

2023-06-18

# BU-CVKit: extendable computer vision framework for species independent tracking and analysis

---

M. Patel, L. Carstensen, Y. Gu, M. Hasselmo, M. Betke. 2023. "BU-CVKit: Extendable Computer Vision Framework for Species Independent Tracking and Analysis" <https://doi.org/10.48550/arXiv.2306.04736>  
<https://hdl.handle.net/2144/48258>

*"Downloaded from OpenBU. Boston University's institutional repository."*

# BU-CVKit: Extendable Computer Vision Framework for Species Independent Tracking and Analysis

Mahir Patel<sup>1,2</sup>, Lucas Carstensen<sup>2</sup>, Yiwen Gu<sup>1</sup>, Michael Hasselmo<sup>2</sup>, and Margrit Betke<sup>1</sup>  
<sup>1</sup> Department of Computer Science and <sup>2</sup> Center for Systems Neuroscience  
 Boston University

## Abstract

*A major bottleneck of interdisciplinary computer vision (CV) research is the lack of a framework that eases the reuse and abstraction of state-of-the-art CV models by CV and non-CV researchers alike. We present here BU-CVKit, a computer vision framework that allows the creation of research pipelines with chainable Processors. The community can create plugins of their work for the framework, hence improving the re-usability, accessibility, and exposure of their work with minimal overhead. Furthermore, we provide MuSeqPose Kit, a user interface for the pose estimation package of BU-CVKit, which automatically scans for installed plugins and programmatically generates an interface for them based on the metadata provided by the user. It also provides software support for standard pose estimation features such as annotations, 3D reconstruction, re-projection, and camera calibration. Finally, we show examples of behavioral neuroscience pipelines created through the sample plugins created for our framework.*

## 1. Introduction

Computer Vision has the potential to become an integral part of interdisciplinary research. The advances in sub-fields such as Pose Estimation, Object Detection, Segmentation, and 3D Reconstruction can directly impact research in sciences as diverse as Conservation Ecology, Neuroscience, Physiology, Psychology, and many more. Integration of state-of-the-art computer vision with applied sciences can be supported by open-source python packages or wrappers like OpenCV [4], TensorFlow [1], PyTorch [13], Scikit-learn [15], OpenPose [5], MMPose [7], DeepLabCut [12], and SLEAP [16] with different levels of abstraction and complexity. However, researchers in the applied sciences usually have to rely on themselves to implement or, at best, adapt computer vision methods for their research pipeline. There is an absence of an extendable framework that provides a high-level abstraction for accessing the

methods implemented through these open-source software. Such a framework could allow for easier integration of these methods in interdisciplinary pipelines, thus benefiting non-computer-vision specialists on the one hand and also providing computer vision researchers the tools to further their research on the other.

We present here BU-CVKit framework, which bridges the accessibility gap to state-of-the-art computer vision research for researchers from diverse backgrounds and application disciplines. The core idea behind BU-CVKit is to provide chainable modules that can be used to create abstract research pipelines which can be easily modified and shared with the community. We also present MuSeqPose Kit, which provides an intuitive user interface to the pose estimation sub-module of the framework.

We demonstrate the potential of our framework by implementing plugins for two state-of-the-art 2D/3D pose estimation methods, DeepLabCut [12] and OptiPose [14]. DeepLabCut is a widely-used feature-rich framework that provides state-of-the-art markerless 2D pose estimation of animals. OptiPose is a denoising auto-encoder that encodes postural dynamics to optimize coarse 3D poses. We use the plugins to create standard behavioral neuroscience pipelines.

## 2. BU-CVKit Framework

BU-CVKit is an extendable framework that provides standard functionalities such as efficient input/output, evaluation metrics, geometric transformations, camera calibration, multi-view reconstruction, and other abstractions.

To illustrate how BU-CVKit can reduce programming overhead and thus potentially accelerate computer vision research, we give the following pipeline example. We denote a 2D pose estimation method as function  $f$ , a 3D reconstruction method as function  $g$ , a 3D pose filtering method as function  $h$ , and a data analysis method as function  $i$ . BU-CVKit enables the user to design a research pipeline that outputs

$$o = i(h(g(f(data)))). \quad (1)$$

Furthermore, consider a single-step 3D-pose estimation method  $j$ , where  $j \approx g * f$ . A user can replace  $f$  and  $g$  without affecting the semantics or flow of the pipeline and can thus explore different methods.

In the remainder of this section, we discuss the three major modules of the framework.

### 2.1. Input/Output Modules

Efficient and intuitive input/output is necessary for the application domains of computer vision. We here describe two types of input/output modules. First, BU-CVKit contains an abstract buffered VideoReader class that can be extended to provide sequential or random access to video frames using different backbone libraries. With our package, we provide buffered implementation of OpenCV [4], Deffcode [17], Decord, and an Image plugin, which supports reading a directory of images and providing them as a video stream to support the data format adopted by the datasets. We compared the throughput of the standard OpenCV video reader versus the buffered modules of BU-CVKit under different CPU loads reading high-resolution video frames, see Table 1. The buffered module performs better than the standard OpenCV video readers and the improvement is much more evident when the CPU is idle. This performance improvement may be beneficial for scientists who need to review videos of animal behavior and their analysis in real-time. Users can also create their own IO plugins, perhaps one that utilizes GPU codecs to achieve even better performance.

The second I/O Module we describe here is the pose estimation data reader module. It provides an abstract DataReader class that can be extended to provide intuitive access to 2D or 3D pose estimation data. Each instance of a person’s or animal’s pose is converted into a Skeleton object which in turn contains multiple Part Objects. Each Part object extends a Numpy array and therefore supports efficient vectorized operations. In addition, the Skeleton object supports further Pose-Estimation-specific features such as behavior annotations and unified arithmetic and geometric

Table 1. Average frames-per-second throughput of baseline OpenCV implementation and different BU-CVKit buffered implementations, when computed over 1,000 1,024x1,024 H264-encoded frames under two CPU work modes on a 6-core Xeon processor.

VideoReader	Throughput [fps]	
	CPU Load	CPU Idle
OpenCV VideoCapture(original)	47.23	162.07
CVReader (buffered OpenCV)	48.59	193.05
Deffcode	49.30	193.09
Decord	64.71	193.42

operations.

To support input/output operations with a wide area of toolkits, the module includes implementations for reading, writing, and translating flattened CSV files, our custom  $n$ -dimensional CVKit files for storing pose data, HDF5 [8] files, Mat [10] files, and 2D DeepLabCut [12] files.

### 2.2. Processor Modules

The pose estimation package provides an abstract Processor class that can be extended to implement plugins for state-of-the-art computer vision methods. The instances of these Processors are chainable and, therefore, can be used to create a pipeline that takes raw data and generate the desired output. The Processors can be classified into three categories - filters, generative, and utility. The filter package contains the Processors that denoise the input data. The included dimension-independent filters are a constant acceleration Kalman filter, linear interpolation, a statistical distance filter, a moving average filter, and a velocity filter. The generative module contains the Processors that generate new data from the provided input. This module includes 3D reconstruction and reprojection processors, a kinematics generator, and plot generators. Finally, the utility module contains a file loader, a file saver, and an input statistics generator. Although these Processors are not directly used in processing or analyzing data, they are required to facilitate chaining and other utilities.

### 2.3. Pose Estimation Utility Modules

The framework provides several pose estimation utility modules. The camera calibration is a standard first step in any pose estimation work. Therefore we provide a calibration module that uses EasyWand [18] to generate Direct Linear Transformation (DLT) [2, 19] coefficients and camera parameters. The 3D reconstruction and reprojection modules use DLT to estimate the 3D positions of matching 2D keypoints and mapping 3D coordinates into 2D camera planes. The metrics module for pose estimation provides the Mean Per Joint Position Error (MPJPE) [11] and dynamic percentage of correct keypoints (PCK@x) [11] metrics. Finally, the geometric transformations module enables 3D rotation, translation, and axes alignment.

## 3. Plugin Sub-system

The abstract code design of the CVKit enables researchers to extend the framework by creating plugins for their own work. The Processor class described in the previous section acts as the main entry point for plugins. We here showcase the potential of the plugin system by extending the capabilities of BU-CVKit through our three plugins. Our DeepLabCut [12] plugin provides a Processor for running a DeepLabCut model on videos and providing 2D pose predictions. Our OptiPose [14] processor takes in the

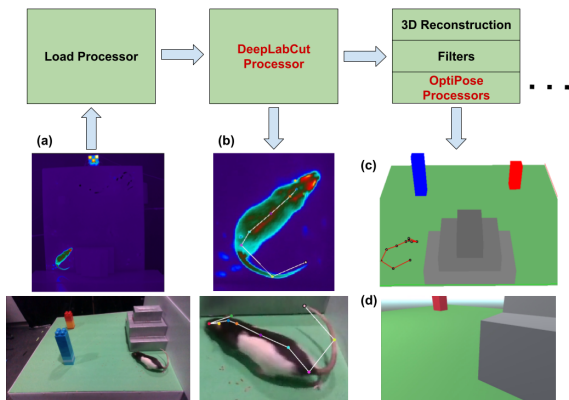


Figure 1. Chainable Processors: A Processor pipeline using the DeepLabCut [12] and OptiPose [14] plugins. (a) The raw input thermal and RGB images. (b) Predicted 2D poses generated by the DeepLabCut plugin. (c) Optimized 3D poses generated by the pose-optimization Processor from the OptiPose plugin. (d) A simulated binocular egocentric view generated by the viewpoint reconstruction Processor from the OptiPose plugin.

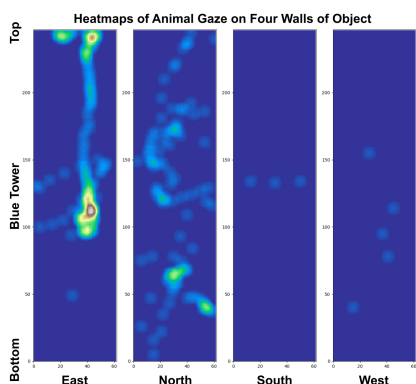


Figure 2. Gaze heatmaps computed by a Generative plugin that uses ray tracing with radially attenuated focus to approximate the gaze of the animal on the blue tower shown in Fig. 1, for each wall of the tower (the west wall of the blue tower faces the red tower). The dimensions of the heatmaps correspond to the dimensions of the walls, i.e., 62 mm (width) and 247 mm (height). The rodent most often looked at the center of the East wall (hot colors).

3D poses generated by the reconstruction Processor and optimizes them through a pre-trained OptiPose model. It also supports generating a view direction vector that can be fed into a game engine to generate a simulated egocentric view of the subject. The resulting pipeline created by BU-CVKit is shown in Fig. 1. It starts with the DeepLabCut plugin to compute 2D predictions, the reconstruction to perform 3D Reconstruction, a series of filters, pose optimization through an OptiPose model, and viewpoint simulation.

The NeuroAnalysis plugin processes the tracking data generated from the OptiPose plugin (not shown in Fig. 1). It provides a single-ray-tracing Processor that is used by other

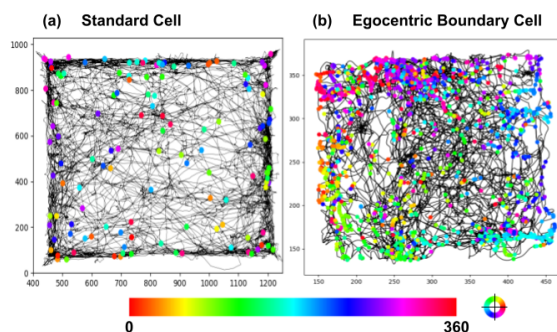


Figure 3. Results of a Generative plugin that combines neural recording with animal tracking to plot cellular activity and their corresponding spatial locations. The black lines indicate the trajectory of the rodent exploring an area (as in Fig. 1 but without objects), and each dot is the location where the cell fired. The color of the dot represents the head direction of the rodent at the time the cell fired. (a) Example of non-deterministic cell firing. (b) Example of an Egocentric Boundary Cell firing.

behavioral analysis generators. With the plugin, the gaze direction of the animal can be analyzed by tracing the view direction vector to surfaces in the animal's environment. Here we make the assumption that the rodent is generally looking at the direction it is facing, i.e., where its snout is pointing. We also model the rodent's attention focus, where the attention score attenuates radially outwards from the assumed gaze point. With these tools, we can localize the rodent's approximate focus of interest on the four walls of the blue tower, shown in Fig. 1. Summative analysis then leads to heat maps that highlight to which areas on the tower the animal seemed to have paid attention most often, see Fig. 2.

The EBC-Processor helps detect Egocentric Boundary Cells (EBC), neurons that exhibit increased firing when a boundary is at a specific distance and direction relative to the subject [3]. Analytical plots, as in Fig. 3, are generally used to distinguish non-EBC patterns from EBCs. Each plot is produced for a specific cell. The plot on the left in Fig. 3 shows no discernible directional or spatial pattern in the cell activity, whereas the plot on the right exemplifies a firing pattern when the boundary is on the left of the rodent.

Animal behavioral analysis studies how an animal explores an environment under a spectrum of novel conditions. The amount of time the rodent spends occupying certain regions and the instances of rearing are used as indicators of exploration and memory [6, 9]. The Occupancy Processor can generate 2D plots that show the distribution of where the time was spent globally as well as locally near the objects of interest (Fig. 4). Similarly, the Rearing Processor detects rearing instances by monitoring selected keypoints to generate 3D bar plots of rearing frequency (Fig. 4).

We stress that a major advantage of the framework architecture is that the processors are easily replaceable. The DeepLabCut Processor and the OptiPose Processor can be

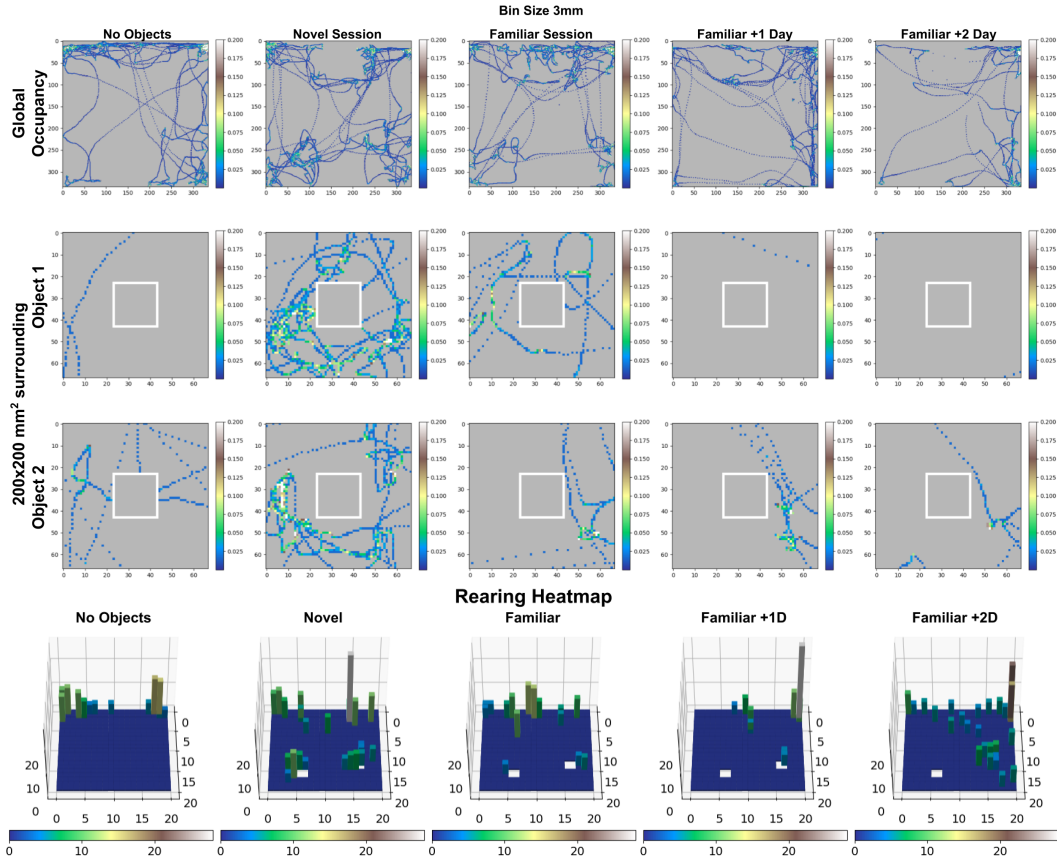


Figure 4. Top: Results of a behavioral analysis plugin that generates Occupancy Maps for comparing the change in behavior corresponding to the novelty of objects. Bottom: Results of a behavioral analysis plugin that generates "rearing" plots to compare exploratory behavior under novel or familiar conditions.

replaced by any 2D pose estimator and 3D pose optimizer, respectively, without affecting the neuro-processors. They can also be bypassed entirely by creating a plugin for one of the single-stage 3D pose estimators. This allows researchers to try different methods without drastically changing their analysis pipeline.

#### 4. MuSeqPose Kit

We provide MuseqPose Kit, a user interface to the pose-estimation sub-package of BU-CVKit. It provides a powerful video annotation widget that supports interpolating annotations across frames. The widget also provides a reprojection toolbox to mitigate the need for annotating every view. The camera calibration widget automatically picks a diverse set of synchronized annotated frames and generates the necessary files for the EasyWand Calibration package. The synchronized playback widget allows for displaying videos with real-time plots. Finally, the Pipeline widget automatically scans for the installed plugin metadata and generates a user interface to access the underlying Processors. Using this widget, users can create, visualize, execute,

and save their research pipelines using this widget.

#### 5. Conclusions

With BU-CVKit Framework, we have provided an extendable framework for simplifying and accelerating interdisciplinary research. We showcase the usability of such a framework by implementing plugins for state-of-the-art computer vision research and using them for behavioral neuroscience. We plan to add extendable Segmentation and Object Detection packages. The plug-and-play aspect of the Processors enables researchers to try different methods without extensively modifying their pipeline. Finally, since non-computer-science researchers may have a varying range of programming skills, we provide a powerful user interface that automatically adapts to the installed plugins while providing pose estimation features.

## References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. **1**
- [2] Y. I. Abdel-Aziz and H. M. Karara. Direct linear transformation from comparator coordinates into object space coordinates in close-range photogrammetry. *Proceedings of the Symposium on Close-Range Photogrammetry*, pages 1–18, 1971. **2**
- [3] Andrew S. Alexander, Lucas C. Carstensen, James R. Hinman, Florian Raudies, G. William Chapman, and Michael E. Hasselmo. Egocentric boundary vector tuning of the retrosplenial cortex. *Science advances*, 6(8):eaaz2322, 2020. **3**
- [4] G. Bradski. The OpenCV Library. *Dr. Dobb’s Journal of Software Tools*, 2000. **1, 2**
- [5] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. **1**
- [6] Lucas C. Carstensen. *Neural representations of environmental features in retrosplenial cortex and 3-dimensional reconstruction of animal pose*. PhD thesis, 2023. **3**
- [7] MMPose Contributors. Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose>, 2020. **1**
- [8] The HDF Group. Hierarchical data format version 5, 2000–2010. **2**
- [9] Lindsey M. Lueptow. Novel object recognition test for the investigation of learning and memory in mice. *Journal of Visualized Experiments: JoVE*, 55718:126, 2017. **3**
- [10] The MathWorks Inc. Matlab, 2023. **2**
- [11] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3D Vision (3DV), 2017 Fifth International Conference on*, IEEE, 2017. **2**
- [12] Tanmay Nath, Alexander Mathis, An Chi Chen, Amir Patel, Matthias Bethge, and Mackenzie Weygandt Mathis. Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *bioRxiv*, 2018. **1, 2, 3**
- [13] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. **1**
- [14] Mahir Patel, Yiwen Gu, Lucas C. Carstensen, Michael E. Hasselmo, and Margrit Betke. Animal pose tracking: 3d multimodal dataset and token-based pose optimization. *International Journal of Computer Vision*, 131(2):514–530, 2022. **1, 2, 3**
- [15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. **1**
- [16] Talmo D. Pereira, Nathaniel Tabris, Arie Matsliah, David M. Turner, Junyu Li, Shruthi Ravindranath, Eleni S. Papadoyannis, Edna Normand, David S. Deutsch, Z. Yan Wang, Grace C. McKenzie-Smith, Catalin C. Mitelut, Marielisa Diez Castro, John D’Uva, Mikhail Kislin, Dan H. Sanes, Sarah D. Kocher, Samuel S.-H. Wang, Annegret L. Falkner, Joshua W. Shaevitz, and Mala Murthy. SLEAP: A deep learning system for multi-animal pose tracking. *Nature Methods*, Apr. 2022. **1**
- [17] Abhishek Singh Thakur. abhitronix/deffcode: v0.2.4, Oct. 2022. **2**
- [18] D. H. Theriault, N. W. Fuller, B. E. Jackson, E. Bluhm, D. Evangelista, Z. Wu, M. Betke, and T. L. Hedrick. A protocol and calibration method for accurate multi-camera field videography. *The Journal of Experimental Biology*, 217:1843–1848, Feb. 2014. Open access online, [pdf](#). **2**
- [19] Kwon Young-Hoo. Dlt method. <http://www.kwon3d.com/theory/dlt/dlt.html>. Accessed April 28, 2023. **2**