–67.
[26] Paul Viola, Michael Jones, et al. 2001. Robust real-time object detection. *International journal of computer vision* 4, 34-47 (2001), 4.
[27] Michael A Wulder, Thomas R Loveland, David P Roy, Christopher J Crawford, Jeffrey G Masek, Curtis E Woodcock, Richard G Allen, Martha C Anderson, Alan S Belward, Warren B Cohen, et al. 2019. Current status of Landsat program, science, and applications. *Remote sensing of environment* 225 (2019), 127–147.
[28] Jia Yu, Jinxuan Wu, and Mohamed Sarwat. 2015. Geospark: A cluster computing framework for processing large-scale spatial data. In *Proceedings of the 23rd SIGSPATIAL international conference on advances in geographic information systems*. 1–4.
[29] Lili Zhang and Chunhui Zhao. 2017. Tensor decomposition-based sparsity divergence index for hyperspectral anomaly detection. *JOSA A* 34, 9 (2017), 1585–1594.
[30] Andreas Züfle. 2021. Uncertain Spatial Data Management: An Overview. *Handbook of Big Geospatial Data. Springer Nature.* (2021), 355=397.
[31] Andreas Züfle, Goce Trajcevski, Dieter Pfoser, and Joon-Seok Kim. 2020. Managing uncertainty in evolving geo-spatial data. In *2020 21st IEEE International Conference on Mobile Data Management (MDM)*. IEEE, 5–8.
[32] Andreas Züfle, Goce Trajcevski, Dieter Pfoser, Matthias Renz, Matthew T Rice, Timothy Leslie, Paul Delamater, and Tobias Emrich. 2017. Handling uncertainty in geo-spatial data. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*. IEEE, 1467–1470.