

Chapter 4

Uncertainty-Aware Enrichment of Animal Movement Trajectories by VGI



Yannick Metz and Daniel A. Keim

Abstract Combining data from different sources and modalities can unlock novel insights that are not available by analyzing single data sources in isolation. We investigate how multimodal user-generated data, consisting of images, videos, or text descriptions, can be used to enrich trajectories of migratory birds, e.g., for research on biodiversity or climate change. Firstly, we present our work on advanced visual analysis of GPS trajectory data. We developed an interactive application that lets domain experts from ornithology naturally explore spatiotemporal data and effectively use their knowledge. Secondly, we discuss work on the integration of general-purpose image data into citizen science platforms. As part of inter-project cooperation, we contribute to the development of a classifier pipeline to semi-automatically extract images that can be integrated with different data sources to vastly increase the number of available records in citizen science platforms. These works are an important foundation for a dynamic matching approach to jointly integrate geospatial trajectory data and user-generated geo-referenced content. Building on this work, we explore the joint visualization of trajectory data and VGI data while considering the uncertainty of observations. *BirdTrace*, a visual analytics approach to enable a multi-scale analysis of trajectory and multimodal user-generated data, is highlighted. Finally, we comment on the possibility to enhance prediction models for trajectories by integrating additional data and domain knowledge.

Keywords Visualization · Trajectory data · User-generated content · Animal tracking

Konstanzer Online-Publikations-System (KOPS)
URL: <http://nbn-resolving.de/urn:nbn:de:bsz:352-2-82mujt57zx1n2>

Y. Metz (✉) · D. A. Keim
Universität Konstanz, Konstanz, Germany
e-mail: yannick.metz@uni-konstanz.de; keim@uni-konstanz.de

© The Author(s) 2024
D. Burghardt et al. (eds.), *Volunteered Geographic Information*,
https://doi.org/10.1007/978-3-031-35374-1_4

4.1 Introduction

Our goal is to integrate and match the two heterogeneous large data sources to enrich the spatial databases with contextual information. Furthermore, we want to investigate in detail the uncertainty that is found in different VGI data sources. In this chapter, we describe our research toward this goal, by focusing on three core contributions:

1. Visual-interactive analysis of GPS trajectory data for domain experts
2. The usage of additional non-verified data sources for VGI in the context of birdwatching
3. A joint approach to analyze GPS trajectory data and VGI contributions considering the uncertainty that, e.g., was introduced by the previous extraction of non-verified data

Additionally, we present work on using visual analytics for the training of deep learning models for movement prediction, which can support future research in domains of geographic information science and forecasting for biologging data.

In the following, we first present the relevant background that motivates our research. Subsequently, we describe pursued research contributing to the aforementioned points, done in conjunction with domain experts or in close collaboration with other partners of the priority program. Finally, we give a brief outlook on possible future research directions.

4.2 Related Work

VGI contains insight that has the potential to solve fundamental and unsolved social and environmental challenges. The big scientific challenge consists of how to investigate and extract value from noisy VGI data sources. Coxen et al. (2017) stressed that there is a general concern about spatial biases in citizen science datasets. Geldmann et al. (2016) also underline that untrained volunteers might be introducing biases, leading to spatial biases toward densely populated areas and easy-to-watch observations. Integrating and assimilation of VGI into scientific models requires a change of paradigm that embraces uncertainty and bias. Some popular citizen science initiatives have already tried to assimilate VGI to improve the results of biodiversity models. Since 2002, the eBird project has been gathering bird observation records from volunteers around the world. The participation of volunteers has shown a rapid increase in recent years with millions of observations submitted each year. Since then, more than 500,000 users have visited the eBird website (Sullivan et al. 2009). Kelling et al. (2012) proposed a Human/Computer Learning Network for Biodiversity Conservation incorporating VGI coming from eBird using an active learning feedback loop to improve the results of the AI algorithms. Fink et al. (2010, 2013) introduced a spatiotemporal exploratory model,

STEM, and AdaSTEM, to study species distribution models. They used massively crowdsourced citizen science data to construct the corresponding models. The STEM model was afterward integrated into a visual analytics system, BirdVis (Ferreira et al. 2011), that allows the ornithologist to analyze abundance models to understand better bird populations. Coxen et al. (2017) compare two species distribution models using satellite tracking data vs. citizen science datasets coming from eBird datasets. Their results showed the effectiveness of citizen science datasets for this particular use. In other disciplines, such as atmospheric sciences, the shift to the user-generated information paradigm is even harder. Chapman et al. (2017) stated that in atmospheric science, high-quality and precise observation is deeply rooted in the essence of the discipline and other data sources with low-quality, bias, and imprecision are hard to accept. Other antecedents to understanding bird migration behavior and patterns are the work of Jain and Dilkina, who constructed a migration network using K-means clustering and Markov chain model (Jain and Dilkina 2015), and the tool of Wood et al., where they studied the seasonal behavior of bird species within a specific location (Wood et al. 2011). There is also previous work to understand the quality of the observations and uncertainty of the models, by characterizing bird watchers. Cole and Scott in 1999 used Texas Conservation Passport holders and members of the American Birding Association to categorize differences between two different groups of wildlife watching as casual wildlife watchers and serious birders (Cole and Scott 1999). These two groups were defined by their skill level at identifying birds, the frequency of participation, expenditures, and bird-watching behavior. Afterward, Scott conducted another study with Thigpen (Scott and Thigpen 2003) to understand bird watchers' behavior. Data were collected from the bird-watching festival in September 1995 at the Seventh Annual HummerBird Celebration in Rockport/Fulton, Texas. In Scott and Shafer (2001), specialization was measured with regard to birdwatcher behavior, level of skill, and commitment. Based on the this, they categorized birders into casual, interested, active, and skilled birders. The outcomes of the study by Scott and Thigpen have supported McFarlane's investigation of birders in Alberta (McFarlane 1994). She revealed that 80% of the general population in her example were casual or novice birders. This information could be very valuable to quantify the confidence of VGI observations. Previous work shows a glimpse of the unprecedented opportunity to materialize a change in the way scientific models treat data, changing the VGI data paradigm to embrace the uncertainty and bias of data provided by humans as part of the scientific investigation process. From now on, a fascinating endeavor comprises for us the development of new techniques that can bridge the gap between social, exact, and natural sciences by using VGI as an interlinkage among biodiversity and human behavior to provide effective and timely answers to societal calls such as climate change and nature preservation.

Previous work has shown how geo-tagged social media data reflects the spatiotemporal distribution of social groups (MacEachren et al. 2011). For example, scatterplots combine event detection and classification in investigating the geo-tagged social media, to enable situation awareness (Thom et al. 2012; Cao et al. 2012) for geospatial information diffusion. Our preliminary work in this area covers

the dynamics of social groups and their expressions on social media. Our work on social media bubbles (Diehl et al. 2018) is the first step toward the structuring of complex social relationships on social media. The main connection point between the social group structure as the scaffolding of society and VGI is the uncertainty that humans introduce into the data and the trustworthiness of systems consumed by humans. This work addresses different aspects of uncertainty from a practical point of view. Our work on the visual assessment for visual abstractions (Sacha et al. 2017) addresses the trustworthiness of the users in the systems for the particular case of soccer data. The works above addressed the uncertainty from the perspective of the producer of VGI and the trustworthiness of the users on the systems. During the last few years, the study of uncertainty and its propagation through the visual analytics workflow have gained popularity. Early, in 2015, MacEachren proposed to consider the propagation of uncertainty through the whole VA workflow rather than just the visualization of the uncertainty at the end of the pipeline (MacEachren et al. 2011). He illustrated current challenges and possible approaches to tackle uncertainty using definitions from decision sciences. Kinkeldey et al. analyzed the impact of visually represented geodata uncertainty on decision-making and addressed possible approaches for the evaluation of uncertainty in visualizations. Previously, we tackled the uncertainty aspects of VGI using a theoretical framework (Diehl et al. 2018), which, for the first time, shapes the human factors of uncertainty in VGI and defines a new term “user uncertainty” to enclose them.

4.3 Analysis of GPS Trajectory Data

We start by looking at an analysis that is possible for stand-alone trajectory data in the biologging context. Analysis tools, both visually and algorithmically, build the foundation for our goal of enabling a joint approach for trajectory data and voluntary geographic information.

4.3.1 *Motivation and Research Gap*

Segmenting biologging time series of animals on multiple temporal scales is an essential step that requires complex techniques with careful parameterization and possibly cross-domain expertise. Yet, there is a lack of visual-interactive tools that strongly support such multi-scale segmentation. To close this gap, we present our MultiSegVA platform for interactively defining segmentation techniques and parameters on multiple temporal scales in our paper *MultiSegVA: Using Visual Analytics to Segment Biologging Time Series on Multiple Scales* (Meschenmoser et al. 2020). MultiSegVA primarily contributes tailored, visual-interactive means and visual analytics paradigms for segmenting unlabeled time series on multiple scales. Further, to flexibly compose the multi-scale segmentation, the platform

contributes a new visual query language that links a variety of segmentation techniques. To illustrate our approach, we present a domain-oriented set of segmentation techniques derived in collaboration with movement ecologists. In the paper, the applicability and usefulness of MultiSegVA are demonstrated in two real-world use cases from movement ecology, related to behavior analysis after environment-aware segmentation and after progressive clustering. Expert feedback from movement ecologists shows the effectiveness of tailored visual-interactive means and visual analytics paradigms in segmenting multi-scale data, enabling them to perform semantically meaningful analyses. Here, we want to highlight two key aspects of the work, the characteristics of biologging time series data and the respective analysis, as well as how we can support this process within the visual analytics framework. For further details, we refer to the paper of Meschenmoser et al. (2020). For our work, we focus on biologging time series of moving animals: these time series have prototypical multi-scale characters and include widely unexplored behaviors, which are hidden in high resolutions and cardinalities. Additionally, biologging-driven movement ecology is an emerging field (Brown et al. 2013; Shepard et al. 2008), triggered by technical advances that enable academia to address open questions in innovative ways. The biologging time series stems from miniaturized tags and gives high-resolution information about, e.g., an animal's location, tri-axial acceleration, and heart rate. Here, semantics are typically distributed on diverse temporal scales, including life stages, seasons, days, day times, and (micro)movement frames. These temporal scales are complemented by spatial scales concerning, e.g., the overall migration range, migration stops, and foraging ranges. There are complex scale- and context-specific conditions (Benhamou 2014; Levin 1992), implying different energy expenditures, driving factors, and decisions for behavior. Hence, segmenting such time series on a single scale with global parameters does not sufficiently address their multi-scale character. The relevance of multi-scale segmentation can be further motivated by three reasons. First, analysts can deepen their understanding of how scales relate to each other: e.g., in terms of nesting relations, next to relative scale sizes and types. A multi-scale perspective can even enable one "to gain an insight on an entire knowledge domain or a relevant sub-part" (Nazemi et al. 2015). Second, even without labeled data or thoroughly parameterized single-scale techniques, it is possible to identify fine-grained patterns that are wrapped by lower-scale, context-yielding patterns. Such fine-grained and context-aware patterns are crucial to enriching existing classification and prediction models. Third, demands for more multi-scale analyses originate from domain literature. Such demands can be found in movement ecology and analysis (Andrienko and Andrienko 2013; Demšar et al. 2015), but also in, e.g., medical sciences (Alber et al. 2019) and social sciences (Cash et al. 2006). However, in practice, segmenting time series on multiple scales is often impeded by several factors. First, multi-scale techniques rely on more in-depth, theoretical foundations and inherent parameters that need to be carefully adapted. Therefore, analysts (e.g., movement ecologists) might require cross-domain expertise in statistical multi-scale time series analysis. Second, even with such expertise, it is difficult to decide on scale properties (e.g., size, dimension, number of scales) and further parameters. Third, we observe a lack of suitable

visual-interactive approaches in related works (Sect. 3.2) that can strongly support and promote segmenting time series on multiple scales.

For MultiSegVa, we defined four requirements in cooperation with domain experts:

1. An application that integrates analysis tools at different time scales without the need to manually combine different algorithms or libraries.
2. Support time series segmentation by revealing the multi-scale structure and addressing its specifics.
3. The analysis should be able to flexibly parameterize segmentation algorithms to the specific context.
4. Visual-interactive features that can help the analysts' work.

4.3.2 Approach

To close the research gap of enabling multi-scale analysis of biologging time series, we present our web-based MultiSegVA platform that allows analysts to visually explore and refine a multi-scale segmentation, which results from a simple way of setting segmentation techniques on multiple scales. In the context of multi-scale segmentation, MultiSegVA primarily contributes to the use of tailored visual-interactive features and established VA paradigms. To flexibly configure segmentation techniques and parameters, MultiSegVA includes a new visual query language (VQL, C2) that links a variety of segmentation techniques across multiple scales. These techniques stem in the present case from a set that was derived together with movement ecologists and covers typical domain use cases. Figure 4.1 shows the main window of the *MultiSegVA* system. Here, the analyst can build visual queries and analyze the existing segmentation results in a hierarchy visualization,



Fig. 4.1 A screenshot from MultiSegVA



Fig. 4.2 The visual query language interface: In the left column, multiple time series segmentation methods can be selected. The hierarchy of applied methods can be changed via an interactive interface, shown in the second column. Finally, detailed settings for each method can be changed in the third column

which in turn is closely linked to one- and multidimensional time series plots. It is also possible to access additional details of segments via a temporal detail window or inspect the underlying trajectories on a map. MultiSegVA implements a feedback loop for iterative analysis and refinement of the segmentation. After importing a time series, e.g., GPS trajectory data of tracked animals via Movebank (Wikelski M 2023), the analyst can start to analyze the time series. After an initial visual inspection of the time series dimensions, an analyst can steer hierarchical time series segmentation using the visual query language (VQL). The interface is shown in Fig. 4.2. The VQL serves three purposes here: (1) Different types of segmentation techniques can be easily arranged across different time scales, (2) the hierarchical application order can be defined by manipulating with building blocks, and (3) the chosen techniques can be interactively parameterized. The query interface first provides a list of available techniques organized by category and can recommend appropriate techniques. To modify the hierarchical application of techniques, selected techniques are arranged as visual building blocks which can be modified by drag-and-drop interactions. This can alleviate some issues that arise with text-based queries. In particular, it avoids changing the ordering of nested queries, which can be tedious and error-prone in text-based queries. The query language also provides several selectors and operations to chain and link different techniques at the same or different scales. The finalized query is then processed in the backend of the application, and the results are visualized via the icicle hierarchy view; see the top of Fig. 4.1. The analyst can choose to adapt the query based on the achieved results to iteratively improve the segmentation results and get a detailed understanding of the time series data.

4.3.3 Results

MultiSegVA enables the comprehensive exploration and refining of multi-scale segmentation by tailored visual-interactive features and VA paradigms. MultiSegVA includes segment tree encoding, subtree highlighting, guidance, density-dependent features, adapted navigation, multi-window support, and a feedback-based workflow. The VQL facilitates exploring and parameterizing different multi-scale structures. Still, a few aspects remain for further reflection. The icicle visualization meets expert requests and has several benefits. Yet, guiding the user by color to interesting parts of the segment tree is a challenging task. We tested global, level-based, and sibling-based guidance variants according to color fills. We chose sibling-based guidance (i.e., all siblings of one hovered segment are colored) that optimally captures local similarities while requiring more navigation effort across levels and nodes. Upcoming works will include an even more effective variant, i.e., guidance to local similarities with little interaction and one fixed color scale. Our VQL makes it trivial to build a multi-scale segmentation. Query building is a play with building blocks that benefits from strong abstraction and simple interactions. Rather, it is difficult to decide which multi-scale structure and building blocks are most appropriate: a decision that depends on data, analysis, and tasks. MultiSegVA facilitates this decision through extensive documentation, technique categorization, few technique parameters, and short processing times in a compact workflow. For further support, we plan predefined queries, and instant responses at query building, next to the parameter and technique suggestions. For suggesting parameters, we will apply estimators (Catarci et al. 1997; Yao 1988) for the number of change points as well as the elbow method for knn-searches. While motif length and HDBSCAN's minPts (Campello et al. 2013) optimally benefit from domain expertise, suggesting other parameters will simplify the interaction and can address another limitation. Now, a technique processes each segment of one scale with the same parameters; thus, slight data-dependent parameter modifications will be examined. For technique suggestions, we envision for each technique a scale-wise relevance score that reflects data properties and is part of a rule-based prioritization, shaped by domain expertise and meaningful hierarchies. It is essential to depict the semantics into which MultiSegVA can provide insights. First of all, MultiSegVA illuminates diverse multi-scale structures and gives insights into how scales relate to each other. Coarse behaviors can be distinguished by relatively simple techniques, motifs show repetitive behaviors, and knn-searches allow the matching with already explored segments. Segment lengths and similarities can be explored, next to local anomalies and spatial contexts. However, with the current techniques, it is difficult to broadly capture deeper, behavioral semantics (e.g., chew, scratch). Hereto more complex or learning techniques (e.g., HMMs, SVMs) will be needed that neither overfill the interface nor delimit generalizability due to the lack of learned patterns. The latter point goes hand in hand with our major limitation and the corresponding implication for upcoming work: integrating even more intelligent methods and automatism. These plans all relate to aspects from above, i.e., better guidance,

technique, and more parameter suggestions, as well as techniques for deeper behavioral semantics. MultiSegVA relies on requirements from movement ecology experts and stands for an iterative, extensively collaborative, and interdisciplinary process. We can gather domain feedback on several stages, derive a domain-oriented set of techniques, and even link MultiSegVA to Movebank with >2.2 billion animal locations. With this application domain-focused, MultiSegVA underpins the value of multi-scale analyses and is certainly another step forward “to empower the animal tracking community and to foster new insight into the ecology and movement of tracked animals” (Spretke et al. 2011). Meanwhile, our third use case shows that MultiSegVA variants for other domains are conceivable, especially with tailored domain-oriented technique sets. This generalizability is promoted by the platform’s I/O features and its ability to handle heterogeneous time series, with >1.2 million records.

4.4 Analysis of VGI Contributor Data

In joint collaborative work with partners of the priority program, we investigate the utility of using a novel pipeline based on a deep learning-based image classifier to integrate images from the social media platform *Flickr* with data from citizen science platforms: *A text and image analysis workflow using citizen science data to extract relevant social media records: combining red kite observations from Flickr, eBird and iNaturalist* (Hartmann et al. 2022). In our research agenda, this work serves a dual role: (1) We explore the characteristics of VGI image data in our chosen domain of migratory birds, as well as automated integration techniques, and (2) we integrate contributions from non-verified data sources, which directly connects to the research topic of uncertainty in data sources. Specifically, the confidence of the developed classification pipeline might be directly used as an uncertainty measure for the matching process we introduce in the following chapter.

4.4.1 Motivation and Research Gap

There is an urgent need to develop new methods to monitor the state of the environment. One potential approach is to use new data sources, such as user-generated content, to augment existing approaches. Despite a wide range of works discussing and demonstrating the potential of new data forms in the creation of indicators, we could not identify previous research which explicitly created a workflow designed to integrate data from different sources and of different modalities. Furthermore, although the properties of different forms of UGC are relatively well understood, they have not been effectively used to develop reproducible workflows. Finally, most studies evaluate the quality of extracted information in isolation through metrics such as precision and recall, but do not explore the added value of integrating data.

In the paper, we propose, implement, and evaluate a workflow taking advantage of citizen science data documenting and recording sightings of birds and more specifically red kites (*Milvus milvus*). Analyzing social media data until recently has often used simple keyword-based methods to perform an initial filtering or search step, meaning that content tagged in other ways was not found. However, improvements in content-based classification now mean that it is also possible to use off-the-shelf, pre-trained algorithms to reliably identify predefined classes such as the presence of buildings, people, or birds in image data with reasonable accuracy.

4.4.2 Approach

We take a new approach, using citizen science projects recording sightings of red kites (*Milvus milvus*) to train and validate a convolutional neural network (CNN) capable of identifying images containing red kites. This CNN is integrated into a sequential workflow that also uses an off-the-shelf bird classifier and text metadata to retrieve observations of red kites in the Chilterns, England. Our workflow reduces an initial set of more than 600,000 images to just 3065 candidate images. Manual inspection of these images shows that our approach has a precision of 0.658. A workflow using only text identifies 14% fewer images than that including image content analysis, and by combining image and text classifiers, we achieve an almost perfect precision of 0.992. Images retrieved from social media records complement those recorded by citizen scientists spatially and temporally, and our workflow is sufficiently generic that it can easily be transferred to other species.

Flickr is a social media site, where individuals can upload photographs and metadata, including tags and locations in the form of coordinates. Flickr's usage has declined in recent years, but it remains very popular in research, mostly because of its well-documented and easy-to-use API, which allows querying using search terms and bounding boxes. Our citizen scientist data came from two platforms: iNaturalist and eBird. iNaturalist allows participants to upload images of organisms such as plants and insects to the platform and use its community to crowdsource taxonomic identification. Currently, according to their website (<https://inaturalist.org>), iNaturalist hosts nearly 100 million observations of over 375000 species and is, therefore, one of the largest and most successful citizen science projects to date (Unger et al. 2020). eBird has similar features to iNaturalist but as a platform is exclusively specialized in bird observations. Their website states (<https://ebird.org>) that "eBird is among the world's largest biodiversity-related science projects, with more than 100 million bird sightings contributed annually." It predominantly hosts observation location data, but also corresponding bird images, as well as bird sounds (Sullivan et al. 2009; Wood et al. 2011).

Since our workflow is designed to be generic, take advantage of the text and image data, and combine records from citizen science reports with social media data, it uses a combination of a simple rule-based approach, existing pre-trained models, and a model trained specifically for our target species. Our approach is designed

to take advantage of what we assume to be high-quality data collected by citizen scientists with an interest in ornithology, use off-the-shelf models where possible, and reduce the initial number of social media posts in a given region to a manageable size for manual verification. In the following, we want to give a summary of the proposed workflow:

1. We identify all geo-tagged social media records in a study area (in our case, the *Chiltern Hills* area in the UK).
2. Out of the identified records, we assign all records that contain the *Latin* name of our target species *Milvus milvus* to our result set. We assume that users familiar with the biological taxonomy are experts and thus treat these records as trustworthy.
3. We use a generic image classification model to identify images that contain birds (with a confidence threshold of p_B 0.5). These retained images are then processed further.
4. For these filtered images, we use metadata such as title or description to identify records that are highly likely Red Kites, e.g., because the description contains the common name in a European language (such as “Red Kite” or “Rotmilan” (German)). We include these images in the result set.
5. We use a secondary image classifier trained on citizen science data to identify images that likely are red kites (with a confidence of $p_{RK} > 50\%$). These are also added to the final set of candidate images.
6. As a final step, we assume that an expert can manually verify the extracted images. As the workflow significantly reduces the set of candidate images, this task becomes feasible and ensures high data quality.

The workflow creates a high-confidence dataset of images that can be integrated with existing citizen science platform data. As part of the paper, we ran a detailed study of the characteristics of the different data sources, namely, *Flickr*, *eBird*, and *iNaturalist*. In the chosen target area, we compare (1) spatial coverage, (2) temporal distribution, (3) contributor patterns, as well as (4) image data quality.

4.4.3 Results

Our workflow aimed to extract relevant images of red kites from Flickr data and to use these to complement citizen science records from eBird and iNaturalist. In the following, we, therefore, explore the following aspects of the results we obtained:

- How effective is our workflow at extracting relevant red kite images, and how much added value is obtained through the use of both text and image content?
- What are the properties of the extracted records within our study area, and do the social media data complement the citizen science platforms?

The workflow returned 3065 candidate images, downsampling the original dataset by 99.5%. These images were then individually inspected to identify true

Table 4.1 Precision using different combinations of the components in the workflow

approach	included posts	true positives	precision
only visual data	2763	1723	0.624
only text data	2215	1946	0.878
initial workflow	3065	2017	0.658
text + visual data (final workflow)	3559	2262	0.636

and false positives and allow us to calculate the precision. Images were marked as true positives if a red kite was identifiable in an image. This meant that images had to be sufficiently clear, such that distinctive features of red kites (e.g., their forked tails or red-brown coloring) were visible. Images where a bird was visible, but not unambiguously identifiable, images showing feathers or pellets, and images that were obviously irrelevant were all marked as false positives. A total of 2017 records were thus identified as true positives, with 1048 false positives and a resulting precision for the complete workflow of 0.658.

To understand the benefits of text and image analysis, we ran the components of the workflow individually and annotated any additional images extracted (Table 4.1).

1. In the textual workflow setting, records were returned if either the Latin name or a common name for red kite (in six language variations) were detected. This approach identified 2215 posts, of which 1946 were true positives and 269 false positives, resulting in a precision of 0.879.
2. In the visual workflow setting, only visual information was considered. A post was considered relevant and included if both the bird model and red kite model return a probability above 50% for the given image. This approach returned 2763 included posts, of which 1723 were true positives and 1040 were false positives, giving a precision of 0.624.

We found 1419 Flickr posts that were included by both settings, of which 1407 were true positives. This means that by only retaining candidate records identified by both textual and image-based information, we can achieve an almost perfect precision of 0.992. We then checked for records that were exclusively identified by either text or image analysis. Five hundred and thirty-nine posts were only detected by the textual analysis (point 1 in the list above), and 316 were only detected by the visual analysis (point 2 in the list above). Combining these results leads to a total of 3559 records, of which 2262 are true positives and 1297 are false positives, and a precision of 0.636. Looking back at the performance of our initial integrated workflow, we note that 245 (12%) additional true positive red kite posts were extracted by merging the results of separately performed textual and visual analysis. This increase in recall is at the cost of a very slight reduction in the precision of 0.02. Summarizing these findings, 62% of true positives were found using either text or image analysis. Twenty-four percent are only correctly classified by textual data, and 14% are missed if no visual analysis is performed.

We find that the workflow functions as a data filter, reducing the data volume by 99.5%. By reducing the data volume, it becomes realistic to analyze the remaining data manually to select true positives. The workflow thus addresses the research gap identified by Burke et al. (2022), using generalizable methods to extract target data from various unverified sources to enrich data.

We found that while keyword matching delivered high precision with little evidence of ambiguity, image analysis returned more potential candidates than textual analysis, but with lower precision. By only retaining posts identified by both textual and visual analysis, they were able to achieve almost perfect precision (0.992), at the cost of a lower recall. By combining the two approaches, they increased the extracted data volume by almost 14% while still downsampling the original dataset by around 99.5% and with a precision 0.636.

The visual distribution of points on the map in the article shows how different sources can complement one another when trying to determine patterns. The locations of Flickr posts tend to cluster around urban areas and points of interest along existing road networks, which suggests that the Flickr observations are often taken opportunistically. eBird and iNaturalist observations, on the other hand, are more heterogeneously allocated and show less obvious relationships to known spatial features, suggesting that birdwatchers go out with a clear intent to observe birds and seek a variety of locations for that purpose.

The study found that the temporal coverage of red kite observations in the Chilterns was different on a yearly and monthly scale. Aggregating data over years showed that the pattern shown by Flickr was different from the ones of eBird and iNaturalist. Year-on-year changes appeared to be more driven by underlying platform dynamics, such as user base and popularity changes. The study found that the rapid drop in Flickr observations from 2012 onward represented a decrease in Flickr popularity rather than a decline in red kites in the Chilterns. On the other hand, the study found that there was a strong increase for eBird and iNaturalist from the year 2016 onward. This could be the result of increased popularity, increased interest in red kites, or increased visits to the study region. Looking at monthly temporal scales showed a trend toward the warmer spring and summer months between March and June. These results may suggest higher visitation rates to the Chilterns in warmer periods, but could also be influenced by specific red kite behavioral patterns. Investigating the number of unique users per data source revealed that representativeness varies between platforms. eBird data was contributed by the fewest individuals, whereas Flickr and iNaturalist offered a more diverse user base. This observation could be attributed to higher platform popularity and an overall larger user base of the latter. Knowing the share of the population represented by a UGC-based analysis is crucial for policymakers to make adequate decisions that reflect the people's opinion (Wang et al. 2019b).

The image quality analysis revealed clear differences between social media data on Flickr and citizen science data in eBird and iNaturalist. Flickr's users are interested in capturing scenic and visually pleasing images, while eBird and iNaturalist users are more concerned about capturing the target species itself as proof of observation and less about the image quality. This discovery may point

to the potential usefulness of social media data for the identification and tracking of individuals.

4.5 BirdTrace: A Visual Analytics Application to Jointly Analyze Trajectory Data and VGI

4.5.1 Motivation and Research Gap

BirdTrace makes use of two primary data sources: GPS data from tagged birds and user-generated content from birders. GPS data typically is of high quality but is only available on a small scale, while user-generated data is more abundant but of variable quality. By combining these two data sources, BirdTrace can provide a more complete picture of bird populations. The system uses a dynamic matching approach to semantically enrich trajectory data with geo-referenced data like images or textual descriptions.

4.5.2 Automated Matching

A key step was the development of semi-automatic methods to extract, integrate, and match data from VGI and tracked spatiotemporal datasets. This has allowed for further knowledge to be gained about individuals and populations of animals, including information about local animal habitats, animal migrations across continents, land-use change, biodiversity loss, invasive species, the spread of diseases, and climate change. There are yet no existing methods and systems that integrate and fuse mixed VGI data from birdwatchers and tracked trajectories of wildlife animals from the ICARUS (Movebank), so we developed them as part of the project.

We have already described the analysis of the trajectories, but now want to find relevant VGI contributions (e.g., images, video, audio, or text descriptions) for trajectories. Here, *relevance* refers to “how well a domain expert can use the found VGI contributions to answer specific questions.” As this implies, the criteria for relevance might therefore depend on the problem. We tackle this problem by giving users the possibility to choose between different matching criteria, e.g., based on spatial or temporal distance, and potential classified behaviors like breeding. By including additional data sources, one might increase the number of possible matching criteria in the future. Let’s look at the way we enable automated matching; see Fig. 4.3:

1. We assume the availability of GPS trajectory data for individuals of a species of interest. As our data source, we utilize Movebank.
2. Secondly, collect and locally store multimedia VGI data from citizen science platforms and potentially Flickr (as described in the previous chapter).




Fig. 4.3 The pipeline and workflow of the BirdTrace system. We combine animal movement trajectory data with VGI data from different citizen science portals. An automated matching approach is used to filter relevant VGI contributions to respective movement trajectories. Both trajectories and VGI data points are then jointly visualized in a shared visual interface. The interface enables aggregating, filtering, or annotating the given data

3. Based on a user query and selected matching criteria, we match individual VGI contributions with GPS trajectory data.
4. Trajectories and VGI contributions are jointly visualized in an interactive visual analytics application, which facilitates analysis by a domain expert.
5. The user can use additional interactive tools to search, filter, and highlight the matched contributions.

To facilitate the matching process, we implemented a data processing pipeline, which applies appropriate preprocessing to both the GPS trajectory data and the VGI contributions. We apply steps like line simplification, motif discovery, and outlier detection to the trajectory data to reduce the size of the data and to simplify the matching computations. VGI contributions from VGI portals like eBird, iNaturalist, and GBIF are collected and processed. To simplify analysis, we use precomputing and caching of data. This enables the efficient clustering of VGI contributions and fast matching of VGI contributions with trajectory data.

4.5.3 A Joint Visual Analytics Workspace

We developed **BirdTrace**,¹ a novel visual analytics method and interfaces to support the semantic annotation of the integrated database consisting of the VGI and the tracked trajectory data. The goal of the application is to add context information semantically from a domain expert (ornithologist) to increase the quality and enrich the integrated data sources previously discussed. The primary challenge is to reduce the uncertainty of the combined data sources and to raise awareness of the remaining uncertainty in our resulting database. Specifically, the tool allows annotating spatiotemporal databases and VGI in their semantic context. We will further support the annotation process with a semi-automatic process to enable the fast and reliable annotation of large datasets. We have to add semi-automatically

¹ Available at <https://birdtrace.dbvis.de>.




Fig. 4.4 A screenshot from BirdTrace

annotation to the dataset, as the domain experts do not have the time to review a large number of matchings.

Using the semantic annotations, we will be able to enrich our joined database with more knowledge from domain experts to increase the data quality of our integrated database. For instance, an ornithologist can verify or oppose the VGI information. The semantic annotation will also assist to clean and prune possibly incorrect merged data records and increase the awareness of uncertainty. Figure 4.4 shows the user interface of BirdTrace, showing a spatial map view on top, and a temporal “timeline” view on the bottom.

4.6 Data-Driven Modeling of Tracked and Observed Animal Behavior

Finally, we explored the challenging tasks of training prediction models, applicable, e.g., to animal trajectory forecasting. To improve uncertainty-aware prediction models for animal trajectories, we explored techniques from deep imitation and reinforcement learning. Specifically, we explored how to use data-driven deep learning methods to predict the movement of fish swarms. A complete presentation of results on predictive models would be beyond the scope of this chapter. We, therefore, want to focus on the workflow and, specifically, how concepts from visual analytics can be used here. We leave a discussion on model implementations for future work. In general, although deep learning-based approaches for related tasks are very promising, we still observe low adoption. Multiple challenges hinder the application of reinforcement learning algorithms in experimental and real-world use cases. Such challenges occur at different stages of the development and deployment of such models. While reinforcement learning workflows share similarities with

machine learning approaches, we argue that distinct challenges can be tackled and overcome using visual analytic concepts. Thus, we propose a comprehensive workflow for reinforcement learning and present an implementation of our workflow incorporating visual analytic concepts integrating tailored views and visualizations for different stages and tasks of the workflow (Metz et al. 2022). In this final section, we would like to shine a light on how our workflow supports experimentation in this space and encourage future research in the application of novel RL-based methods for a wide range of problems in the context of geoinformatics and VGI, e.g., trajectory forecasting.

4.6.1 *Motivation and Research Gap*

Recently, there have been notable examples of the capabilities of reinforcement learning (RL) in diverse fields like robotics (Nguyen and La 2019), physics (Martín-Guerrero and Lamata 2021), or even video compression (Mandhane et al. 2022). Despite these successes, the application and evaluation of recent deep reinforcement and imitation learning techniques in real-world scenarios are still limited. Existing research almost exclusively focuses on synthetic benchmarks and use cases (Bellemare et al. 2013). We argue that the usage and evaluation in realistic scenarios is a mandatory step in assessing the capabilities of current approaches and identifying existing weaknesses and possibilities for further development. In this chapter, we present a visual analytics workflow and an instantiation of the approach that facilitates the application of state-of-the-art algorithms to various scenarios. Our presented approach is designed specifically to support domain experts, with basic knowledge of core concepts in reinforcement learning, who are interested in applying RL algorithms to domain-specific sequential decision-making tasks. The goal is to enable the effective application of their knowledge to (1) design agents and simulation environments including reward functions and (2) a detailed assessment of trained agents' capabilities in terms of performance, robustness, and traceability. A structured and well-defined approach can also help to critically investigate and combat some fundamental difficulties of reinforcement learning like *brittleness*, *generalization* to new tasks and environments, and issues of *reproducibility* (Dulac-Arnold et al. 2020; Henderson et al. 2017).

Outside of reinforcement and imitation learning, there exists a wide range of workflows and interactive visual analytics (VA) tools for the training and evaluation of ML models (Amershi et al. 2015; Endert et al. 2018; Spinner et al. 2020). Compared to other fields of machine learning, there has been less work on applying visual analytics in the space of reinforcement and especially imitation learning. A large number of necessary decisions and the existence of interconnected tasks make the application of interactive machine learning, with a close coupling of model and human, especially valuable for reinforcement learning.

Existing work such as *DQNViz* by Wang et al. (2019a) enables the analysis of spatial behavior patterns of agents in Atari environments like breakout (see arcade learning environment (Bellemare et al. 2013)) using visual analytics. He et al.

present *DynamicsExplorer* (He et al. 2020) to evaluate and diagnose a trained policy in a robotics use case, which incorporates views to track the trajectories of a ball in the maze during episodes. The application enables the inspection of the effect of real-world conditions for trained agents. Saldanha et al. (2019) showcase an application that supports data scientists during experimentation by increasing situational awareness. Key elements are thumbnails summarizing agent performance during episodes and specialized views to understand the connection between particular hyperparameter settings and training performance.

Compared to the existing approaches, we (1) extend the existing frameworks to encompass a holistic view of the relevant stages of the reinforcement learning process instead of just sub-tasks; (2) present a generic, easily adaptable application, which can be instantiated to specific use cases; (3) explicitly consider imitation learning, due to the frequent use in conjunction with reinforcement learning; and (4) apply our framework in a novel, custom real-world use case instead of an existing benchmark environment.

4.6.2 Approach

There has not been a comprehensive workflow for the experimentation and application of reinforcement learning tightly incorporating users. This leaves both researchers and practitioners to loosely defined best practices. In the following chapter, we outline a conceptual workflow for developers and researchers, which we base on guides, projects, and popular open-source libraries. We follow the terminology used, e.g., in the *Gym* package (Brockman et al. 2016). As a starting point, we consider the fundamental workflow from Sacha et al. (2019) that is aimed at generic ML tasks:

1. Prepare-Data: Data selection, cleaning, and transformations; detection of faulty or missing data
2. Prepare-Learning: Specification of an initial model, preparation of training, selection of algorithms, and training parameters
3. Model-Learning: Training of the actual model, monitoring, and supervision
4. Evaluate-Model: Apply the model to testing data, selecting and analyzing quality metrics, and understanding the model

We are interested in highlighting steps and tasks that are specific and critical to reinforcement and imitation learning and which have not been captured previously by more generic workflows. Figure 4.5 summarizes our proposed workflow described in this chapter. In the paper, we discuss the specific stages for imitation and reinforcement learning mirroring the workflow. Specifically, we highlight user tasks during (1) setup and design of the environment which corresponds to mapping a domain-specific problem to a setup applicable to RL algorithms, (2) model training and supervision, and finally (3) evaluation and understanding of trained models. For each of these steps, we present further detailed user tasks and highlight how

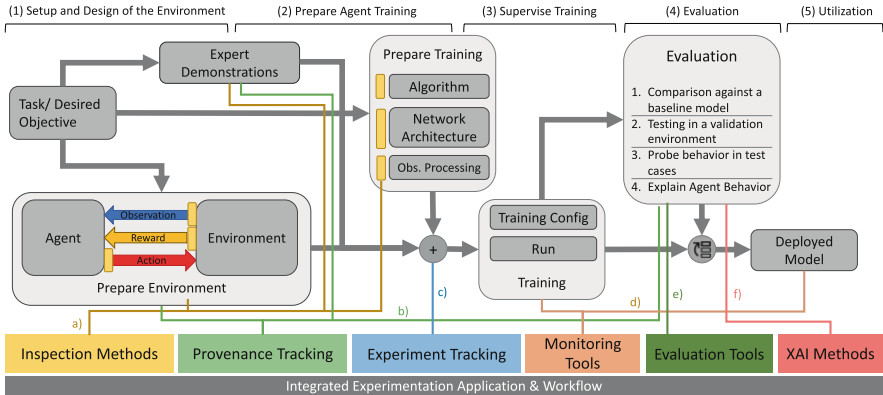


Fig. 4.5 Overview of the **RIVA** (**R**einforcement and **I**mitation Learning with **V**isual **A**nalitics) workflow. **RIVA** is an integrated experimentation workflow and application that provides a range of tools to support all major critical steps: **(a)** inspecting observations, actions, and rewards and ensuring matching values between simulation, expert demonstrations, and architectures, **(b)** provenance tracking of interesting states to enable targeted case-based evaluation, **(c)** tracking of parameters and settings to ensure reproducibility and understand the effect of design decision, **(d)** interactively monitor training and final performance beyond reward, **(e)** enable effective evaluation by integrating multiple evaluation tools, and **(f)** explain behavior by natively integrating XAI methods like input attribution techniques

visual analytics concepts are applicable. We apply the proposed framework and developed an application in the use case of imitation and reinforcement learning for collective behavior: data-driven learning of the behavior of fish schools (collective movement of fish swarms). We cooperated with a domain expert throughout the entire process, from designing custom environments and agents, training, to final evaluation. Modeling the behavior of individual actors in swarm systems has been a long-standing problem in biology (Reynolds 1987; Sumpter 2006; Calovi et al. 2013). Learning individual policies that lead to coordinated collective behavior via both reinforcement learning and imitation learning from recorded trajectories is an exciting application that promises to overcome existing simplifications in hand-crafted models.

4.6.3 Results and Discussion

The use case can be well integrated into our workflow and application with minimal modifications. Noticeably, a custom interactive rendering of the environment was added. We utilize the modularity of the software to integrate additional components like custom visualizations. During the design phase, the inspection views were used to ensure consistency between environment, agent, and dataset, e.g., to spot premature episode termination. Our workflow was highly effective in maintaining

a high level of productivity and consistency through an iterative design process, in which we experimented with different observation space designs, reward functions, network types, and hyperparameter configurations. The set of evaluation tools is used both for internal evaluation and external presentation.

4.7 Discussion and Conclusion

In this chapter, we have given an overview of research contributing to the overall goal of enriching high-quality sparse and curated data with VGI contributions to enable different applications like analysis of species distribution or prediction modeling of movement. In particular, we highlighted the potential of visual analytics solutions for different stages of curation, analysis, and model building using VGI data. Visualizations can be especially suited to present data of varying quality and express uncertainty. Both our joint collaborative work on integrating Flickr images with citizen science data and the BirdTrace platform highlight the potential of integrating different data sources, ranking from citizen science platforms, social media, to professional data collection efforts.

Acknowledgments This research was supported by the German Research Foundation DFG within Priority Research Program 1894 *Volunteered Geographic Information: Interpretation, Visualization, and Social Computing* (VGIscience). In this chapter, results from the project *Uncertainty-Aware Enrichment of Animal Movement Trajectories by VGI* (BirdTrace, Project number 314671965) are reviewed.

References

- Alber M, Buganza Tepole A, Cannon WR, De S, Dura-Bernal S et al (2019) Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *NPJ Digit Med* 2(1):1–11
- Amershi S, Chickering M, Drucker SM, Lee B, Simard P et al (2015) Modeltracker: redesigning performance analysis tools for machine learning. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, New York, NY, USA, 2015. Association for Computing Machinery, pp 337–346. ISBN 9781450331456. <https://doi.org/10.1145/2702123.2702509>
- Andrienko N, Andrienko G (2013) Visual analytics of movement: an overview of methods, tools and procedures. *Inform Visualization* 12(1):3–24
- Bellemare MG, Naddaf Y, Veness J, Bowling M (2013) The arcade learning environment: An evaluation platform for general agents. *J Artif Intell Res* 47:253–279
- Benhamou S (2014) Of scales and stationarity in animal movements. *Ecol Lett* 17(3):261–272
- Brockman G, Cheung V, Pettersson L, Schneider J, Schulman J et al (2016) Openai gym
- Brown DD, Kays R, Wikelski M, Wilson R, Klimley AP (2013) Observing the unwatchable through acceleration logging of animal behavior. *Animal Biotelem* 1(1):1–16
- Burke HM, Tingley R, Dorin A (2022) Tag frequency difference: rapid estimation of image set relevance for species occurrence data using general-purpose image classifiers. *Ecol Inform* 69:101598. ISSN 1574-9541. <https://doi.org/10.1016/j.ecoinf.2022.101598>

- Calovi D, Lopez U, Ngo S, Sire C, Chaté H et al (2013) Swarming, schooling, milling: phase diagram of a data-driven fish school model. *New J Phys* 16:015026. <https://doi.org/10.1088/1367-2630/16/1/015026>
- Campello RJ, Moulavi D, Sander J (2013) Density-based clustering based on hierarchical density estimates. In: *Advances in Knowledge Discovery and Data Mining: 17th Pacific-Asia Conference, PAKDD 2013, Gold Coast, Australia, April 14–17, 2013, Proceedings, Part II* 17. Springer, Berlin, pp 160–172
- Cao N, Lin Y-R, Sun X, Lazer D, Liu S et al (2012) Whisper: tracing the spatiotemporal process of information diffusion in real time. *IEEE Trans Visualization Comput Graph* 18(12):2649–2658
- Cash DW, Adger WN, Berkes F, Garden P, Lebel L et al (2006) Scale and cross-scale dynamics: governance and information in a multilevel world. *Ecol Soc* 11(2)
- Catarci T, Costabile MF, Levaldi S, Batini C (1997) Visual query systems for databases: a survey. *J Visual Lang Comput* 8(2):215–260
- Chapman L, Bell C, Bell S (2017) Can the crowdsourcing data paradigm take atmospheric science to a new level? A case study of the urban heat island of London quantified using Netatmo weather stations. *Int J Climatol* 37(9):3597–3605
- Cole JS, Scott D (1999) Segmenting participation in wildlife watching: a comparison of casual wildlife watchers and serious birders. *Hum Dimens Wildl* 4(4):44–61
- Coxen CL, Frey JK, Carleton SA, Collins DP (2017) Species distribution models for a migratory bird based on citizen science and satellite tracking data. *Global Ecol Conserv* 11:298–311
- Demšar U, Buchin K, Cagnacci F, Safi K, Speckmann B et al (2015) Analysis and visualisation of movement: an interdisciplinary review. *Movement Ecol* 3(1):1–24
- Diehl A, Yang B, Das RD, Chen S, Andrienko G et al (2018) User-uncertainty: a human-centred uncertainty taxonomy for VGI through the visual analytics workflow
- Dulac-Arnold G, Levine N, Mankowitz DJ, Li J, Paduraru C et al (2020) An empirical investigation of the challenges of real-world reinforcement learning. arXiv e-prints, art. arXiv:2003.11881
- Endert A, Ribarsky W, Turkyay C, Wong W, Nabney I et al (2018) The state of the art in integrating machine learning into visual analytics. arXiv e-prints, art. arXiv:1802.07954
- Ferreira N, Lins L, Fink D, Kelling S, Wood C et al (2011) Birdvis: visualizing and understanding bird populations. *IEEE Trans Visualization Comput Graph* 17(12):2374–2383
- Fink D, Hochachka WM, Zuckerberg B, Winkler DW, Shaby B et al (2010) Spatiotemporal exploratory models for broad-scale survey data. *Ecol Appl* 20(8):2131–2147
- Fink D, Damoulas T, Dave J (2013) Adaptive spatio-temporal exploratory models: Hemisphere-wide species distributions from massively crowdsourced ebird data. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol 27, pp 1284–1290
- Geldmann J, Heilmann-Clausen J, Holm TE, Levinsky I, Markussen B et al (2016) What determines spatial bias in citizen science? Exploring four recording schemes with different proficiency requirements. *Diversity Distrib* 22(11):1139–1149
- Hartmann MC, Schott M, Dsouza A, Metz Y, Volpi M et al (2022) A text and image analysis workflow using citizen science data to extract relevant social media records: combining red kite observations from flickr, ebird and inaturalist. *Ecol Inform* 71:101782
- He W, Lee T, van Baar J, Wittenburg K, Shen H (2020) Dynamicexplorer: visual analytics for robot control tasks involving dynamics and LSTM-based control policies. In: *2020 IEEE Pacific Visualization Symposium (PacificVis)*, pp 36–45. <https://doi.org/10.1109/PacificVis48177.2020.7127>
- Henderson P, Islam R, Bachman P, Pineau J, Precup D et al (2017) Deep reinforcement learning that matters. In: *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp 3207–3214. <http://arxiv.org/abs/1709.06560>.
- Jain N, Dilkina B (2015) Coarse models for bird migrations using clustering and non-stationary Markov chains. In: *AAAI Workshop: Computational Sustainability*
- Kelling S, Gerbracht J, Fink D, Lagoze C, Wong W-K et al (2012) ebird: a human/computer learning network for biodiversity conservation and research. In: *Twenty-Fourth AAAI Conference*
- Levin SA (1992) The problem of pattern and scale in ecology: the robert h. macarthur award lecture. *Ecology* 73(6):1943–1967

- MacEachren AM, Jaiswal A, Robinson AC, Pezanowski S, Savelyev A et al (2011) Senseplace2: Geotwitter analytics support for situational awareness. In: 2011 IEEE conference on visual analytics science and technology (VAST). IEEE, pp 181–190
- Mandhane A et al (2022) MuZero with self-competition for rate control in VP9 video compression. arXiv e-prints, art. arXiv:2202.06626
- Martín-Guerrero, JD, Lamata L (2021) Reinforcement learning and physics. *Appl Sci* 11(18). ISSN 2076-3417. <https://doi.org/10.3390/app11188589>
- McFarlane BL (1994) Specialization and motivations of birdwatchers. *Wildl Soc Bull* 361–370
- Meschenmoser P, Buchmüller JF, Seebacher D, Wikelski M, Keim DA (2020) Multisegva: using visual analytics to segment biologging time series on multiple scales. *IEEE Trans Visualization Comput Graph* 27(2):1623–1633
- Metz Y, Schlegel U, Seebacher D, El-Assady M, Keim DA (2022) A comprehensive workflow for effective imitation and reinforcement learning with visual analytics. In: 13th International EuroVis Workshop on Visual Analytics (EuroVA 2022), pp 19–23
- Nazemi K, Burkhardt D, Ginters E, Kohlhammer J (2015) Semantics visualization—definition, approaches and challenges. *Proc Comput Sci* 75:75–83
- Nguyen H, La H (2019) Review of deep reinforcement learning for robot manipulation. In: 2019 Third IEEE International Conference on Robotic Computing (IRC), pp 590–595. <https://doi.org/10.1109/IRC.2019.00120>
- Reynolds CW (1987) Flocks, herds and schools: a distributed behavioral model. In: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '87, New York, NY, USA, 1987. Association for Computing Machinery, pp 25–34. ISBN 0897912276. <https://doi.org/10.1145/37401.37406>
- Sacha D, Al-Masoudi F, Stein M, Schreck T, Keim DA et al (2017) Dynamic visual abstraction of soccer movement. In: *Computer Graphics Forum*, vol 36. Wiley Online Library, pp 305–315
- Sacha D, Kraus M, Keim DA, Chen M (2019) VIS4ML: an ontology for visual analytics assisted machine learning. *IEEE Trans Vis Comput Graph* 25(1):385–395. <https://doi.org/10.1109/TVCG.2018.2864838>
- Saldanha E, Praggastis B, Billow T, Arendt D (2019) ReLVis : visual analytics for situational awareness during reinforcement learning experimentation. In: EuroVis (Short Papers). Eurographics Association, pp 43–47. <https://doi.org/10.2312/evs.20191168>
- Scott D, Shafer CS (2001) Recreational specialization: a critical look at the construct. *J Leisure Res* 33(3):319–343
- Scott D, Thigpen J (2003) Understanding the birder as tourist: segmenting visitors to the Texas hummer/bird celebration. *Hum Dimensions Wildl* 8(3):199–218
- Shepard EL, Wilson RP, Quintana F, Laich AG, Liebsch N et al (2008) Identification of animal movement patterns using tri-axial accelerometry. *Endangered Species Res* 10:47–60
- Spinner T, Schlegel U, Schäfer H, El-Assady M (2020) explainer: a visual analytics framework for interactive and explainable machine learning. *IEEE Trans Vis Comput Graph* 26(1):1064–1074. <https://doi.org/10.1109/TVCG.2019.2934629>
- Spretke D, Bak P, Janetzko H, Kranstauber B, Mansmann F, Davidson S (2011) Exploration through enrichment: a visual analytics approach for animal movement. In: Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, pp 421–424.
- Sullivan BL, Wood CL, Ilf MJ, Bonney RE, Fink D et al (2009) ebird: a citizen-based bird observation network in the biological sciences. *Biol Conserv* 142(10):2282–2292
- Sumpter D (2006) The principles of collective animal behavior. *Philos Trans R Soc Lond Series B: Biol Sci* 361:5–22. <https://doi.org/10.1098/rstb.2005.1733>
- Thom D, Bosch H, Koch S, Wörner M, Ertl T (2012) Spatiotemporal anomaly detection through visual analysis of geolocated twitter messages. In: 2012 IEEE Pacific Visualization Symposium. IEEE, pp 41–48
- Unger S, Rollins M, Tietz A, Dumais H (2020) inaturalist as an engaging tool for identifying organisms in outdoor activities. *J Biol Educ* 0(0):1–11. <https://doi.org/10.1080/00219266.2020.1739114>

- Wang J, Gou L, Shen HW, Yang H (2019a) DQNViz: a visual analytics approach to understand deep Q-networks. *IEEE Trans Vis Comput Graph* 25(1):288–298. ISSN 19410506. <https://doi.org/10.1109/TVCG.2018.2864504>
- Wang Z, Hale S, Adelani DI, Grabowicz P, Hartman T et al (2019b) Demographic inference and representative population estimates from multilingual social media data. In: *The World Wide Web Conference, WWW '19*, New York, NY, USA, 2019b. Association for Computing Machinery, pp 2056–2067. ISBN 9781450366748. <https://doi.org/10.1145/3308558.3313684>
- Wikelski M, Davidson SC (2023) Movebank: archive, analysis and sharing of animal movement data. Hosted by the max planck institute of animal behavior. www.movebank.org
- Wood C, Sullivan B, Iliff M, Fink D, Kelling S (2011) ebird: engaging birders in science and conservation. *PLoS Biol* 9(12):e1001220
- Yao Y-C (1988) Estimating the number of change-points via Schwarz' criterion. *Stat Probab Lett* 6(3):181–189

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

