Sustainability: Energy-efficient intelligent bicycle sensors to promote mobility transition

Paula Scharf, Benjamin Karic, Luca Hesse, and Thomas Bartoschek

Abstract: Cycling plays a key role in pushing towards more environmentally friendly forms of mobility and ensuring the safety of cyclists, which is crucial to motivate more people to choose the bike as their preferred mode of transportation. The still car-centred infrastructure calls for innovative methods to ensure a safe environment for cyclists. In Germany, despite the legal requirement of a minimum distance of 1.5 metres when cars overtake bicycles, in practice, compliance is often inadequate. In this work, we propose and prototype a mobile, low-cost, and low-energy system that allows cyclists to monitor the occurrence of dangerously close takeovers. This system employs a time-of-flight sensor that measures distances in an 8x8 pixel matrix and uses tiny machine-learning models to detect dangerous takeovers. When included in citizen science campaigns, the collected data can be used for urban mobility planning and safety strategies.

Keywords: sensors, citizen science, machine learning, urban mobility, sustainability

1. Introduction

The transformation of transportation is a major societal challenge in achieving climate neutrality. Technological innovations like improved renewable energy production and battery technologies are critical steps toward this goal. But just as important is reducing the use of private fuel-powered vehicles and increasing the use of more climate-friendly modes such as bicycles, pedelecs, and public transportation whenever possible to promote green transportation.

While almost half of all Germans would prefer to ride their bicycles more often (Sinus Marktund Sozialforschung GmbH, 2023), cars are still predominantly used – even for short commutes under 5km (Statistisches Bundesamt (Destatis), 2022). The low usage of bicycles, even for short distances, can be attributed to the lack of proper infrastructure: 40% of cyclists state that they do not feel safe while cycling due to excessive traffic, reckless driving by motor vehicles, and a lack of designated bike lanes (Sinus Markt- und Sozialforschung GmbH, 2023). Cars performing takeover manoeuvres on streets without separated bike lanes pose a typical safety threat to cyclists. Despite the German Road Traffic Regulations (StVO) stipulating a minimum passing distance of 1.5 metres between cars and cyclists in urban areas (§ 5 Abs. 4 S. 3 StVO), compliance with this requirement is often lacking in practice. Such situations not only pose an immediate safety risk but also have a significant impact on people's willingness to use bicycles as an everyday means of transportation. In fact, in 2022, a comprehensive study found that 42% of German respondents would choose a car over a bicycle, even for short journeys, due to the unsafe cycling infrastructure (Pyhel, 2022).

To improve cycling infrastructure, it is crucial to identify high-risk areas. Subsequently, implementing safety measures and promoting cycling as a safe and sustainable mode of transportation can help overcome the negative safety perceptions associated with cycling. However, identifying high-risk locations in Germany's extensive transportation infrastructure presents a significant challenge to municipalities. The task at hand presents a two-fold challenge: 1) detecting high-risk locations in cycling infrastructure and 2) ensuring the scalability of the approach. To address this challenge, we propose a concept using a citizen science approach to identify safety gaps in urban cycling infrastructure that result in dangerous overtaking manoeuvres.

We present a prototype based on the proposed concept that combines cutting-edge time-of-flight sensor technology, neural networks, and a senseBox (Wirwahn et al., 2015) to automatically identify high-risk locations. The method uses low-cost and energy-efficient technologies that can be implemented in a citizen science project to generate large-scale data on the safety of cycling infrastructure while engaging cyclists in improving cycling safety. The objective is to establish a sustainable data collection method to support proactive accident prevention, enhance road safety, and improve urban cycling infrastructure. Research in this area is highly relevant as it positively impacts the sustainability of urban mobility planning and safety strategies.

Existing efforts in this field often rely on manual, qualitative incident reporting. Prominent examples of this approach include bikemaps.org (Nelson et al., 2015) and the Bike Barometer (Storme et al., 2022). However, data collection is often exhaustive and sporadic for such solutions, leading to a lack of broad spatial and temporal coverage.

CycleAI[1] has taken a different approach to bicycle road safety. They have trained an AI algorithm on how ‘bikeable’ a road is based on static images (from Google Street View) of that stretch of road. Based on these ratings, a routing algorithm suggests safer routes. This algorithm considers only the infrastructure depicted in the images and does not take into account additional environmental data or a temporal component. Despite achieving better spatial coverage, it does not consider varying safety at different times of the day or with changing weather conditions.

The OpenBikeSensor project[2] offers a semi-automated solution for recording violations of the minimum takeover distance. They use an ultrasonic proximity sensor to measure the distance to passing cars when the cyclist presses a button (e.g. Hauenstein et al., 2023). However, it can be challenging to trigger the device reliably in hazardous traffic situations, which may result in forgetting to capture data or recording inaccurate data.

The projects ‘Sicher überholt!’ (Blees, 2021) and ‘Radmesser’ (Lehmann et al., 2018) propose a fully automatic system that uses a combination of a smartphone camera and an ultrasonic proximity sensor. The smartphone automatically captures images to confirm a car takeover triggered by ultrasonic sensor measurements. Nonetheless, the use of cameras raises privacy concerns and increases energy consumption.

The senseBox:bike[3] project involves using a senseBox to measure distances to the left side of a bicycle with an ultrasonic proximity sensor. However, the primary purpose of the project is to serve as a mobile environmental monitoring station for bicycles. Although distance values are captured, they are not processed to further determine if a car has actually overtaken the bicycle.

2. Concept

In this work, we propose a mobile, low-cost, and low-energy takeover detection system as schematised in Figure 1. At the heart of our system is a microcontroller unit (MCU) – a compact unit of processor, memory,

1 https://cycleai.net/ (last visited 22.04.2024)
2 https://www.openbikesensor.org/ (last visited 22.04.2024)
3 https://sensebox.de/de/products-bike (last visited 22.04.2024)

and periphery, combining all essential parts of a computer on a small scale with low requisites in terms of cost and power. Specifically, an MCU of the senseBox family is chosen. The senseBox is a versatile electronics kit designed specifically for citizen science projects and educational initiatives, focusing on environmental monitoring and data collection (Witte et al., 2023). Originally backed by the German Federal Ministry of Education and Research, it was designed to promote understanding of environmental science and technological applications. In addition to providing hardware and software components, the senseBox enables users to gather, analyse, and disseminate their own environmental data. Its anchoring in open-source technology and compatibility with the popular Arduino platform[4] for programming and prototyping electronics make it easily accessible and adaptable for a wide range of users and use cases (Bartoschek et al., 2019).

4 https://www.arduino.cc/en/Guide/Introduction (last visited 25.04.2024)

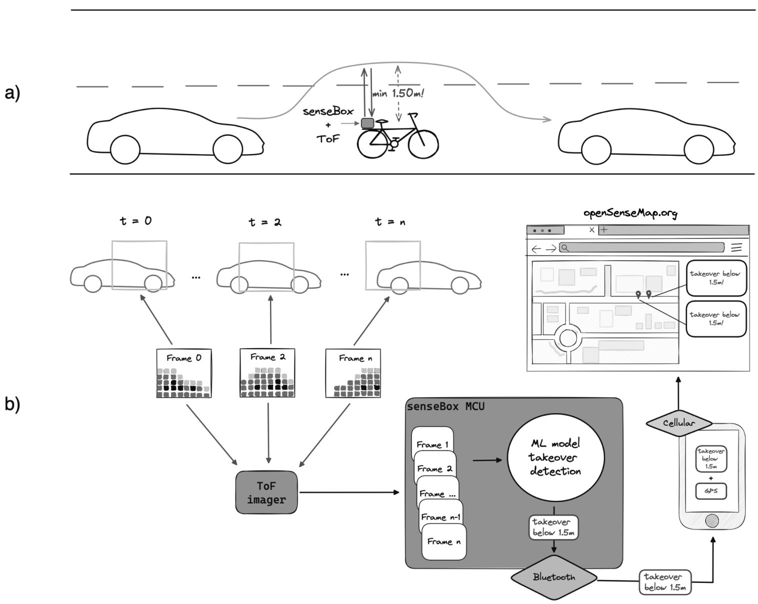

Figure 1. A conceptual model of the takeover detection system. a) A car taking over a bicycle monitored by an automatic detection system; b) Architecture and flow of the detection system. A time series of depth images captured by a Time-of-Flight (ToF) imager is classified on a senseBox MCU using machine learning. Results are transmitted via Bluetooth to a mobile application for geo-location and mobile data transfer to the open data platform openSenseMap (Pfeil et al., 2015).

A proximity sensor is utilised to sense passing traffic, particularly one that can capture very low-resolution depth images. Depth images, as opposed to single distance readings, enable the system to perceive not only the presence of objects but also their approximate shape and dimensions. In contrast to cameras, they offer a high level of privacy (no images of faces or licence plates can be acquired) and consume less power. Especially in light of the reduced processing power of MCUs, it is sensible to keep the amount of captured information low, including the image resolution. Additionally, since determining the distance of a passing car is crucial for assessing the safety of a manoeuvre, depth images are preferred over standard colour images. In this context, the VL53L8CX multizone time-of-flight (ToF) sensor[5] poses a perfect fit for our use case. ToF sensors capture distances/depths by emitting infrared light and measuring the time it takes for the light to be reflected. Compared to other options for recording depth images, the VL53L8CX sensor has particularly low requisites in terms of size, cost, and energy consumption. More detailed specifications of the sensor will be given in the following section.

The low-resolution depth images captured by the ToF sensor are well-suited for machine learning applications, being decently complex and often noisy. Furthermore, the usage of machine learning algorithms allows for flexibility in problem definition, for example, by detecting the type of the passing vehicle (car, truck, bicycle) and its speed and distance all in a single algorithm. To detect takeover manoeuvres, sequences of recorded depth images and their temporal order need to be considered. Recurrent neural networks (RNNs) offer a suitable machine learning methodology for this combination of classification problems and available data. They allow current input data (e.g. a single depth image) to be put in the context of previously recorded input data. This ability enables the analysis of the entire movement process of an object, such as a car, over time, which is crucial for distinguishing takeovers from other situations, such as oncoming traffic.

To put the detected takeovers into a geographic context and upload the resulting data, we adopt existing solutions built within the senseBox:bike project. Previous iterations of the senseBox:bike relied on a dedicated GPS module for georeferencing and utilised WiFi for data upload whenever a signal from preconfigured networks was available, typically the cyclist's home network. While functional, newer versions of the senseBox:bike improve this workflow by utilising a smartphone app through Bluetooth connection, in which both the position is determined and data is uploaded to the open data platform openSenseMap (Pfeil et al., 2015). The app streamlines the process, minimises code and configuration requirements, reduces the risk of data loss, and allows cyclists to begin recording immediately rather than waiting for the dedicated GPS module to establish a signal. The data uploaded to the openSenseMap is publicly available and can be explored through a web application or accessed for further analysis.

5 https://www.st.com/en/imaging-and-photonics-solutions/vl53l8cx.html (last visited 22.04.2024)

3. Prototype

As a demonstration of our proposed concept for a mobile, low-cost, and low-energy takeover detection system, we built a prototype, which entailed the assembly of hardware components and the development and deployment of its software. The latest iteration of the senseBox features an MCU based on an ESP32S2,[6] which boasts a single-core processor with 4 MB of flash memory and 320 KB of RAM, which is enough for small machine-learning applications.

To perceive passing traffic, the VL53L8CX multizone ToF sensor is chosen. The 'multizone' attribute indicates that the sensor measures distance not just at a single point but across 8x8 zones. Under perfect conditions without any ambient light whatsoever, the VL53L8CX can measure distances up to 400 cm. However, in regular ambient daylight, the maximum measurable distance is reduced to approximately 200 cm, which is just enough to detect when a car does not comply with the 150 cm minimum distance. Figure 2 provides illustrations of what the sensor detects compared to a camera capturing visible light.

a)

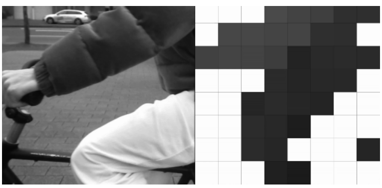

b)

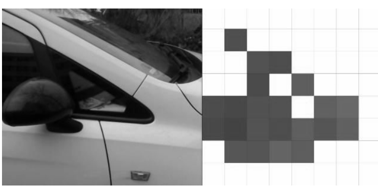

Figure 2. Two examples of what a normal webcam sees (top) and what the VL53L8CX sees (bottom). One of a cyclist (a) and one of a car (b), both at close distance

6 https://www.espressif.com/sites/default/files/documentation/esp32-s2_datasheet_en.pdf (last visited 22.04.2024)

To mount the sensor and the microcontroller on a bicycle, we 3D printed a simple case that can be attached to the carrier of a bicycle (see Figure 3). Alternatively, the case could also be attached underneath the saddle of the bicycle if it does not have a carrier.

Figure 3. Sensor (in black) in a 3D printed case together with the senseBox MCU S2 on a bicycle. The sensor is oriented to the left side of the bicycle.

As this prototype serves as a proof-of-concept rather than a finalised production system, the classification problem to be solved is deliberately simplified: sequences of distance recordings shall be either classified as takeover or no takeover, without any regard for distance, speed, or type of vehicle. To solve this problem, a simple RNN is trained and deployed.

Given the straightforward nature of the chosen classification problem, we maintain a simplistic structure for the neural network employed. The network comprises a sequence of only two layers: one consisting of long short-term memory (LSTM) neurons and the other of dense neurons. LSTM neurons are a type of recurrent neural network unit capable of capturing long-range dependencies in sequential data and retaining information over extended time periods. The network is configured to take a series of 20 frames as input for each classification task. Given that the sensor records data at a rate of 15 Hz, this equates to approximately 1.25 seconds of data, in which a takeover manoeuvre should fully take place.

For the purpose of training the proposed neural network, it was implemented in Python using TensorFlow,[7] a widely used open-source machine learning framework developed by Google. Given the high compatibility and support of the TensorFlow framework, the resulting trained model can be directly ported to and run on an MCU in the Arduino programming language. The model is first converted from the original TensorFlow version to a smaller and more efficient TensorFlow Lite version. This is then converted to a byte array and reconstructed in Arduino with TensorFlow Lite for Microcontrollers.

The network undergoes supervised training, relying on manually classified training data to learn effectively. As a matter of fact, a significant portion of the prototype's development effort was dedicated to recording and formatting suitable training data. We decided on a realistic approach for recording data by mounting the ToF sensor on a bicycle and driving around the streets of Münster. Data collection was impeded by the already well-developed bicycle infrastructure in this city, leaving only a few suitable streets with flowing traffic but without dedicated bike lanes. Data was recorded in those places by driving up and down the street during peak rush hour. To simplify data collection and not further endanger the cyclist recording the data, it was chosen to augment data by driving by parked cars and reversing the recorded video. Additionally, sequences in which other cyclists passed at short range were also considered valid takeovers.

A normal webcam was also mounted on the bicycle and configured to record simultaneously with the ToF. Using timestamps, the webcam videos were later used to label the sequences of depth frames from the ToF sensor as takeover or no takeover. The resulting labelled data was then divided into subsets for training, validating, and testing the neural network.

The data and the code for training and deployment are openly accessible on GitHub, along with more detailed descriptions and explanations of the procedure: https://github.com/TinyAIoT/tof-takeover-detection.

7 https://www.tensorflow.org/ (last visited 22.04.2024)

4. Applicability

The ultimately trained model performs incredibly well on the recorded data, with an accuracy of 99% or 95.5% for the model converted to TensorFlow Lite. However, considering that the recorded no-takeover sequences lack variety (they are mainly composed of empty frames), this is not immediately translatable to positive performance in the real world. The trained model struggles, especially by classifying many false positives, meaning that it detects takeovers when there are none. Nonetheless, the results are valuable as a proof-of-concept of the proposed method.

The size of the model comes out at a little over 15 KB, while memory reservations of a little over 8 KB are necessary to store the data going through the model. Meanwhile, its inference time on the senseBox based on ESP32S2 reaches 26 microseconds on average. Due to this incredibly low inference time, the model can easily be applied to each new sequence of frames in real time. The VL53L8CX ToF sensor can, at maximum, be configured to 15Hz, meaning it records a new frame every 66.7 milliseconds. These frames are continuously recorded in a ring buffer, and the most recent 20 frames are fed to the model whenever a new frame comes in.

Regarding the overall applicability of the VL53L8CX sensor, this study offers a very promising outlook on its performance. While manually reviewing the recordings that were made with the sensor attached to a bike, discernible differences among passing vehicles – whether cars, bicycles, scooters, or pedestrians – become immediately evident. These differences in dimensions and textures, clearly perceptible to the human eye, suggest the feasibility of similar distinctions by neural networks. Consequently, the cost-effective time-of-flight technology has the potential to contribute to traffic technologies. Although the lighting conditions impaired the quality to a certain extent, the captured images still resulted in clearly recognisable structures of vehicles and other objects. However, a notable issue persists in accurately capturing dark or highly reflective objects. Dark-coloured vehicles usually only yield patchy images containing mostly noise, a recognised limitation of ToF sensors (Baek et al., 2020; Zieliński, 2021).

In terms of future work, the available memory and RAM on the senseBox MCU leave room to increase the model size and complexity. It could, for example, be interesting to add convolutional layers to the network. Measurements of the accelerometer onboard the senseBox MCU could also aid the model performance and enable the determination of the speed of both the cyclist and passing vehicle at the moment of takeover. Priority for improving the model performance should, however, be the enhancement of the training data. Primarily, efforts should focus on diversifying the dataset by incorporating a wider array of non-takeover sequences to mitigate the prevalent issue of false positives. Additionally, existing data can be enriched through artificial augmentation techniques, such as introducing random noise, translating or shifting sequences, or adjusting the speed of recordings. Such augmentation strategies generally contribute to the robustness and generalisation capabilities of a model.

The georeferencing, uploading, and provision of collected data was not implemented in the prototype but has already been deployed in the senseBox:bike project, as previously mentioned. The existing workflow can be largely mirrored for our use case, with the exception that the interface of the smartphone app would have to be adjusted in line with the uploaded data. Depending on the detection solution adopted, uploaded data could, for example, include the distance, time of incident, type of vehicle, and also the current speed of the cyclist and passing vehicle, in case readings from the onboard accelerometer are added. To facilitate further analysis and provide additional visualisation tools, it would also be interesting to develop our own web interface tailored to the specific data instead of using the openSenseMap.

5. Conclusion

In this work, we present a concept and an initial prototype demonstrating the detection of dangerous overtaking manoeuvres using low-cost ToF sensor technology combined with machine learning methods. By employing ToF technology instead of visible light cameras and running the classification on-device, we effectively reduce the size of collected data and communication, thus saving energy. The two-fold significance of sustainability in this work stems from the fact that the proposed system is motivated to support a mobility transition towards more sustainable forms of transportation overall but also intrinsically aims for sustainability through the energy and resource-efficient, on-device processing of data.

The trained model showcases its capability to generate valuable classifications of takeover manoeuvres. It proves the feasibility of the concept while demonstrating the necessity for further collection of training data to offer a more robust solution. An unsolved challenge persists in detecting black or highly reflective cars with the chosen proximity sensor. Previous projects with the senseBox have already paved the way for georeferencing and uploading data, which have not yet been implemented in this described prototype. When combined with location information, the data can show spatial and temporal accumulations of dangerous overtaking manoeuvres, and consideration can then be given to adjusting the local cycling infrastructure. Ultimately, the result of this work is not intended to represent a final system but rather to serve as an essential preliminary step for further research and development with the knowledge gained and the software developed.

Looking ahead, our work holds promise for applications in the field of traffic and transport planning. For instance, the detected and localised dangerous overtaking manoeuvres allow for the identification of high-risk areas on roads, which can then be addressed by local authorities through targeted infrastructure improvements such as adding physical barriers between bicycles and cars on shared road segments. By analysing broad datasets collected by citizens, traffic planners can gain valuable insights into traffic patterns and the frequency of dangerous manoeuvres, allowing for data-driven decisions to enhance road safety. This participatory approach not only increases the amount of data available for analysis but also fosters community engagement in traffic safety initiatives.

Further use of devices applying our approach can significantly impact the field of education. By employing the device in workshops and initiatives in schools, students can learn about machine learning and sensor technology through hands-on experience. Schools can participate in data acquisition and analysis projects, enabling students to contribute to real-world datasets as part of citizen science initiatives. This involvement promotes spatial citizenship, as described by Gryl (2012), by empowering students to understand and influence their local environment. This educational approach not only enriches the learning experience but also instils a sense of civic responsibility and technological literacy in young learners.

Acknowledgements

This research has been partly funded by the Federal Ministry for the Environment, Nature Conservation, Nuclear Safety and Consumer Protection through the funding programme KI-Leuchttürme in the project 'TinyAIoT Energie und ressourceneffiziente künstliche Intelligenz für moderne IoT-Anwendungen' and partly by the Federal Ministry for Education and Research through the funding programme EUREKA-Globalstars-Kooperation

mit Brasilien in the project ‘atrAI Bikes Sichere Fahrradstädte durch Datenwissenschaft und intelligente Fahrräder’.

References

Baek, E.-T., Yang, H.-J., Kim, S.-H., Lee, G., & Jeong, H. (2020). Distance error correction in time-of-flight cameras using asynchronous integration time. *Sensors*, *20*(4), 1156–1156. https://doi.org/10.3390/s20041156

Bartoschek, T., Fehrenbach D., Fehrenbach T., Pesch M., & Steinmann, L. (2019). *Das senseBox-Buch*. dpunkt.verlag

Blees, V. (2021) Sicher Überholt. Radverkehrsinfrastruktur – Baustein Der Verkehrswende. Gemeinsame Abschlusspublikation Des NRVP-Forschungsbegleitkreises ‘Förderschwerpunkt Infrastruktur’, 24–31.

Lehmann, H., Meidinger, D., Wittlich, H., & Gegg, M. (2018, August 20). Fahrradverkehr in Berlin: Das Projekt Radmesser. *Der Tagesspiegel Online*. https://www.tagesspiegel.de/gesellschaft/medien/das-projekt-radmesser-3980056.html.

Statistisches Bundesamt (Destatis). (2022, January 31). Erwerbstätige nach Stellung im Beruf, Entfernung, Zeitaufwand und benutztem Verkehrsmittel für den Hinweg zur Arbeitsstätte 2020 in %. https://www.destatis.de/DE/Themen/Arbeit/Arbeitsmarkt/Erwerbstaetigkeit/Tabellen/pendler1.html.

Gryl, I., & Jekel, T. (2012). Re-centering geoinformation in secondary education: Toward a spatial citizenship approach. *Cartographica*, *47*(1), 18–28. https://doi.org/10.3138/carto.47.1.18.

Hauenstein, J., Eckart, J., Zeile, P., & Merk, J. (2023). The effect of overtaking distances on the stress occurrence of cyclists in urban areas. *LET IT GROW, LET US PLAN, LET IT GROW. Nature-based Solutions for Sustainable Resilient Smart Green and Blue Cities. Proceedings of REAL CORP 2023, 28th International Conference on Urban Development, Regional Planning and Information Society* (pp. 699–708). https://doi.org/10.48494/REALCORP2023.9045.

Jurczok, F., Gensheimer, T., & Specht, F. (2023). *Fahrrad-Monitor Deutschland 2023*. Sinus Marktund Sozialforschung GmbH. https://bmdv.bund.de/SharedDocs/DE/Anlage/StV/fahrradmonitor-langfassung.pdf?__blob=publicationFile.

Nelson, T. A., Denouden, T., Jestico, B., Laberee, K., & Winters, M. (2015). BikeMaps.Org: A global tool for collision and near miss mapping. *Frontiers in Public Health*, 3. https://doi.org/10.3389/fpubh.2015.00053.

Nobis, C. (2019). *Mobilität in Deutschland – MiD Analysen Zum Radverkehr Und Fußverkehr*.infas, DLR, IVT and infas 360 on behalf of the German Ministry of Transport. https://www.mobilitaet-in-deutschland.de/archive/pdf/MiD2017_Analyse_zum_Rad_und_Fussverkehr.pdf

Pfeil, M., Bartoschek, T., & Wirwahn, J. A. (2015). OPENSENSEMAP: A citizen science platform for publishing and exploring sensor data as open data. *Free and Open Source Software for Geospatial (FOSS4G) Conference Proceedings*, *15*(39). https://doi.org/10.7275/R56971SW.

Pyhel, J. (2022). *Studie Zum Weltfahrradtag: Sicherheitsbedenken Halten Menschen Vom Radfahren Ab*.Ipsos GmbH. https://www.ipsos.com/sites/default/files/ct/news/documents/2022-05/Ipsos-PI_Cycling_2022-05-24.pdf.

Storme, T., Benoit, S., Van De Weghe, N., Mertens, L., Van Dyck, D., Brondeel, R., Witlox, F., Zwartjes, L., & Cardon, G. (YEAR). Citizen science and the potential for mobility policy – Introducing the bike barometer. *Case Studies on Transport Policy*, *10*(3), 1539–1549. https://doi.org/10.1016/j.cstp.2022.05.013.

Wirwahn, J. A., & Bartoschek, T. (2015). Usability Engineering For Successful Open Citizen Science. *Free and Open Source Software for Geospatial (FOSS4G) Conference Proceedings*, *15*(1),. https://doi.org/10.7275/R54J0C9W.

Witte, V., Schwering, A., Bartoschek, T., & Pesch, M. (2023). Zukunftsweisender MINT-Unterricht Mit Dem senseBox-Ökosystem Die Plattform Für Partizipative Data Science Mit Physical Computing. *MNU Journal / Verband Zur Förderung Des MINT-Unterrichts*, *76*(4), 296–301.

Zieliński, B. (2020). A comparison of proximity sensors for a bicycle-to-car distance rangefinder. *International Journal of Electronics and Telecommunications*, 277–282. https://doi.org/10.24425/ijet.2021.135976.