Computers & Graphics 103 (2022) 121–128

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Section on EnvirVis

Visualising geospatial time series datasets in realtime with the Digital Earth Viewer

Valentin Buck [a,*,1], Flemming Stäbler [a,1], Jochen Mohrmann [b], Everardo González [a], Jens Greinert [b,c]

[a] Digital Earth, GEOMAR Helmholtz Centre for Ocean Research Kiel, Kiel, Germany
[b] DeepSea Monitoring/Marine Geosystems, GEOMAR Helmholtz Centre for Ocean Research Kiel, Kiel, Germany
[c] Christian-Albrechts University Kiel, Institute of Geosciences, Ludewig-Meyn-Str. 10-12, 24098 Kiel, Germany

ARTICLE INFO

Article history:
Received 30 July 2021
Received in revised form 27 January 2022
Accepted 31 January 2022
Available online 3 February 2022

Keywords:
GIS-software
3D
4D
Data-visualisation
Data-exploration

ABSTRACT

A comprehensive study of the Earth System and its different environments requires understanding of multi-dimensional data acquired with a multitude of different sensors or produced by various models. Here we present a component-wise scalable web-based framework for simultaneous visualisation of multiple data sources. It helps contextualise mixed observation and simulation data in time and space. This work is an extended version of the conference paper (Buck et al., 2021).

1. Introduction

The Earth functions as a complex system where diverse processes of physical, chemical, and biological nature are interlinked on very different spatial and temporal scales. Environmental processes and phenomena can cover entire oceans and span millions of years, or can take place on a meter scale and happen within seconds [1]. Many integrative geoscience studies rely on the analysis of such heterogeneous datasets and such studies are arguably the most diverse compilations of data encountered in any scientific discipline.

The technological advances of the last decades with regards to field sensors, robotics, and modelling have led to an explosion in data collection and generation [2]. If the mining of such data, the understanding of relationships between data sets, and the exploration of models that lead to knowledge should all grow at a comparable speed, geoscientists must benefit from the similar technological advances in the field of data driven science. One way in which digital technologies can support the cognitive tasks of scientists is through the creation of visualisation tools [3]. However, the development of these tools is only justified if they will provide a cost vs. revenue ratio better than that of existing solutions [4].

The Digital Earth Viewer is a web-based visualisation platform capable of ingesting data from heterogeneous sources and performing spatial and temporal contextualisation upon them. The viewer was developed within the context of the Digital Earth Project[2] to allow the parallel navigation of different, large, and diverse datasets in a virtual 4D environment. The examination of further aspects of this data e.g. by using classical 2D graph overlays and user defined operations (for example combining single datasets into new data products) are also implemented.

In particular, we argue that these capabilities make the Digital Earth Viewer perform the tasks of map-like visualisation (identify location, retrieve values, assess distances, and trace paths) [5] with higher effectiveness than other existing software packages. At the same time, online and offline deployment, cross-platform implementation, and a comprehensive graphical user interface, are all capabilities that make the Digital Earth Viewer particularly accessible to geoscience users in the geosciences.

This work is an extended version of the conference paper [6]. The methods section of this previous publication has been expanded in order to give a more comprehensive overview of the software's architecture and data pipeline. The results section was extended using individual instances of the Digital Earth Viewer (including [7] from the same conference) to illustrate the capabilities of the software.

* Correspondence to: GEOMAR Helmholtz Centre for Ocean Research, Wischhofstraße 1-3 24148 Kiel, Germany.
E-mail address: vbuck@geomar.de (V. Buck).

[1] Equal contribution.

[2] https://www.digitalearth-hgf.de/en

https://doi.org/10.1016/j.cag.2022.01.010

2. Related work

2.1. Sensor-specific software

Typically, manufacturers of commercially available sensors provide software tools for data access and visualisation. Such tools are typically proprietary solutions applicable only to outputs from a small range of devices. RDI's WinADCP, a tool for managing and exporting data captured with acoustic Doppler current profilers (ADCPs) manufactured by Teledyne, the SeaSave V7 for oceanography instruments from Sea-Bird Scientific, and the SIS 5 multibeam echosounder software from Kongsberg are examples of this. Using these kinds of applications makes the contextualisation of different sensor data a cumbersome task that requires the synchronisation of multiple unrelated visualisation tools.

2.2. Commercial and open-source GIS software

Many geographic information system (GIS) software packages (e.g. ArcGIS, QGIS, Saga) are used to analyse and visualise data in geographic context. They provide strong analysis capabilities and can access most of the data types commonly used in the geosciences. However, their 3D visualisation capabilities are limited or only enabled through 3rd party plugins (e.g. Qgis2threejs plugin for QGIS [8]). Most of these applications do not treat time as a true dimension, some are only able to do so after lengthy setup operations.

WebGIS applications, like the maps@awi interface, exemplify the benefits of uncomplicated, cross-platform organisation and display of geographical data. In our experience such WebGIS applications are typically limited to a two-dimensional data representation too. One example is Ocean Data View (ODV), a software package for the analysis and visualisation of marine environmental datasets typically acquired through discrete water sampling or CTD (Conductivity, Temperature, Depth) water column profiling during oceanographic expeditions. Its ability to deal with multiple environmental variables, its compatibility with many common marine data formats, and its ability to create visual transformations in the form of 2D graphs and charts, have led to a wide adoption by the scientific community [9].

2.3. Standard GIS servers

Web Map Service (WMS, implemented in ArcGIS Online, Geoserver, …), Unidata's Thematic Real-time Environmental Distributed Data Services (THREDDS, implemented in the TDS [10]) and similar data access protocols provide unified data access for both Web- and regular GIS applications and permit selective loading of data from specific geographic areas and moments in time. WMS in particular can provide projected map data generated from raw data sources, simplifying generation of interactive maps.

2.4. 3D GIS software libraries

Earth [11] is a 2.5D interactive visualisation tool for weather data. It dynamically reprojects georeferenced data into configurable projected coordinate systems at runtime and displays winds, ocean currents and wave direction vector fields as animated particles. The software was developed for the popular visualisation at [12] and its source code has been made available under the MIT licence.

The Cesium engine is a commercial product based on open-source software components which can be used to visualise preprocessed geospatial data through a web browser. CesiumJS applications usually need to be supported by powerful backend servers that preprocess source data into an intermediate format called "3D-tiles". These tiles consist of 3D geometry data and texture information. By default, they can only be scaled uniformly. This helps with performance in regional visualisations, but we found applying effects such as the exaggeration of terrain at global scales to be difficult.

Another product which has seen widespread use for geospatial visualisation is the NASA WorldWind engine. It was originally developed for the use in the NASA Earth Observing System Data and Information System Project and allows for 3D visualisation of WMS-delivered maps. It offers a basic map navigation functionality and serves as a way to present preprocessed maps to an audience in an appealing way. Interactivity is however by default mostly limited to buttons and sliders external to the 3D view; user inputs for precise values for e.g. terrain height exaggeration are not implemented.

The GLOBE (GLobal Oceanographic Bathymetry Explorer) developed at Ifremer is an example of a hybrid application. A Java-based backend is used to convert different data sources to formats readable by an embedded instance of the NASA WorldWind engine. While mostly geared at processing bathymetric and sonar data, some facilities for the import and display of other datatypes (for example point-wise measurements) exist. Gridded datasets are converted into a tiled and pyramided representation on import, leading to relatively good interactive performance. The interactive experience however quickly degrades when more than a small number of point measurements are added and temporal interaction functionality is not easily accessible.

2.5. Conclusion

We found the core capabilities of the visualisation tools mentioned to be limited in one or multiple ways in respect to our use cases. Many of these tools focus on the display of two- and three-dimensional data but few are capable of spatiotemporal exploration of multiple heterogeneous data sources. The 3D and 4D tools that fulfil this expectation, are often geared towards a completely online deployment, which is not always attainable for example during ships expeditions or field campaigns in remote areas.

3. Materials and methods

The Digital Earth Viewer is a hybrid desktop/web application that is conceptually split into a server backend and a web viewer frontend. Both components are bundled into a single native executable which can be launched either as a local desktop-web app (where the application runs completely on the user's machine, but the graphical user interface is displayed by a web browser) or as a server that is accessible from a local area network or the internet. The modes can be configured through command line arguments. Native in this context means that the application is compiled and optimised for the processor architecture it is run on, as compared to a web application, which accepts certain performance restrictions in order to be compatible with all devices.

The server backend is built as a plugin-oriented architecture around a central dispatcher and written in the rust programming language. It integrates a high-throughput multi-threaded asynchronous work scheduler with a set of statically linked plugins which handle data extraction from various file formats. Secondary tasks such as re-gridding the data into viewable areas or caching already extracted sections are also handled in this component. An interface to the client component is provided through an embedded web server which offers a stateless HTTP interface [13].

The client is built as a dynamic Model-View-ViewModel [14] application written in TypeScript. It consists of a WebGL render loop and a VueJS frontend, which are both connected through a service pattern singleton [15].

Since the server backend can be used for either a single user instance or hosted deployment, the optimisation of processing large numbers of complex requests was a development priority. For this purpose, both a multi-threaded asynchronous work scheduler and a caching mechanism are employed. For having the caching mechanism work efficiently, the surface of the earth is divided into a tileset in which each tile covers an approximately equal area and has an aspect ratio close to that of a square. Dividing the earth into tiles makes it more likely that two client instances request the same area. Maintaining an approximately equal area and aspect ratio means that tiles contain similar amounts of data assuming an equally distributed dataset. The cache implementation chosen for the Digital Earth Viewer uses an SQLite on-disk database in which the results from the data extraction requests are stored in a compressed form. Using a statically linked SQLite implementation as a backing database for the cache means that no external database has to be set up and maintained, while search and access speeds in the cache are still high. Furthermore, it allows the cache to be kept in a single file.

Data from heterogenous sources is represented by presentation layers in the user interface. A scene setup containing layers and references to the data sources used can be created in the frontend or loaded from the server. This scene setup is used for data display and runs a data selection algorithm running in the client that determines the sections of data which will be visible to the user from a given point of view and requests this data from the server. In the server, the requests is received and passed to the scheduler, which, after querying its internal cache, passes on the request to the appropriate plugin. The plugin can then either load the data from a file, generate it internally (for example through evaluation of a model) or compute a derivative of the output of another plugin by handing over a second request to the scheduler. When the data is ready, the request terminates with the data being transmitted to the client.

As the client caches the data internally, not every change in the scene setup or viewing parameters triggers a new request, since only new data needs to be loaded. After loading the data into a buffer on the client device's graphics processing unit (GPU), the rendering loop uses a set of shaders specifically developed for the Digital Earth Viewer to create the final display image of the data. For example, in the case of tiled scalar data, the shader pipeline will first raster a floating point texture of the data in the correct three-dimensional projection. This floating-point texture is then mapped to the colour scale selected by the user. Finally, visual effects such as gamma correction or lambert shading are applied by the shader pipeline. In other cases, additional computations are done by the GPU, for example when differences between scalar data layers are computed in the viewport (Fig. 1).

The application's graphical user interface (GUI) consists of a 3D visualisation surface and a series of graphical control elements like buttons and sliders. The visualisation area displays a three-dimensional environment which can be explored by panning, tilting, rotating, and zooming in and out. A time slider component enables the selection of time intervals, its range is automatically adjusted to the temporal boundaries of a selected data layer. The buttons on the top left corner provide access to contextual pop-up windows where the user can control the application and layer settings, add new data layers, set an automatic playback over time, or save and load scene setups. Further GUI components include a colorbar to provide a key for the meaning of colours of the selected data layer, and an interactive wind rose-type North arrow for geographical orientation (Fig. 2).

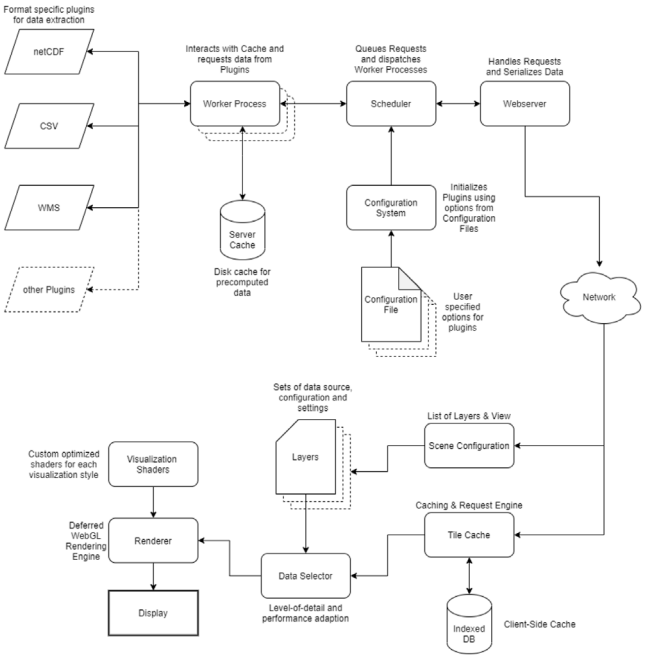

Fig. 1. Runtime dataflow in the Digital Earth Viewer: Data from plugins is sampled through a scheduling mechanism and sent to the client by a web server. The client uses this data as well as a scene configuration and visualisation shaders to create a rendered image.

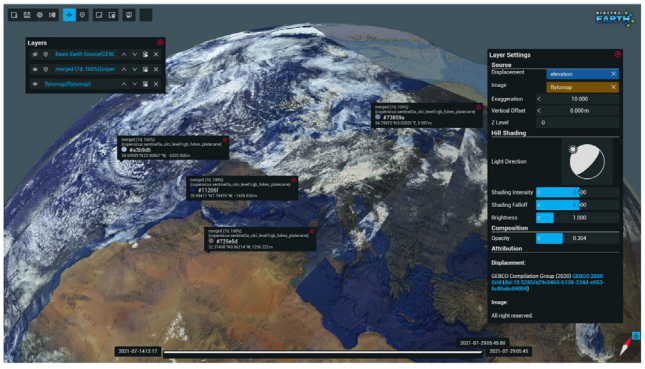

Fig. 2. The Digital Earth Viewer GUI displaying a 3D model of the Earth's surface. The window on the left lists active data layers and the right window contains the settings for the selected layer properties. A set of pins on the surface provide information of the selected layer at this point.

Aside from rendering the graphical representation for viewing on a normal screen, specialised rendering modes exist for the use with projection domes or 3D glasses. This is achieved through post-processing of the 3D scene with fragment shaders that can assemble multiple camera perspectives into one output image. For example, for an anaglyph 3D view, images are rendered from two virtual cameras that are arranged in a pattern like human eyes. These are then colourised with a red tint for the left channel and a blue tint for the right channel and combined into one output image which is then shown. Scenes that have been customised by users can be saved and accessed via a shareable link. Furthermore, scene actions such as movements, temporal selection, or modification of the layers can be synchronised with other users via network through a websocket connection which is specific to the current scene.

The Digital Earth Viewer is an open source project licensed under the EUPL [16] and its source code is available at GEOMAR's GitLab.[3]

4. Results

The Digital Earth Viewer takes advantage of the interplay between server and client technologies. The server component can handle native code for extracting data from complex file formats in an efficient manner, while the client component uses standardised interfaces for graphic display and interaction. This enables the software to display 4D data in real time. It provides a 3D rendered representation of multiple heterogeneous data sources and its navigation treats each spatial dimension and the temporal component as first class citizens. Server executables can be compiled for all three mayor operative systems and thus allows to deploy locally hosted instances for offline use.

The application was tested on different computers to gauge compatibility with various system configurations. We found that 4 GB of random access memory and a multi-core 64-bit CPU running at least 2 GHz, as well as a video accelerator with 2 GB of dedicated memory and a web browser with complete WebGL 1.0 [17] support. We recommend a system with a GPU similar to the NVIDIA GeForce GTX 1060 or the AMD Radeon RX 570, a quad-core processor and 16 GB of RAM for an optimum viewing experience.

This section will illustrate the implications of these capabilities with the help of a series individual use cases of the software.

4.1. GLODAP instance

The Global Ocean Data Analysis Project (GLODAP) is agglomerating chemical water analysis data from almost 1000 research cruises and provides a quality controlled and normalised product containing 12 core variables [18]. The GLODAP dataset encompasses over a million measurements from almost 50 years of scientific exploration. The wide variety of sources that compose the dataset coming from widely differing historical periods and with varying degrees of accuracy, calls for several stages of data cleanup and consistency check [19]. While some of these checks are automated, one important component is still a visual inspection of the data.

Data for the 12 core variables (e.g. O_2, phosphate, nitrate, but also temperature and pressure) comes in form of individual point values tied to point coordinates given by latitude, longitude, and depth and the point in time when the physical sample was obtained. Measurements and samples were undertaken all around the globe starting in the year 1972. Another data product provided by the GLODA-project is gridded data of the same 12 core variables.

To give a 3D context of the ocean, the General Bathymetric Chart of the Oceans (GEBCO) data [20] is used to visualise a global three-dimensional relief of the land and the oceans in this instance. The entire collection of physical samples from the GLODA-project is imported and displayed as point data (Fig. 3). Individual variables can be imported as separate layers and displayed side by side. This three-dimensional representation makes certain outliers in the data evident. Selecting the entire time span provides an overview of the global distribution of all samples and reveals areas that are under- and over-represented. Navigating a shorter time selection uncovers temporal behaviours which include seasonal patterns (e.g. scarce sampling of the Poles during their respective winter months) and preferred sampling locations across decades.

[3] https://git.geomar.de/digital-earth/digital-earth-viewer

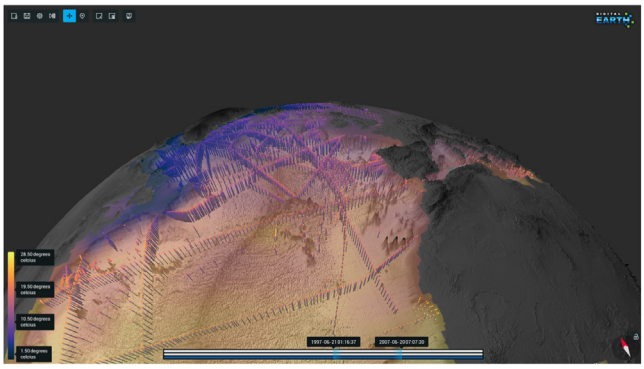

Fig. 3. GLODAP point layers and gridded data layers in the North Atlantic region for the decade between 1997 and 2007.

Point data from single measurements can be directly compared to the global gridded data products to corroborate the precision of the later at different geographical regions and depths. A live version of this instance can be found here: https://digitalearthviewer-glodap.geomar.de.

4.2. Plume tracking instance

Deep sea mining for metals such as nickel, copper, and cobalt is currently being developed further e.g. with respect to technological and economical feasibility. In general it poses the threat of destructing the sensitive ecosystems in which the extraction operations take place. To evaluate the impact of manganese-nodule mining activities at a depth of 4180 m, the JPI Oceans project Mining Impact II [21] studied the distribution of a sediment plume created by a seafloor disturbance. In 2019 a ship-towed dredge was used to simulate a mining operation and monitored the resulting sediment plume. [22] The in situ observations of this small-scale experiment were compared to the results from an ocean current and sediment deposition model.

The challenges at hand consisted on the spatial and temporal contextualisation of multiple heterogeneous data sources:

- A high definition bathymetric relief of the experiment grounds obtained with a multi-beam echo sounder (MBES) mounted on an autonomous underwater vehicle (AUV).
- An array of measuring devices was placed along two transects on the seafloor area. It consisted of seven individual platforms each housing an Aquadopp current profiler and an optical backscatter turbidity sensor (OBS). The turbidity measured by the calibrated OBS sensors served as a proxy for the amount of suspended sediment particles that resulted from the dredging activity [23]. The current velocity of the Aquadopp data helped to characterise the dynamics of the water at the time of the dredging event.
- The Massachusetts Institute of Technology general circulation model (MITgcm) was coupled to a sediment transport module [24,25] in order to recreate the plume dispersion along the dredging tracks.

The sensor array data, the bathymetric data, as well as the sediment transport and deposition model results were imported into the Digital Earth Viewer. The software automatically assigns the data values to objects in a virtual 3D environment and filters time dependent values through the time-slider. In result, the sensor platforms and related sensor readings are correctly placed on the bathymetric terrain, resulting in a visual confirmation of the array's geometry. The output of the deposition model is projected onto the seafloor in the same way. In situ measurements of

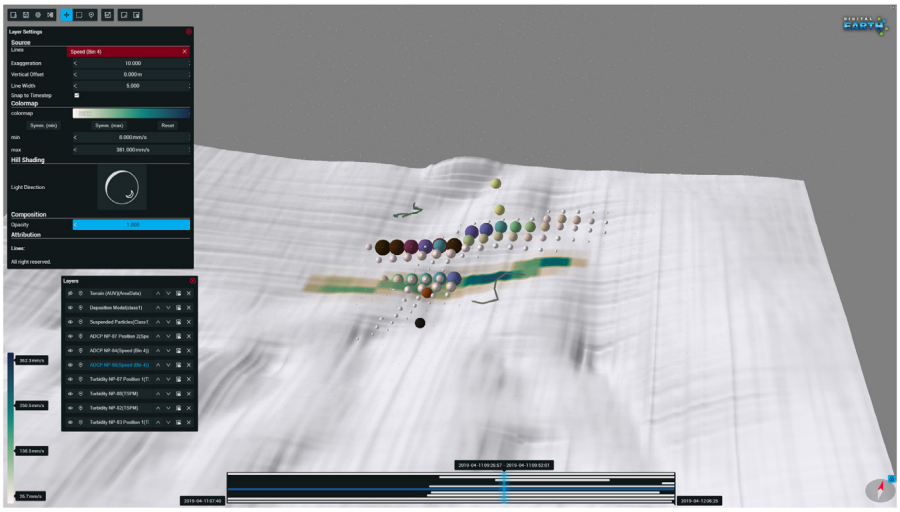

Fig. 4. View of the application displaying a three-dimensional bathymetric map with a sediment deposition map projected on top, a plume dispersion model, and turbidity sensors.

current velocities and water turbidity share the 3D space in the water column with the outputs of the sediment transport model (Fig. 4).

Exploring this virtual environment along the temporal axis reveals the change of the data sources over time. Before the begin of the dredging experiment, current trajectories change with the tides but turbidity sensors, modelled suspended particles, and modelled sediment deposition all show minimal values. Once the experiment started, the spreading of the sediment plume is observed downstream from the dredge track, and values measured by the turbidity sensors begin to spike as the plume is reaching their location. Shortly after, the values of the cumulative sediment deposition begin being noticeable and keep on increasing long after the actual dredging has concluded.

This visualisation allows to compare and validate model results with in situ observations, to make a quick assessment in order to determine at which location what type of sensor is needed to capture the plume dispersion in the best possible way, and, finally, to present the results to non-experts.

A live version of this instance can be found under the following link: https://digitalearthviewer-plume.geomar.de.

4.3. Offline onsite deployment

The experiment outlined in Section 4.2 was conducted using a dredge. For an even better understanding of the environmental impact of "real" deep-sea mining two additional experiment were performed using the mining vehicle Patania II from the Belgian company Global Sea Mineral Resources (GSR). During the MANGAN2021 cruise in the Central North-eastern Pacific Patania II collected Mn nodules in water depths of about 4500 m. (As of the 30th of July, this cruise report is still unpublished).

In order to monitor the suspended particles in the water column a variety of sensors were fixed to 24 different platforms including acoustic and optical turbidity sensors, hydrophones and cameras, and a Kongsberg Hugin 6000 AUV, which was equipped with a Kongsberg EM2040 MBES and a Sea-Bird Scientific Wet Lab FLNTU optical turbidity sensor.

At the start of the experiment the AUV moved in circles at different altitudes (5 m, 10 m, 30 m and 50 m) with a distance of 500 m to the outer edge of the mining area. This was complemented by a far-field lawn mower pattern at 5 m and 10 m altitude. Directly after the AUV was recovered during the cruise, the turbidity data from the FLNTU sensor was exported as a .csv

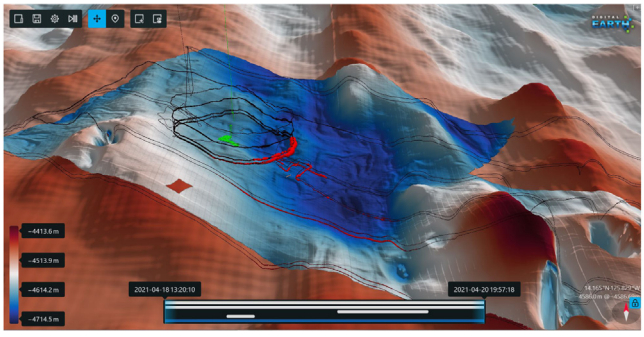

Fig. 5. In the figure a high-resolution bathymetric map is displayed overlaying the lower resolution ship based bathymetry. The Patania II navigation is shown in green and the AUV turbidity in black to red with increasing concentrations.

file including timestamp, a 3D position and the turbidity value. The file was loaded into the Digital Earth Viewer running locally on a notebook, without the need for a dedicated server. The data was complemented by a previously acquired high resolution AUV multi-beam map of the working area available as geoTIFF and the recorded mining vehicle navigation. The point-cloud from the turbidity data was configured using a heat map and by scaling the point diameters with the turbidity value. Using such visualisation settings the AUV path is visible in 3D while making recordings with high turbidity values very prominent along its path. Following a classical approach of using maps and data plotted in QGIS did show the general sediment plume impacted region, but a detailed inspection was not possible. Both, time and height dimensions are crucial for this task, since the AUV visited the same location several times but at different altitude. By moving through time using the Digital Earth Viewer it was possible to compare the mining vehicle location with the recorded plume strength around the area and immediately show the displacement direction and extend of the sediment plume. By combining this information with the 3D terrain, a more in depth discussion about the driving factors causing the plume movement (e.g. current vs. gravity flow) were possible (Fig. 5).

The on board visualisation enabled a fast data exploration and led to immediate actions e.g. for the placement of sensors in a secondary experiment. Subsequent Remotely Operated Vehicle dives to investigate the plume and distribute the different sensor

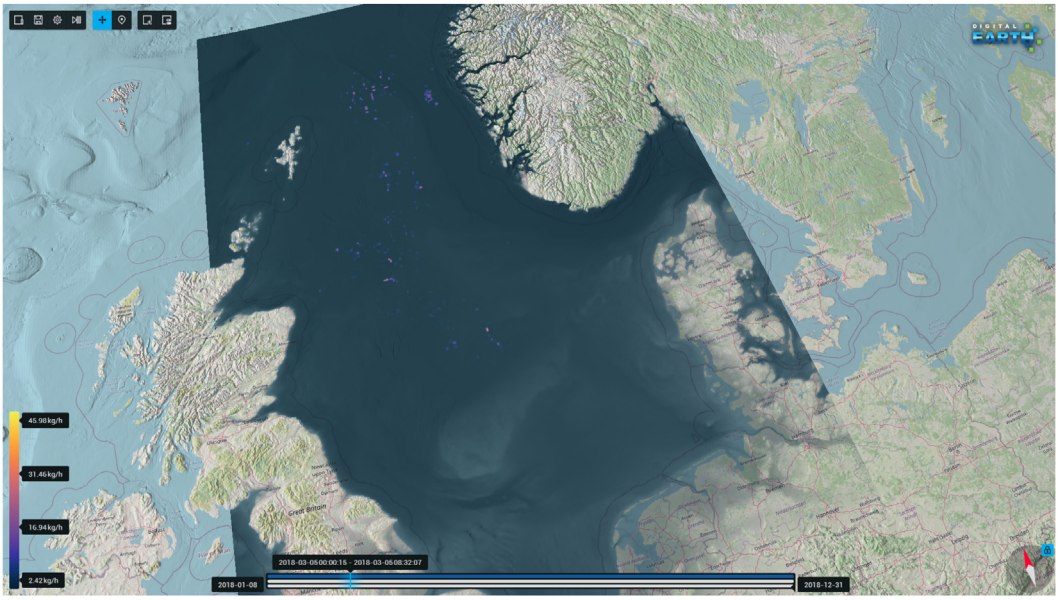

Fig. 6. Top view of the North Sea region with gridded data of methane flows into the atmosphere.

platforms were planned according to the Digital Earth Viewer enabled discussions and secured a better informed secondary experiment. As these data are still being processed to give calibrated values, there is now instance of the currently published.

4.4. North sea methane instance

Atmospheric methane (CH_4) is of great interest to the Environmental Science community because of its role as a greenhouse gas with a greenhouse warming potential (GWP) which is at least 25 times larger than that of CO_2 [26]. The concentration of methane in the atmosphere is a dynamic process which is regulated by a complex interplay of sources (of natural and anthropogenic character) and sinks [27]. While agriculture, agricultural waste, and biomass burning make up for the dominating anthropogenic methane sources, leakage of methane from oil and gas exploration on land and at sea are important sources of unquantified emissions that lead to large differences between top-down and bottom-up estimates [26].

As part of the Digital Earth project, one aim was to highlight these discrepancies for the North Sea region where a series of anthropogenic underwater sources, namely abandoned oil and gas wells, is put into context through comparison with oceanic and atmospheric models.

The 3D relief in this viewer instance (Fig. 6) is composed of OpenStreetMap tiles and the corresponding altitude value from the GEBCO Gridded Bathymetry Data [20]. A set of coordinates of oil and gas wells is placed across the region. Water current, salinity, temperature and the depth of the pycnocline were derived from the NEMO ocean engine [28], and are used to characterise the North Sea waters in 2018. Wind speed and direction, temperature, pressure, and the concentration of certain atmospheric chemicals including CH_4 are obtained from the ICON-ART [29] model for the same year. Finally, CH_4 fluxes from the water column into the atmosphere at the wells are calculated for different temperature, salinity and bubble size scenarios and transformed into gridded data of the sea-atmosphere and cross-pycnocline flux.

Water currents and pycnocline depths at well locations are compared to the gridded CH_4 data above and below the pycnocline revealing dependencies and strong seasonal patterns. The gridded CH_4 data for different flow rates can be visualised on top of each other or in form of difference layers. This shows the strong impact of the initial methane flow quantity on the final gridded methane flux products.

The marine methane fluxes can finally be compared to the wind, temperature, and atmospheric CH_4 models across the entire time span. This provides a sense scale for the local CH_4 emissions while at the same time revealing discrepancies between the bottom-up and top-down approaches.

A live version of this instance can be found under the following link: https://digitalearthviewer-methane.geomar.de.

5. Discussion

The Digital Earth Viewer was compared to other open source software packages to gauge its effectiveness in supporting the visualisation of 4D data.

An often performed operation is the visual exploration of a set of point-wise measurements over a given terrain. For this evaluation, worldwide water temperature measurements from the GLODAP dataset were visualised over the GEBCO ocean terrain data set. The temperature measurements are available as comma separated value files from https://www.glodap.info, while the GEBCO dataset is available both as a netCDF file as well as georeferenced TIFF tiles from https://www.gebco.net.

In QGIS, both datasets are relatively easily loaded and visualised into a 2D view. This view can be interacted with through zooming and panning. In theory, QGIS also supports displaying only points referencing measurements in a specific time range, but since the time fields in GLODAP are separated into component fields for year, month, day, etc., this is non-trivial for this dataset. Furthermore, zooming and panning is relatively slow, since for every movement, the view needs to be rebuilt point by point which takes several seconds. The 3D rendering using the 3rd party pluggin Qgis2Threejs[4] was unsuccessful.

Loading the netCDF data into ParaView was unsuccessful, since the application terminated each time when attempting to load the netCDF dataset after consuming more than 50 GB of working memory.

The Globe Viewer sidesteps this problem by first creating a set of pyramided images for each data layer. While this computation

[4] https://qgis2threejs.readthedocs.io/en/docs/

takes time, it leads to a 3D view that can be interacted with in real time. In our experiments this pyramiding however did not extend to point datasets. While the Globe viewer was able to load the GLODAP .csv file containing around 1.2 million point measurements, it did not draw any of them to the 3D view. On interaction the globe slowed to less than one frame rendered per second.

Both netCDF and CSV are natively supported file formats in the Digital Earth Viewer thus no further conversion of input data was necessary. Two source blocks were added, one integrating the terrain data and one for the temperature measurements. After starting the viewer, a scalar layer was added for the terrain data with its displacement set to the same dataset. The temperature measurements could then be overlaid as a point layer. Since level-of-detail refinement is done at runtime in the Digital Earth Viewer, when zoom or pan actions are undertaken the viewer sometimes needs several seconds to extract new data. In the meantime, upsampled data from lower zoom levels for the correct position is shown.

Aside from these capabilities which make the Digital Earth Viewer a better tool for the display of 4D geoscientific data, further features like such as multi-user interactivity, display on projection domes, and enabling the use of stereographic glasses, are uncommon in other open source GIS software packages.

6. Conclusion

The Digital Earth Viewer's deployment as a hybrid desktop-web application split into a native server backend and a web viewer frontend effectively harnesses the strengths of each technology stack. This makes it possible to access and visualise diverse data types commonly used in Earth Sciences from any modern computer while forgoing the cumbersome task of synchronising multiple unrelated visualisation tools. Its cross-platform nature makes the application accessible to a wide range of users.

The software is capable of addressing the map-like visualisation tasks of location identification, distance assessment, and path tracing with high efficacy. This is because it captures the spatiotemporal nature of geoscientific data more faithfully. The task of data retrieval is augmented by the contextualisation of heterogeneous data sources and by the creation of new data products from user defined combinations.

All of these features lead to an improved visualisation of effects, possibilities, and consequences, which leads to better mental models of 'what is happening' and enables a discussion with others to better understand the Earth System.

CRediT authorship contribution statement

Valentin Buck: Methodology, Software, Writing – original draft, Visualisation. **Flemming Stäbler:** Methodology, Software, Writing – original draft, Visualisation. **Jochen Mohrmann:** Investigation, Writing – review & editing. **Everardo González:** Conceptualisation, Writing – original draft, Writing – review & editing, Investigation. **Jens Greinert:** Conceptualisation, Resources, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] Ruban D. Episodic events in long-term geological processes: A new classification and its applications. Geosci Front 2018;9:377–89. http://dx.doi.org/10.1016/j.gsf.2017.11.004.

[2] Fraser S, Dickson B. Data mining geoscientific data sets using self organizing maps. 2006.

[3] Crapo A, Waisel L, Wallace W, Willemain T. Visualization and the process of modeling: A cognitive-theoretic view. In: Proceeding of the sixth ACM SIGKDD international conference on knowledge discovery and data mining. 2000, p. 218–26. http://dx.doi.org/10.1145/347090.347129.

[4] van Wijk J. The value of visualization. In: VIS 05. IEEE Visualization, 2005. 2005, p. 79–86. http://dx.doi.org/10.1109/VISUAL.2005.1532781.

[5] Hogräfer M, Heitzler M, Schulz H-J. The state of the art in map-like visualization. Comput Graph Forum 2020;39:647–74. http://dx.doi.org/10.1111/cgf.14031.

[6] Buck V, Stäbler F, González E, Greinert J. Digital earth viewer: a 4D visualisation platform for geoscience datasets. In: Dutta S, Feige K, Rink K, Zeckzer D, editors. Workshop on visualisation in environmental sciences. The Eurographics Association; 2021, http://dx.doi.org/10.2312/envirvis.20211081.

[7] González E, Purkiani K, Buck V, Stäbler F, Greinert J. Spatiotemporal visualisation of a deep sea sediment plume dispersion experiment. In: Dutta S, Feige K, Rink K, Zeckzer D, editors. Workshop on visualisation in environmental sciences. The Eurographics Association; 2021, http://dx.doi.org/10.2312/envirvis.20211082.

[8] Akagi M, Casagrande L, Dalang O, K J, Cudini S, Gaspar J, et al. Qgis2threejs plugin. 2013, URL: https://github.com/minorua/Qgis2threejs.

[9] Schlitzer R. Interactive analysis and visualization of geoscience data with ocean data view. Comput Geosci 2002;28:1211–8. http://dx.doi.org/10.1016/S0098-3004(02)00040-7.

[10] THREDDS Data Server, Unidata, https://doi.org/10.5065/D6N014KG, URL: https://www.unidata.ucar.edu/software/tds/.

[11] Beccario C, Irish P. Earth. 2012, URL: https://github.com/cambecc/earth.

[12] Beccario C. Earth :: a global map of wind, weather and ocean conditions. 2012, URL: https://earth.nullschool.net/.

[13] Fielding RT. Architectural styles and the design of network-based software architectures. 2000.

[14] Baillie G, Armour B, Allan D, Milne R, Connolly TM, Beeby R. A model-view-DynamicViewModel and its performance in a web-based component architecture. In: SEKE. 2011.

[15] Mauro C, Leimeister JM, Krcmar H. Service oriented device integration - an analysis of SOA design patterns. In: 2010 43rd Hawaii international conference on system sciences. 2010, p. 1–10. http://dx.doi.org/10.1109/HICSS.2010.336.

[16] European union public licence (EUPL). 2017, European Commission, URL: https://joinup.ec.europa.eu/collection/eupl/eupl-text-eupl-12.

[17] Jackson D, Gilbert J. Webgl specification. 2015, Khronos Group, URL: https://www.khronos.org/registry/webgl/specs/latest/1.0/.

[18] Olsen A, Key RM, van Heuven S, Lauvset SK, Velo A, Lin X, et al. The global ocean data analysis project version 2 (GLODAPv2) – an internally consistent data product for the world ocean. Earth Syst Sci Data 2016;8(2):297–323. http://dx.doi.org/10.5194/essd-8-297-2016, URL: https://essd.copernicus.org/articles/8/297/2016/.

[19] Lauvset SK, Lange N, Tanhua T, Bittig HC, Olsen A, Kozyr A, et al. An updated version of the global interior ocean biogeochemical data product, GLODAPv2.2021. Earth Syst Sci Data Discuss 2021;2021:1–32. http://dx.doi.org/10.5194/essd-2021-234, URL: https://essd.copernicus.org/preprints/essd-2021-234/.

[20] Group GC. GEBCO 2020 grid. 2020, http://dx.doi.org/10.5285/a29c5465-b138-234d-e053-6c86abc040b9, URL: https://www.bodc.ac.uk/data/published_data_library/catalogue/10.5285/a29c5465-b138-234d-e053-6c86abc040b9/.

[21] Gillard B, Purkiani K, Chatzievang, Linke P, Haeckel M. Short cruise report RV SONNE SO268/1+2. 2020, URL: https://www.ldf.uni-hamburg.de/sonne/wochenberichte/wochenberichte-sonne/so267-2-268-3/so268-scr.pdf.

[22] Purkiani K, Gillard B, Paul A, Haeckel M, Haalboom S, Greinert J, et al. Numerical simulation of deep-sea sediment transport induced by a dredge experiment in the northeastern Pacific ocean. Front Marine Sci 2021;8:1175. http://dx.doi.org/10.3389/fmars.2021.719463, URL: https://www.frontiersin.org/article/10.3389/fmars.2021.719463.

[23] Davies-Colley RJ, Smith DG. Turbidity suspenied sediment, and water clarity: A review1. JAWRA J Am Water Resour Assoc 2001;37(5):1085–101. http://dx.doi.org/10.1111/j.1752-1688.2001.tb03624.x, URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1752-1688.2001.tb03624.x.

[24] Marshall J, Adcroft A, Hill C, Perelman L, Heisey C. A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. J Geophys Res Oceans 1997;102(C3):5753–66. http://dx.doi.org/10.1029/96JC02775, URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/96JC02775.

[25] Marshall J, Hill C, Perelman L, Adcroft A. Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. J Geophys Res Oceans 1997;102(C3):5733–52. http://dx.doi.org/10.1029/96JC02776, URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/96JC02776.

[26] Stocker TF, Qin D, Plattner G-K, Tignor MM, Allen SK, Boschung J, et al. Climate change 2013. The physical science basis. Working group i contribution to the fifth assessment report of the intergovernmental panel on climate change - abstract for decision-makers; changements climatiques 2013. les elements scientifiques. Contribution du groupe de travail I au cinquieme rapport d'evaluation du groupe d'experts intergouvernemental sur l'evolution du CLIMAT - resume a l'intention des decideurs. 2013.

[27] Kirschke S, Bousquet P, Ciais P, Saunois M, Canadell JG, Dlugokencky EJ, et al. Three decades of global methane sources and sinks. Nat Geosci 2013;6:813–23. http://dx.doi.org/10.1038/ngeo1955.

[28] Madec G. NEMO Ocean Engine. Project Report, 27, Institut Pierre-Simon Laplace (IPSL); 2008, URL: https://eprints.soton.ac.uk/64324/.

[29] Rieger D, Bangert M, Bischoff-Gauss I, Förstner J, Lundgren K, Reinert D, et al. Icon–ART 1.0 – a new online-coupled model system from the global to regional scale. Geosci Model Dev 2015;8(6):1659–76. http://dx.doi.org/10.5194/gmd-8-1659-2015, URL: https://gmd.copernicus.org/articles/8/1659/2015/.