
On Coordinate Decoding for Keypoint Estimation Tasks

Anonymous Author(s)

Affiliation

Address

email

Reproducibility Summary

1

2 Scope of Reproducibility

3 A series of 2D (and 3D) keypoint estimation tasks are built upon *heatmap* coordinate representation, *i.e.* a probability
4 map that allows for learnable and spatially aware encoding and decoding of keypoint coordinates on grids, even allowing
5 for sub-pixel coordinate accuracy. In this report, we aim to reproduce the findings of DARK [1] that investigated the 2D
6 heatmap representation by highlighting the importance of the encoding of the ground truth heatmap and the decoding of
7 the predicted heatmap to keypoint coordinates. The authors claim that **a**) a more principled distribution-aware coordinate
8 decoding method overcomes the limitations of the standard techniques widely used in the literature, and **b**), that the
9 reconstruction of heatmaps from ground-truth coordinates by generating accurate and continuous (non-quantized)
10 heatmap distributions lead to unbiased model training, contrary to the standard coordinate encoding process that
11 quantizes the keypoint coordinates on the resolution of the input image grid.

12 Methodology

13 To reproduce DARK, we thoroughly studied and followed the official implementation provided by the authors. We
14 partially used the publicly available code and integrated it to a model development kit to accelerate the implementation
15 and execution of our experiments. We conducted the experiments on a recent human pose and depth dataset aiming to
16 validate the claimed concepts on different a different setting than the one originally used and assessed. Our experiments
17 were conducted on an Nvidia 2080 Ti 12 GB GPU with an average training time of 20 hours.

18 Results

19 We were able to reproduce the findings of the original work. The proposed distribution-aware sub-pixel coordinate
20 decoding improves the keypoint estimation performance, outperforming existing, standard decoding techniques. The
21 elimination of coordinate quantization after re-scaling to the down-scaled input image resolution before the coordinate
22 encoding (heatmap reconstruction), indeed leads to unbiased model training. Going a step beyond, we assess the
23 performance of other sub-pixel decoding approaches, such as *Center of Mass* [2, 3], which also show increased
24 performance.

25 What was easy

26 The DARK paper is well-structured and -presented aiding reproducibility. Further, the official code of the paper is
27 publicly available, ensuring that the method and its results are reproducible. Given that DARK serves as a model-
28 agnostic head that relies on post-inference processing only during testing, it is overall straightforward to implement and
29 verify its effectiveness.

30 What was difficult

31 A potential issue has to do with the availability of new datasets, beyond the ones used in the paper, that provide
32 coordinates in higher resolutions.

33 Communication with original authors

34 No communication was required, as the implementation was openly available and the verification is straightforward.

1 Introduction

Keypoint coordinate estimation in N -dimensional grids is a common task related to a variety of computer vision domains such as human and object pose estimation, face tracking, crowd scene analysis and beyond. For modern data-driven methods, the heatmap representation, *i.e.* a N -dimensional probability distribution (Gaussian in most cases) centred at each keypoint coordinates, is considered the *de facto* coordinate representation for keypoint estimation. Heatmap representations enable the use of efficient fully N -dimensional convolutional networks, learning dense spatial features around the ground-truth location, and lending to a more natural way of estimation coordinates.

Despite the popularity of heatmap coordinate representation, DARK is the first research work that investigates heatmaps in depth, meaning that the authors explore and assess the standard *coordinate-to-heatmap encoding* and the *heatmap-to-coordinate decoding* techniques present in the literature, and formulate and assess a novel approach for distribution-aware coordinate representation.

2 Scope of reproducibility

We aim to reproduce the Distribution-Aware coordinate Representation of Keypoints (DARK) [1] method to assess its benefits in keypoint coordinate estimation. In particular, we reproduce and assess DARK in 2D human pose estimation with the use of state-of-the-art fully convolutional networks, such as *Stacked Hourglass* [4] and *HRNet* [5], similarly to the original work. A recent dataset, HUMAN4D [6], including multi-view depth data and 3D poses captured with the use of a high-quality VICON Motion Capture [6] setup, is used to train our models from scratch.

In a nutshell, the main claims of DARK paper, and those that we seek to validate, are summarised below:

- Claim #1: *Efficient Taylor-expansion based coordinate decoding improves coordinate estimation accuracy.*
- Claim #2: *Unbiased sub-pixel centred coordinate encoding results in optimal supervision and upgraded model performance.*

3 Methodology

Our approach in reproducing the work was to carefully re-implement every single step of the approach, following the presented methodology as well as studying the code of the official repository of the authors¹. Beyond that, the code that we used as is from the authors' implementation is the Taylor-expansion based coordinate decoding. The rest of the implementation for all our experiments was developed in a model development kit [7], oriented towards reproduction.

The code and its documentation are submitted and published along with this report.

3.1 Models

Following the original paper, we selected and used the following models for predicting coordinate heatmap distributions in images:

- *Stacked Hourglass (SH)* [4]; with feature maps of 128 width and 4 stages and 18 output heatmap layers.
- *HRNet* [5]; with feature maps of 32 width and 4 stages and 18 output heatmap layers. The 2nd, 3rd and 4th stages contain 1 exchange block containing 4 residual units each.

3.2 Coordinate representation

3.2.1 Coordinate decoding

We consider a heatmap \mathbf{H}^j that encodes the coordinates of each j -joint 2D position as predicted by our trained models. In our experiments, we decode the coordinate from \mathbf{H}^j with 4 different decoding methods. In the following description, N_u and N_v are the width and height of the heatmap, and $\mathbf{p} \in \Omega$ represents a pixel coordinate in the heatmap domain Ω .

- *ArgMax*; by directly using the maximal activation location m by:

$$m = (u, v)_j = \arg \max_{\mathbf{p} \in \Omega} (\mathbf{H}^j(\mathbf{p})) \quad (1)$$

¹<https://github.com/ilovepose/DarkPose/tree/612fad594eddd022b8a162a2c7274f9ee8d06d2c>

74 • **Standard**; by slightly shifting (sub-pixel shifting) the estimated coordinates between the maximal m and
 75 second maximal s activation by:

$$p = (u, v)_j = m + 0.25 \cdot \frac{s - m}{\|s - m\|_2} \quad (2)$$

76 This method constitutes one of the standard coordinate decoding techniques that predicts the maximal activation
 77 with a greedy sub-pixel shifting (0.25 pixels) towards the second maximal activation in the heatmap space Ω .

78 • **Taylor-expansion (DARK)**; exploring the distribution structure of the predicted heatmap to infer the underlying
 79 maximum activation assuming the predicted heatmap signal follows a 2D Gaussian distribution, similarly to
 80 the ground-truth heatmap during supervision by:

$$\mu = (u, v)_j = m - (\mathcal{D}''(m))^{-1} \mathcal{D}'(m) \quad (3)$$

81 where $\mathcal{D}'(m)$ and $\mathcal{D}''(m)$ denote the first and second derivative (i.e. Hessian) of the predicted heatmap after
 82 applying logarithmic transformation.

83 • **CoM**; considering heatmap \mathbf{H}^j as a 2D grid of point masses determined by the corresponding heatmap pixel
 84 values, we calculate one single 2D point whose mass is the total mass of the grid by:

$$c = (u, v)_j = \left(\frac{1}{N_u}, \frac{1}{N_v} \right) \circ \sum_{\mathbf{p} \in \Omega} \mathbf{H}^j(\mathbf{p}) \cdot \mathbf{p} \quad (4)$$

85 where \circ denotes element-wise multiplication. CoM constitutes by nature a sub-pixel coordinate decoding
 86 method that is not evaluated in the original paper.

87 3.2.2 Coordinate encoding

88 The second claim is about coordinate encoding where the existing standard methods down-sample the original input
 89 images into the model input resolution, quantizing the keypoint coordinates accordingly (**w/ quant**). Thus, the ground-
 90 truth heatmaps reconstructed based on the quantized coordinates are erroneous leading to sub-optimal supervision and
 91 degraded model performance.

92 Instead, DARK authors propose the heatmap generation placing the heatmap centre at the unbiased, non-quantized
 93 keypoint coordinate (**w/o quant**), claiming that, independently of the network architecture and the coordinate decoding
 94 method, the models perform consistently better.

95 3.3 Datasets

96 3.3.1 HUMAN4D

97 We investigate the benefits that DARK introduces in heatmap-based keypoint estimation by conducting our experiments
 98 on a recent dataset, HUMAN4D, that contains, among other data, pose annotated depth maps captured with the use of
 99 recent consumer-grade depth sensing devices. Beyond the experimentation on a dataset different from the ones used in
 100 the original work, the rationale behind this selection is twofold:

- 101 • The dataset provides the original 3D poses allowing for their projection as non-quantized 2D coordinates on
 102 the 2D depth map grids and, thereby, enabling the comparison between **w/ quant** and **w/o quant** coordinate
 103 encoding.
- 104 • The training and evaluation of baseline pose estimation data-driven models on depth data which are relatively
 105 limited in comparison with models trained on color data.

106 The activities included are the following: *running, jumping_jack, bending, punching_n_kicking, basketball_dribbling,*
 107 *laying_down, sitting_down, sitting_on_a_stool, talking, object_dropping_n_picking, stretching_n_talking, walk-*
 108 *ing_n_talking, watching_scary_movie* and *in-flight_safety_announcement*.

109 3.3.2 Preparation

110 With the use of HUMAN4D 3D poses, we project the joint 3D positions to the 2D depth map grids to obtain the 2D
 111 non-quantized poses. Using the single-person part of the HUMAN4D dataset, clipping the first 100 frames from each
 112 sequence (to exclude T-Pose calibration frames) and sub-sampling the remaining part of the sequence, we create a set of
 113 98, 864 single-view samples from 14 single-person activities, including all 4 depth maps per multi-view frame.

114 We use the data of the first three subjects performing all the activities for training, while we split the activities performed
 115 by the fourth subject (unseen) in two halves for validation and testing. We decide to split the data that way in order
 116 to assess the models on an unseen subject with different body structure, also keeping different activities between the
 117 validation and testing sets to provide fair conclusions with respect to the models’ performance. The training, validation
 118 and testing data sets consist of 73, 248, 12, 236 and 13, 380 samples, respectively.

119 Finally, beyond the application of center cropping on depth maps from 320×180 to 320×160 ($10px$ from top and
 120 bottom) to better fit to the multi-stage models, no further pre-processing steps were applied on the original data of the
 121 dataset.

122 3.4 Hyperparameters

123 All the models are trained for **30** epochs, picking for testing the best performing ones based on validation metrics across
 124 the epochs. Given that our aim is to assess the heatmap coordinate representation, not purely the effectiveness of the
 125 models in pose estimation, we use constant and common training hyperparameters for the various models, without
 126 further investigation for finetuning. To this end, both for *SH* and *HRNet*, we used the **Adam** optimizer with a learning
 127 rate of $2e - 3$, betas of values 0.9 and 0.999, without weight decay. The seed remained constant during all models’
 128 training, equal to 1314, ensuring reproducibility. We used simple MSE loss as in the original work that easily drove all
 129 models to convergence.




Figure 1: Interactive plots depicting the progress of the models’ training and the validation metrics in the form of a local web page can be found in the supplementary material of the report.

130 3.5 Experimental setup and code

131 As already noted, carefully following the authors’ description and code, we built our implementation upon *moai*
 132 [7], a PyTorch Model Development Kit for rapid development of Deep Learning workflows². In this framework,
 133 we implemented our models as well as the heatmap coordinate encoding and decoding operations with the use of a
 134 configuration file that allows for the definition and configuration of the various components and hyperparameters.

135 Similarly to DARK paper, we report the following standard pose estimation metrics:

136 **RMSE**: represents the square root of the squared differences between estimated and ground truth keypoint coordinates
 137 in pixels.

138 **PCKh**: considers an estimated keypoint location correct only if the euclidean distance error is lower than a percentage
 139 threshold α of the euclidean distance between the head body part segment, i.e. from neck to head joint positions. In our
 140 experiments, we use *PCKh-0.1* and *PCKh-0.5* with $\alpha = 10\%$ and $\alpha = 50\%$, respectively.

²<https://github.com/ai-in-motion/moai>

141 **3.6 Computational requirements**

142 The training of the models was executed on a workstation machine with the hardware specifications figured in Table 1.

Table 1: Hardware Specifications

| | |
|----------------|-------------------------------|
| <i>OS</i> | Windows Microsoft Pro (x64) |
| <i>Storage</i> | 3TB Toshiba HDD |
| <i>CPU</i> | Intel i9-7900X (4.30 GHz) |
| <i>GPU</i> | GeForce RTX 2080 Ti (12 GB) |
| <i>RAM</i> | 4 x 16 GB Kingston (2666 MHz) |

143 **4 Results**

144 Our experimental results validate both claims of DARK [1], as presented in Table 2. In particular, **w/o quant** models
 145 outperform the corresponding **w/ quant** ones, while *DARK* coordinate decoding method shows the best performance
 146 compared to the other competitive decoding methods.

147 **4.1 Results reproducing original paper**

148 The results presented in Table 2 support both claims of the paper, allowing the reader to easily study and relatively
 149 assess the coordinate decoding and encoding achievements of each method.

Table 2: The overall *PCKh-0.1*, *PCKh-0.5* and *RMSE* results of **SH_4S** and **HRNet_4S** on HUMAN4D trained **w/ quant** and **w/o quant** coordinate encoding methods. The final keypoint estimation is assessed with the use of various coordinate decoding methods.

| Model + Coordinate Decoding | PCKh-0.1 ↑ | PCKh-0.5 ↑ | RMSE ↓ |
|---|-------------------|-------------------|---------------|
| SH_4S (w/ quant) + <i>ArgMax</i> | 35.38 | 68.14 | 4.12 |
| SH_4S (w/o quant) + <i>ArgMax</i> | 39.80 | 68.78 | 3.75 |
| SH_4S (w/ quant) + <i>Standard</i> | 36.37 | 68.36 | 4.07 |
| SH_4S (w/o quant) + <i>Standard</i> | 41.18 | 69.05 | 3.69 |
| SH_4S (w/ quant) + <i>CoM</i> | 37.77 | 70.35 | 3.87 |
| SH_4S (w/o quant) + <i>CoM</i> | 44.24 | 71.45 | 3.33 |
| SH_4S (w/ quant) + <i>DARK</i> | 39.54 | 70.75 | 3.87 |
| SH_4S (w/o quant) + <i>DARK</i> | 45.17 | 71.65 | 3.47 |
| HRNet_4S (w/ quant) + <i>ArgMax</i> | 37.29 | 68.10 | 3.93 |
| HRNet_4S (w/o quant) + <i>ArgMax</i> | 39.19 | 68.60 | 3.85 |
| HRNet_4S (w/ quant) + <i>Standard</i> | 38.48 | 68.35 | 3.87 |
| HRNet_4S (w/o quant) + <i>Standard</i> | 40.50 | 68.90 | 3.79 |
| HRNet_4S (w/ quant) + <i>CoM</i> | 33.95 | 68.38 | 3.94 |
| HRNet_4S (w/o quant) + <i>CoM</i> | 36.77 | 70.21 | 3.79 |
| HRNet_4S (w/ quant) + <i>DARK</i> | 41.73 | 70.29 | 3.65 |
| HRNet_4S (w/o quant) + <i>DARK</i> | 44.83 | 71.52 | 3.55 |

150 As mentioned above, our main aim is to highlight the effectiveness of the distribution-aware heatmap representation,
 151 instead of assessing the performance and comparison between the models. Nevertheless, in our experiments, a general
 152 comment could be that **SH_4S** performs better than **HRNet_4S** almost in all combinations.

153 **4.1.1 Taylor-expansion coordinate decoding - Claim #1**

154 Assessing the various coordinate decoding methods we implemented and evaluated, as presented in Sec. 3.2.1, Claim #1
 155 is fully supported since the Taylor-expansion coordinate decoding does improve the coordinate estimation performance.
 156 Considering the predicted heatmap as a Gaussian heatmap and applying Taylor-expansion coordinate decoding shows

157 remarkable results against the rest of the methods. In particular, as shown in Table 2, DARK decoding performs better in
 158 all experiments for all metrics, with the only exception the RMSE error in “*SH_4S (w/o quant) + CoM*” experiment. In
 159 that case, DARK is more accurate than CoM, but CoM exhibits a smaller average error. This means that DARK suffers
 160 from larger outliers than CoM, a partly reasonable case given their nature. CoM uses the entire heatmap to reach a
 161 consensus about the coordinate, while DARK uses local heatmap statistics around the maximum activation, but is more
 162 susceptible to heatmap noise. When applying DARK, a low-pass Gaussian filter is first used, which is standard practise
 163 before the application of derivative extraction, in order to suppress noise and capture only the signal’s high frequency
 164 information. However, this might not always be the case in out-of-distribution data predictions. When the predictions
 165 follow the underlying assumptions of DARK, namely a 2D Gaussian distribution, its finer-grained localisation will
 166 be effective. In the case of distribution noise, the globalised nature of CoM will be more robust. Evidently, the SH
 167 experiment is one such case, but nonetheless, DARK’s accuracy indicates that it still offers favourable behaviour.

168 4.1.2 Non-quantized coordinate encoding - Claim #2

169 Observing again the results in Table 2 from another perspective, we can assess the validity of Claim #2, i.e. that the
 170 coordinate quantization before heatmap generation results in inaccurate and biased ground-truth signals that reduce the
 171 model performance. Indeed, for both architectures, *SH* and *HRNet*, and for every coordinate decoding method, *w/o*
 172 *quant* models show higher *PCKh* accuracy and lower *RMSE* errors, totally supporting the second claim of the authors.




Figure 2: Pose estimation qualitative results with the use of DARK heatmap decoding on HUMAN4D testing dataset (unseen subject). Cyan and red colours indicate the ground-truth and the predicted poses, respectively.

173 4.2 Results beyond original paper

174 An experiment we conducted, beyond the experiments of the present paper, is the comparison between DARK and
 175 spatial coordinate regression methods for keypoint estimation, such as Center of Mass (CoM) [2, 3].

176 4.2.1 DARK against Center of Mass

177 Observing Table 2, DARK obviously outperforms CoM almost in all experiments. Nevertheless, it is worth-noting
178 that focusing on the results of the best performing model, as shown in Table 3, CoM demonstrates comparably to
179 DARK. Especially in RMSE, which is a continuous metric, CoM presents lower errors than DARK, while, generally, its
180 performance is higher when decoding heatmaps predicted by SH_4S.

Table 3: *PCKh-0.1*, *PCKh-0.5* and *RMSE* results of best performing model (SH_4S (w/o quant)) with the use of DARK and CoM coordinate decoding methods.

| Model + Coordinate Decoding | PCKh-0.1 \uparrow | PCKh-0.5 \uparrow | RMSE \downarrow |
|-----------------------------|---------------------|---------------------|-------------------|
| SH_4S (w/o quant) + CoM | 44.24 | 71.45 | 3.33 |
| SH_4S (w/o quant) + DARK | 45.17 | 71.65 | 3.47 |

181 5 Discussion

182 Summarizing the present reproduction study, we conclude that the DARK paper introduces a novel, simple, lightweight
183 and highly effective method for heatmap coordinate encoding/decoding that boosts the keypoint estimation accuracy.

184 DARK outperforms standard greedy sub-pixel coordinate decoding methods. Furthermore, it shows higher accuracy
185 than spatial coordinate regression methods, such as CoM, which also shows lower erroneous predictions than standard
186 methods. On top of those, DARK can be considered a significant improvement for multi-person pose estimation. That
187 is due to the fact that spatial coordinate regression methods, though highly effective and fully differentiable, struggle in
188 the regression of more than one keypoint location when more than one heatmap in the same heatmap layer, i.e. specific
189 joint, are present in the same predicted layer, a common case in multi-person datasets.

190 With regards to identified gaps in the present heatmap representation study, we believe that spatial coordinate regression
191 methods, such as CoM or MoM (median of mass) [8], should have been taken into account, benchmarked and reported
192 in this paper. Several recent heatmap-based keypoint estimation data-driven methods have been developed with the use
193 of spatial coordinate regression, which beyond its effectiveness, allows the design of end-to-end, fully differentiable
194 architectures with direct supervision on the actual task of keypoint estimation, i.e. the coordinate distance error between
195 the predictions and ground-truth.

196 5.1 What was easy

197 Implementing most of the code was straightforward as authors provide their source code on GitHub. Apart from that,
198 past GitHub issues were another source of retrieving information, clarifying parts of the paper when needed.

199 5.2 What was difficult

200 All and all, this reproduction study was smooth without bottlenecks, delays or other barriers. A potential difficulty
201 that can be considered was the availability of human pose datasets, beyond the ones used in the paper, that provide
202 coordinates in higher resolutions, in order to further investigate the quantization error effects.

203 5.3 Communication with original authors

204 As implementing the code of the paper was straightforward and the ideas presented in the paper are clear and well-
205 defined, we did not communicate with the authors.

206 References

- 207 [1] Feng Zhang, Xiatian Zhu, Hanbin Dai, Mao Ye, and Ce Zhu. Distribution-aware coordinate representation for
208 human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
209 pages 7093–7102, 2020.
- 210 [2] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings*
211 *of the European Conference on Computer Vision (ECCV)*, pages 529–545, 2018.

- 212 [3] Aiden Nibali, Zhen He, Stuart Morgan, and Luke Prendergast. Numerical coordinate regression with convolutional
213 neural networks. *arXiv preprint arXiv:1801.07372*, 2018.
- 214 [4] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European
215 conference on computer vision*, pages 483–499. Springer, 2016.
- 216 [5] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose
217 estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages
218 5693–5703, 2019.
- 219 [6] Anargyros Chatzitofis, Leonidas Saroglou, Prodromos Boutis, Petros Drakoulis, Nikolaos Zioulis, Shishir Subra-
220 manyam, Bart Kevelham, Caecilia Charbonnier, Pablo Cesar, Dimitrios Zarpalas, et al. Human4d: A human-centric
221 multimodal dataset for motions and immersive media. *IEEE Access*, 8:176241–176262, 2020.
- 222 [7] *moai*: Accelerating modern data-driven workflows. <https://github.com/ai-in-motion/moai>, 2021.
- 223 [8] Christopher Tensmeyer and Tony Martinez. Robust keypoint detection. In *2019 International Conference on
224 Document Analysis and Recognition Workshops (ICDARW)*, volume 5, pages 1–7. IEEE, 2019.