arXiv:2411.10334v1 [cs.CV] 15 Nov 2024

Y-MAP-Net: Real-time depth, normals, segmentation, multi-label captioning and 2D human pose in RGB images

Ammar Qammaz
Institute of Computer Science, FORTH and
Computer Science Department, University of Crete
ammarkov@ics.forth.gr

Nikolaos Vasilikopoulos
Computer Science Department, University of Crete, and
Institute of Computer Science, FORTH
nvasilik@ics.forth.gr

Iason Oikonomidis
Institute of Computer Science, FORTH,
N. Plastira 100, Vassilika Vouton,
GR70013, Heraklion, Crete, Greece
oikonom@ics.forth.gr

Antonis A. Argyros
Institute of Computer Science, FORTH and
Computer Science Department, University of Crete
argyros@ics.forth.gr

Abstract

We present Y-MAP-Net, a Y-shaped neural network architecture designed for real-time multi-task learning on RGB images. Y-MAP-Net simultaneously predicts depth, surface normals, human pose, semantic segmentation, and generates multi-label captions—all from a single network evaluation. To achieve this, we adopt a multi-teacher, single-student training paradigm, where task-specific foundation models supervise the network's learning, enabling it to distill their capabilities into a lightweight architecture suitable for real-time applications. Y-MAP-Net exhibits strong generalization, simplicity, and computational efficiency, making it ideal for robotics and other practical scenarios. To support future research, we will release our code publicly.

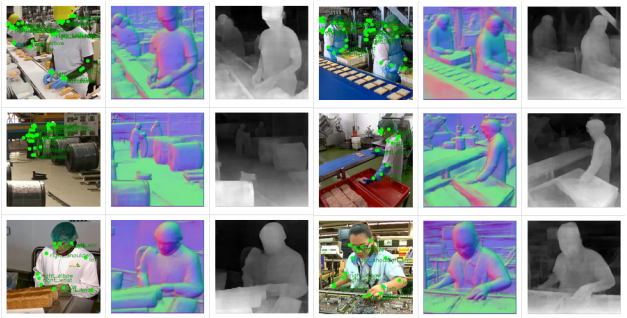

Figure 1. Given an RGB frame, Y-MAP-Net estimates human pose, depth, surface normals, segmentation and image captioning in real-time. The figure shows pose keypoints, depth and normal estimations on publicly available images from factory floors.

1. Introduction

Decades of research have yielded powerful methods that robustly tackle long standing problems for computer vision and pattern recognition. These so-called foundation models stand out for their exceptional generalization ability, achieved through immense scale and complexity. These models, with billions of parameters, are trained on vast datasets, enabling zero-shot problem-solving across diverse tasks and robust performance on in-the-wild data. However this sheer size comes at a cost: they require PetaFLOP-scale computational resources and data centers equipped with thousands of GPGPU accelerators for a single training session, hindering reproducibility. Finally, they are expensive to deploy and impractical for real-time applications.

Meanwhile, we observe a class of much sparser convolutional models, such as the successful family of You Only Look Once (YOLO) [48] models, which are profoundly influential in many different real-world applications. This motivates us to develop a framework that bridges the gap between real-time, multifaceted scene understanding and high-fidelity dense output. Once provided, it can be used as a building block to solve other general problems, in vision, robotics, human-computer interaction, surveillance, etc, by providing a unified, streamlined source of multi-modal data.

Given this motivation, in this work, we propose Y-MAP-Net, a **Y**-shaped **M**ulti **A**ttribute **P**rediction neural network

(NN) model that can be used as a closed-loop processor for RGB streams supplying human pose, depth, surface normals, image segmentation and a textual multiclass caption for each incoming frame (Fig. 6). To achieve this, we adopt a multi-teacher, single-student training paradigm [4, 74], where Y-MAP-Net learns from task-specific foundation models supervising its predictions.

This approach distills large models into a compact architecture, enabling efficient multi-task learning without sacrificing generalization. In particular, Y-MAP-Net accepts monocular RGB input and delivers 44 heatmap / image outputs and 8 caption tokens.

In summary, our main contributions are the following:

- Y-MAP-Net is the first convolutional NN method to achieve simultaneous depth, normal and human pose estimation while also providing scene segmentation and captioning from monocular RGB in a monolithic network.
- Due to its efficiency, the method we present performs in real-time and thus opens the way to many practical applications, especially in the field of robotics, complementing existing models for the various tasks we undertake.
- The proposed Y-MAP-Net topology offers a novel, simple and pragmatic framework for multi-task learning.

2. Related Work

The proposed method attempts to simultaneously tackle multiple mature sub-fields of the computer vision research landscape. To achieve such a generalist technique we studied the various domains, identified their commonalities, best practices, useful data primitives and then gradually integrated them under a common unified formulation.

Image Classification and Captioning: AlexNET [33] was the first method to successfully deploy a large scale convolutional neural network, train it with a large dataset and perform 1000-way softmax classification. Although more than a decade has passed since this seminal work [1], GPGPU acceleration and techniques like data augmentation, dropout, and ReLU activations have become mainstream and are used with small changes in our work. The VGG network family [53] demonstrated the importance of increasing depth, with depth ranging from 11 to 19 layers, also leveraging 3×3 convolutions. GoogLeNet [57] introduced an even deeper architecture featuring multiple data paths with different kernel sizes between layers, pioneering 1×1 convolution kernels. Finally, ResNETs [20] introduced residual connections that improved gradient flows throughout the network, allowing NNs to grow up to 1202 layers deep, dealing with the same classification task. Meanwhile, works like DenseCap [24] moved away from labels to dense text captioning while in the Natural Language Processing (NLP) field Word2Vec [40, 41], GloVe [44] and Fasttext [26] embeddings encoded text in a multi-dimensional representation with conceptual relations between vectors. The advent of transformers [60], a technique built from the ground up to capture textual relations, replaced other LSTM [22] or GRU [8] based attempts. ConvCAP [2] was the last major purely convolutional dense captioning work. Vision transformers (ViTs) [13] inspired development of techniques like CLIP [46] and DINO [6] that directly associated image and text embeddings, paving the way for state of the art VLLM foundation models like GPT4 [15] and LLAMA 3.2 [14] that successfully tackle the problem at its core.

Scene Segmentation: An alternative approach to scene captioning is annotating visible parts of the scene. YOLO [48] served as a seminal work for image region classification. Building on this, Mask R-CNN [21] and Faster R-CNN [49] advanced dense segmentation and per-pixel label classification, inspiring subsequent works like Mask YOLO [55]—fusing earlier approaches [21, 48]—as well as Detectron 2 [64], DPText [71], and Segment Anything [31].

Pose Estimation: AlexNET [33] directly inspired Human Pose Estimation (HPE) methods [58, 59] that adapted it for joint heatmap regression. VGGs [70] and ResNETs [36, 54], followed by stacked hourglass networks [42] were among the first new NN topologies for the HPE task. OpenPose was a landmark work that introduced part affinity fields (PAFs) [5]. During the same year, LCR-Net [50] and Coarse-to-fine 3D human pose [43] also made significant contributions in terms of 3D accuracy. DensePose [19] explored an alternative formulation of the problem by regressing UV-mapped coordinates to a 3D model, thus providing a much denser result than 3D joint positions. HRNET [56] modified U-NETs [51] by featuring a pyramid of different input resolutions and achieved new levels of accuracy. This inspired a host of more approaches [7, 61, 67, 75]. The advent of ViTs [13] shifted focus to transformer-based architectures [11], with the Sapiens [28] foundation model being the first to achieve 2D joint, depth and normal estimation, as well as part segmentation but focusing only on humans.

Depth and Normals Estimation: MiDaS [47] was a major push towards monocular depth estimation showcasing the power of relative depth. ZoeDepth [3], Metric3D [73] and Marigold [27] further refined accuracy, while Dust3R [63], UniDepth [45] and DepthAnything (DA) [68] scaled in both model size and training data leading to DAv2 [69] a mature framework we use as the teacher for our depth training.

Training Datasets: A wide variety of data sources like LAION-5B [52], Instagram-3.5B [37], JFT-300M [54], and V.Genome [32], take advantage of big-data effectiveness in deep learning, providing enormous training corpora, especially fitting for image captioning. Humans-300M [30], Panoptic [25], AMASS [38] and Motion-X [34] are the largest HPE datasets. Common Objects in Context (COCO) [35] remains a balanced middle-ground for captioning, scene segmentation and HPE. It has sufficient size and scope to facilitate reproducible research without being overly encumbering in terms of required training resources. Our approach stands on the "shoulders of giants" of prior research, drawing from a substantial body of computer vision work. At its core, our NN topology has similarities to U-NETs [51], however being pushed to new limits interms of multiple concurrent tasks. Although Y-shaped architectures featuring two output branches have been proposed to tackle specialized medical applications [39, 62],

our Y-MAP-Net formulation is novel by providing captions output using GloVe [44] embeddings. In contrast to other caption works like ConvCAP [2], it does not iteratively query the network for each captioning token. Additionally, our model simultaneously regresses 2D joint heatmaps and PAFs, inspired by OpenPose [5], and predicts depth and surface normals supervised by DepthAnything [68], with conceptual parallels to DensePose [19]. Moreover, supervised by Detectron 2 [64] and DP-Text [71] our work is equipped with per pixel object and scene segmentation capabilities.

A method with multi-modal output and conceptual similarities to Y-MAP-Net, though not directly comparable, is the Sapiens [30] foundation model. Unlike our method, Sapiens focuses solely on the human body, not producing results for the entire input image. While Sapiens uses a 500M training corpus, our approach relies on just 0.1M samples. Architecturally, Sapiens is transformer-based, whereas we adopt a convolutional approach. Instead of Sapiens' separate encoder models for each modality [29]—ranging from 0.3B to 2.0B parameters—our monolithic model is $4\times$ smaller than the aggregate size of the minimum Sapiens setup to tackle all tasks. In our monolithic model, weights are shared between tasks, in contrast to the sequential approach adopted in Sapiens. Finally, we include captioning on the same model, in contrast to Sapiens [30] which lacks this functionality. All these differences render our approach not directly comparable to Sapiens.

3. Y-MAP-Net Model Architecture

Y-MAP-Net has a single **input** or **encoder** branch, marked with green in Figure 2. Its outputs diverge in two dedicated paths. The first, marked with magenta, is related to **pictorial** or **decoder output**: depth, normals, heatmaps, PAFs, and segmentation masks. The second, marked with blue, is related to **token/caption output**.

These two paths are separated at the **bridge** section (orange color in Fig. 2). This structure effectively combines a U-NET-like [51] architecture for pixel-level related tasks with a densely connected token encoder ensemble for caption generation.

Encoder branch: This branch gradually distills lower dimensional feature embeddings until arriving at the bridge section. This is achieved through a series of convolutions with 3×3 kernels followed by average pooling (see Fig. 3).

Bridge segment: Unlike regular U-NETs [51], our bridge segment consists of three densely connected layers, instead of only one that is usually employed. This is designed to handle the multi-modal GloVe [44] embeddings of the token branch as well as the multi-faceted pictorial output.

Pictorial decoder branch: The pictorial decoder branch of the network produces 44 image channels containing pixel-level output. We observed that adding a 1×1 convolution layer with a very high number of filters prior to the output layer increases the fidelity of these images. We use a $1 \times 1 \times 1500$ layer, a decision affected by available H/W resources (we train our network using an NVIDIA RTX A6000 GPGPU with $\approx$ 49GB VRAM). We attribute the

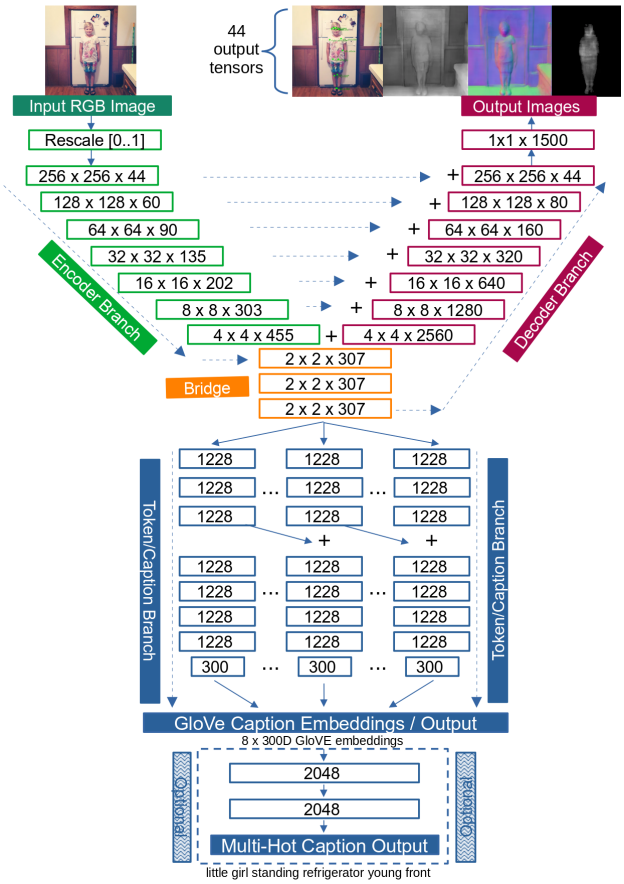

Figure 2. Architecture of the proposed Y-MAP-Net. The flow of data from input to outputs is highlighted with cyan arrows. Residual connections are indicated with + signs, while rectangles give a dimensionality overview for each layer. Green layers (top left) signify encoder blocks originating from 1 RGB image. The network bridge layers appear in orange color (middle). Blue layers (bottom) depict 8 captioning output token GloVe [44] vectors. These can be optionally converted to multi-hot labels with the addition of 2 dense layers for some applications. Magenta (top right) signifies decoder blocks that lead to 44 multi-modal outputs. Detailed encoder and decoder layer architecture is provided in Figure 3.

benefits of this "pixel-wise" layer to increased capacity for a wider range of activated features directly before the final layer of the network.

Both the pictorial/image output and the GloVe token embeddings use hyperbolic tangent ($tanh$) activations in the output layer. This choice forces output in the range of [-1..1], effectively controlling gradient flow. Remaining layers utilize Leaky ReLU [65] activations that also allow a small gradient when inactive to mitigate vanishing gradients. Towards the same end, GloVe vectors from captioning output are also normalized to [-1..1], while the first layer after RGB input also rescales RGB values to [0..1]. As seen in Fig. 3, the encoder blocks of the network feature residual connections once again as a measure to diminish vanishing gradients. Residual connections are also featured, di-

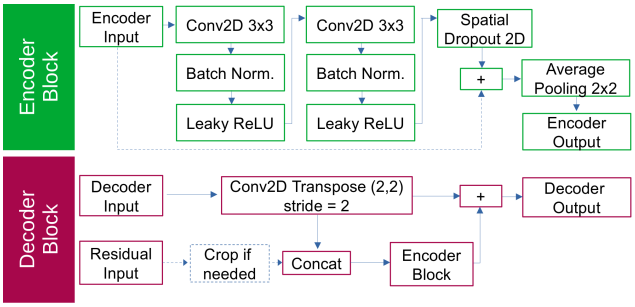

Figure 3. Encoder (green) and decoder (magenta) blocks of Y-MAP-Net (Fig. 2) are tasked with down-scaling input images to the bridge representation and then up-scaling to pictorial outputs. They consist of layers and skip connections shown in this figure.

rectly connecting encoder blocks to their respective same-dimensional decoder counterparts.

Token/caption output branch: The token/caption output branch departs from previous designs. In the past, typical captioning networks used LSTMs [22] or GRUs [8] featuring recurrent connections. ConvCAP [2] maintained a purely convolutional approach. However, recurrence was still maintained because the network expected as input the previously regressed tokens in order to iteratively produce the next ones until forming a full sentence. This would be a prohibitive design choice for Y-MAP-Net because it would immensely slow down the processing framerate. It is worth noting that although ConvCAP experimented with an early pairwise mechanism similar to transformer attention, this proved to have a very small positive impact upon the core ConvCAP architecture. Similarly, modern transformer architectures rely on positional encodings [60] that vastly increase the complexity of the network by calculating attention scores for all possible output pairs. Architectures like differential transformers [72] are now specifically built to tackle this problem by canceling out attention scores on "noisy"/unimportant [72] token pairs, while research is being conducted on the attention heads themselves [76] and the roles they are called to accomplish. In contrast to all this complexity, we follow a much more straight forward approach by eliminating the so called stop-word [16] tokens, using GloVe as our token vector space. This helps with ordering the various responses of the network while training it. Lastly, residual connections from each token output to the next are incorporated in an effort to form encoding chains between tokens that build on previous tokens.

It should be noted that the Y-MAP-Net captioning branch features Layer- instead of Batch normalization between layers. 20% dropout is applied after each token layer and gets conversely reduced as we go through from the first to the last token. Residual connections from one token to the next are scaled to 30% of their initial magnitude to tackle the network producing the same token multiple times. Finally, output dimensions can vary depending on the task and computational resources. Overall, the full network we present outputs 44 heatmaps and 8 tokens: Heatmaps #0 to #16 encode 2D joints, #17-#28 PAFs [5], heatmap #29 depth, #30-#32 normals and #33 to #43 segmentation data.

3.1. 2D Human Pose Estimation (HPE)

HPE can be formulated as a series of transformations of an input RGB image to heatmaps that are active only in the vicinity of a specific human body joint. Typical skeletons like the one used in COCO include 17 joints that contain the eyes, ears, shoulders, elbows, hands, hips, knees and ankles of a person. Y-MAP-Net can scale to more than COCO body joints. However, in an effort to constrain its already very broad scope, we opted to just focus on these 17 joints.

Being able to return 2D keypoints is very useful for detecting the presence of humans in the scene. However, connecting the dots of joints to form complete skeletons is equally important. The output of a 2D HPE module should not only be 2D joints but a list of humans with each joint associated to a specific skeleton. PAFs [5] were proposed to tackle this exact problem, indicating connected joint pairs. Therefore, we include regression of PAFs in outputs #17-#28 which are heatmaps that are active in the areas between neighboring joints and zero everywhere else.

3.2. Depth and Normals Estimation

COCO17 [35], much like virtually all non-synthetic, in-the-wild datasets, does not contain depth data for training. For this reason, we selected Depth Anything V2 (DAv2) [68, 69] as a teacher model to generate them. DAv2 offers state of the art accuracy, and with it we can produce a depthmap for each RGB training sample of COCO17. Although DAv2 produces FP32 depth (similar to Y-MAP-Net) its implementation's .png output is encoded as 8-bit values. We therefore modify it to encode depth as 16-bit. DAv2 features 2 families of models, extracting metric and normalized range depth results. It also features multiple model sizes. We use the biggest one publicly available (vitl 335M) and use normalized depth ranges since they produce rich details and are a superior choice in terms of visual fidelity [47].

Given a depth image D_I for each training sample, we compute the partial gradients $\frac{\partial D_I}{\partial x}$ and $\frac{\partial D_I}{\partial y}$ by applying Sobel filters along the x and y directions, respectively. After performing gradient calculations, we compute the norm of the vectors, also adding a small value ϵ to avoid division by zero. We then compute the normal images n_x, n_y and n_z.

At first glance, including normal data as an output of our NN seems counter-intuitive, since a network regressing depth already encodes normals in the depthmap. In practice, we observe that normal output has better fidelity compared to depth output. This is attributed to neighboring data that have a similar appearance in RGB maintaining similar n_x, n_y and n_z in their local neighborhoods, effectively reducing noise. We thus leverage this behavior, using the output normals to refine output depth as seen in Figure 4.

Iterative depth refinement: Assuming an output depth image D_O and normal maps N_x, N_y and N_z, we compute gra-

Figure 4. NN output normals (right) enforced on NN depth output (left) through our iterative algorithm can improve depth (mid).

dients in x and y directions by locally sampling depth:

$$G_x = \frac{\partial D_O}{\partial x}, \quad G_y = \frac{\partial D_O}{\partial y}.$$

We normalize gradients with their magnitude M to obtain the normalized gradient maps G_x^{nor}, G_y^{nor}, G_z^{nor} and we calculate target normals and the gradients differences of the normalized depth map to find inconsistencies:

$$\Delta_x = N_x - G_x^{\text{nor}}, \quad \Delta_y = N_y - G_y^{\text{nor}}, \quad \Delta_z = N_z - G_z^{\text{nor}}.$$

We finally update the improved output depth map D_O by minimizing the normal difference:

$$D_O = D_O + \alpha \left(\Delta_x \cdot G_x + \Delta_y \cdot G_y + \Delta_z \right).$$

This process is executed iteratively for 35 iterations with $\alpha = 0.01$. We mask-out low depth (far) areas to skip calculations and respect our real-time target. As seen in Figure 4, this process produces noticeable improvements, sharpening resulting depthmaps and suppressing noisy values.

3.3. Image Segmentation

COCO 17 contains 183 segmentation label categories containing vehicles, animals, objects and other elements commonly visible in digital photographs. To perform per pixel segmentation, we need to arrange data in a way where, for each pixel, there is an associated output label. Encoding all possible discrete values into a single heatmap proved infeasible, since the neural network predictably lacks the capacity to produce such finely detailed activations at the segmentation output layer. We thus need to decouple various labels as different heatmaps, similarly to how 2D joints are treated to make the task feasible for HPE.

However, using 183 separate 256×256 heatmaps, results in broadcasting $\approx$ 12 million elements per training sample, leading to excessively long training times. We attempt to encode all classes in one image yielding only 65K for transit to the GPGPU per training sample, and perform one hot encoding inside the loss function. Although this restores training times, with some expressivity loss, the capacity of the output layer doesn't allow this to work in practice. To overcome this problem, we group segmentation classes in 11 broad categories, namely Persons, Vehicles, Animals, Objects, Furniture, Appliances, Materials, Obstacles, Building, Nature and Text. Each label is assigned to its respective group yielding 11 segmentation images with the union of all visible classes. This balanced, divide & conquer approach helps the network successfully tackle a large number of labels while minimizing overheads.

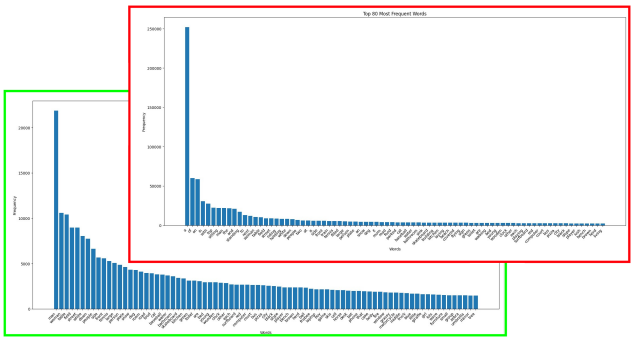

Figure 5. The frequency of tokens encountered while captioning an image follows a very heavy-tail distribution. **Red:** Top-80 frequency words, when not using stop-words [16]. Token 'a' appears several orders of magnitude more than e.g. the token 'cat'. This very heavy class imbalance negatively impacts training. **Green:** After removing tokens: *'(', ')', '.', 'a', 'an', 's', 'of', 'on', 'and', 'I', 'in', 'the', 'is', 'it', 'at', 'to', 'with', 'for' and 'from'*, we get the second distribution (green frame) which is more balanced.

3.4. Multi-Label Captioning

Captioning an image can be posed as selecting words that describe it and putting them in a sentence. However, as seen in Fig. 5 there's severe class imbalance among different tokens/labels. Works like [23] highlight the intricate differences between multi-label, multi-instance, and regular captioning, motivating us to follow the former. To improve data distribution and conserve NN capacity, we remove stop-words (Fig. 5). After limiting tokens to the 2048 most popular ones, initially, we naively attempted to perform multi-hot encoding for 8 output tokens, yielding 8×2048 captioning outputs for each sample. This approach however did not work since token order is very important and the task proves too difficult for a lightweight convolutional approach.

Performing the union over these tokens and emitting a multi-hot 2048 class vector with 8 active labels per sample showed promise. This training approach effectively removed token ordering, significantly simplifying the task. Despite performing class weighting based on the frequencies shown in Figure 5 however, the model tended to only produce results for a subset of tokens, and frequently made errors. This behavior is understandable: tokens like man, boy, woman, girl, person, and human have visually similar representations in images. Yet, despite their conceptual similarity, misclassifying a person as a man versus a cat incurs the same loss penalty, leading to problematic training.

To improve upon this, we employed Global Vectors (GloVe) [44] for word representation. GloVe offers different pre-trained databases. The vocabulary and source material for training can also vary ranging from Wikipedia, Gigaword, Common Crawl and Twitter/X for the pretrained

models. We select GloVe 6B which is uncased, is trained on Wikipedia 2014 and Gigaword 5, with 6B tokens, a 400K vocabulary, and offers 50, 100, 200 or 300 dimensional embeddings for each word. Employing GloVe we lose the ability to perform a union (or sum) of words as a combined output. However, attempting to predict 8×300 captioning outputs, this time using a MSE loss, showed signs of promise even from very early experiments. Observed output artifacts included having multiple outputs regress similar embedding values. During runtime, to transform the GloVe 300D output back to words we follow three steps. First we threshold the norm of the 300D GloVe vector against zero, filtering it out if it is very close to zero. We then select the word token that has the highest cosine similarity with our output. As a final step we can perform a second threshold operation filtering out tokens with low cosine similarities.

In order to improve token sequences we introduced residual connections from each token to the next. We observed that this further increased token repetition. Using tanh output activations, normalizing GloVe embeddings to [-1..1], employing dropout, and layer instead of batch normalization, as well as scaling down the magnitude of the residual connections to 30% mitigated this. As reported by methods like ConvCAP [2], GloVe embeddings can be used as intermediate output supervision. Adding a final 2-layer densely connected tail (bottom of Figure 2) with sigmoid or softmax activations yields easy-to-use outputs for applications. This eliminates the need for extra computations, since thresholding the returned pseudo-probabilities is sufficient.

3.5. Data Augmentation

COCO17 contains 118K training samples. To enhance Y-MAP-Net's generalization with this relatively limited amount of data and allow for a multi-epoch training regime without over-fitting, we employ aggressive data augmentation. The average sample dimension of COCO17 is close to VGA resolution. Our input and output images, however, have a 256×256 size. We thus resize the images while respecting aspect ratios, and with black padding. For each image we also apply a randomization scheme with the following properties. Each augmentation has an independent chance to be activated thus allowing multiple augmentations on an image. A sample has 45% chance to be panned & zoomed. It has 50% chance to have its brightness/contrast altered, out of which, for half of the samples we perform uniform brightness/contrast changes across all channels, while the rest have R,G,B channels individually altered with up to a $\pm 20\%$ change. There is a 15% chance of Gaussian noise perturbation, 50% chance of adding up to 10 burned pixels (up to 1/6553 corruption in 256×256 images). Finally, if a sample does not contain any persons there is a 1% chance of completely replacing the image with random noise. This ensures not hurting human pose due to our limited training data while at the same time conditions the network to filter out erratic input. An important detail of the pan & zoom augmentation is that it has a strict limit of up to 110%. If zooming is left unchecked, captioned areas of the image might be cropped out. This can cause concept drift, with tokens like surfboards associated with e.g. the sea. Captions describe the whole image and erroneous zooming may e.g. crop-out visible surfing boards which occupy less space than the sea. Finally, if persons are present, our randomization ensures at least 30% of existing joints of depicted persons visible inside the cropped image.

3.6. Model Training

We train our model using Keras 3.1.1 with Tensorflow 2.16.1 as a backend. We develop a data loader written from scratch in C and thoroughly profile and optimize it. We use a batch size of 30, corresponding to 30 threads concurrently executing on an AMD Ryzen Threadripper PRO 3955WX 16-Cores/32-Thread CPU. Back propagation is executed on an NVIDIA RTX A6000 GPU, where training occupies 47.2GB of the available 49.1GB VRAM. We also utilize a RAMFS to reduce I/O overhead. Our model weights are in FP32 precision; however, to enhance training throughput, we use INT8 representations for images during their transfer to the GPU. All training sessions start from weights following the Glorot/Xavier [18] initialization. We use the ADAM [12] optimizer with a learning rate starting at 2.4×10^{-4} and decaying linearly to 1.0×10^{-4} after 100 epochs. In total, we allow 250 training epochs, with checkpoints [9] and early stopping [10] with 65 epoch patience.

Heatmap target outputs also have a decay schedule with joints having a Gaussian size of 23×23 during the start of training and decaying down to a target limit of 6×6. In a similar manner, PAFs start from a line width of 6 and decay down to 2. This reduction happens linearly with the size being reduced by one step every 20 epochs. This helps the NN both identify joints early on during training since their magnitude is comparable to other output modalities, but also ensures better localized joint responses for maturely trained models. The loss function $\mathcal{L}$ we use is the following:

$$\mathcal{L} = w \cdot \sum_i g_i \cdot \frac{1}{N} \sum_{j=1}^{N} \left(y_{\text{true},\, i,j} - y_{\text{pred},\, i,j} \right)^2 .$$

In the above equation, i runs over 7 different terms. Each MSE term represents the mean squared error between the true and predicted values for the respective channels, weighted by its corresponding gain g_i. We experimentally determined g_i values that promote harder aspects of the problem by increasing their relative contribution to the total loss. These 7 values are: $g_G = 10.0$ - for GloVe tokens (chan. 0-7), $g_J = 2.4$ - for joint heatmaps (chan. 0-16), $g_{\text{PAF}} = 0.8$ - for PAFs (chan. 17-28), $g_D = 1.0$ (for Depth - chan. 29), $g_N = 1.0$ (for Normals - chan. 30-32), $g_T = 1.0$ (for Text - chan. 33) and $g_S = 3.0$ (for Segmentation - chan. ≥ 34). We use cosine similarity to monitor GloVe tokens. For image output we define **Heatmap Distance Metric (HDM)**. Essentially, HDM measures the fraction of all predicted pixels within relative threshold $\mathcal{T}$ from their target value. In experiments (Tables 1, 2), $\mathcal{T}$ = 0.1 unless stated otherwise. Assuming that y_{true} and y_{pred} are true and

Exp. No.	Layers	M.Size	In.Dim.	Out.Dim.	Pix.wise	Res.con.	Pos.Dpth	Normals	PAFs	Captions	Class Grp.	HDM train	HDM val.
1	4	46M	110	48	-	-	✓	-	-	-	-	0.69	0.56
2	4	31M	220	96	-	-	✓	-	-	-	-	0.70	0.71
3	4	39M	200	100	-	-	✓	-	-	-	-	0.70	0.69
4	5	48M	256	256	-	-	✓	-	-	-	-	0.70	0.70
5	5	54M	300	288	-	-	✓	-	-	-	-	0.82	**0.84**
6	6	70M	256	256	-	-	✓	-	-	-	-	0.66	0.67
7	7	231M	128	128	-	-	✓	-	-	-	-	0.42	0.52
8	7	136M	420	384	-	-	✓	✓	-	-	-	0.75	0.76
9	7	278M	420	384	-	-	✓	✓	-	-	19	0.74	0.77
10	8	401M	420	256	-	-	✓	✓	-	-	19	0.74	0.77
11	7	315M	256	256	-	-	✓	✓	✓	-	3	0.56	0.61
12	7	316M	256	256	1.6K	-	✓	✓	✓	-	3	0.62	0.67
13	7	172M	256	256	1.6K	✓	✓	✓	✓	-	3	0.66	0.72
14	7	213M	256	256	1.7K	✓	✓	✓	✓	-	2	0.53	0.59
15	7	298M	256	256	1.5K	✓	✓	✓	✓	✓	11	0.64	**0.72**

Table 1. Ablation study for Y-MAP-Net. We record enc/dec branch layer depth, model size (millions of weights), input and output dimensionality (images are rectangular so 256 is equivalent to 256x256), NN features such as pixelwise layer before output (Figure 2), residual connections in enc/dec blocks (Figure 3), output modalities enabled in an experiment (PAFs, Normals, Class segmentation) and the training/validation HDM score for each configuration. **Up:** Limiting output to joint+depth modalities we experimentally find that shallow networks perform better than more complex ones that over-fit. **Down:** As more modalities are introduced the problem becomes harder and tensor sizes need to be reduced to maintain real-time performance. The final row contains the Y-MAP-Net 1-8-44 configuration. Its weights are increased compared to previous experiments due to the added captioning branch. Each experiment requires $\approx$ 1 week to run.

predicted heatmaps in the [0..1] range, HDM is defined as:

$$Y = \begin{cases} 1 & \text{if } |y_{\text{true}} - y_{\text{pred}}| \leq \mathcal{T}, \\ 0 & \text{otherwise.} \end{cases}$$

$$\text{HDM} = \sum \mathbf{1}(Y = 1)/(\#\text{pixels}).$$

4. Experiments

Our method addresses a broad range of computer vision sub-topics, making it the first to jointly tackle them. It uses DAv2 [68], Detectron 2 [64], DP-Text [71] and VisionGPT2 [17] as teacher models in the multi-teacher, single-student setting we adopt. It is trained on COCO17 with a training set of 118K samples, which is orders of magnitude smaller than the millions of samples training corpora of our teacher models. Furthermore, in order to provide diverse functionality in real-time, our model is less than half the size compared to the sum of weights of its teachers (Table 3). Last but not least, our method uses the same weights to simultaneously deal with all different output modalities, while each of its teachers only deals with a single problem. We thus expect our method to produce "distilled" output of reduced yet comparable accuracy, compared to its teachers. As discussed in detail in the paragraph concluding Section 2 the conceptually closest method, Sapiens, is not directly comparable to ours. In the absence of other methods combining the functionality we propose, we base our experimental comparisons primarily against our teacher models.

Modality	Teacher	HDM 0.1	HDM 0.3	HDM 0.5	HDM 0.8	MSE
Joints	COCO GT	0.95	0.95	0.95	0.95	2.22
PAFs [5]	COCO GT	0.96	0.96	0.96	0.96	1.45
Depth	DAv2 [68, 69]	0.27	0.37	0.38	0.42	59.74
Normals	DAv2 [68, 69]	0.53	0.64	0.68	0.75	58.79
Text Segm.	DPText [71]	0.94	0.94	0.94	0.94	1.29
Class Segm.	Detectron 2 [64]	0.91	0.91	0.91	0.92	4.74

Table 2. Quantitative results for specific NN output modalities using MSE and HDM for various $\mathcal{T}$ thresholds (0.1 - 0.8).

4.1. Ablation Study

We conduct a series of ablation experiments with different I/O sizes, different relative depths of the network, gradually adding output modalities and broadening the size and scope of our NN until successfully accommodating pose, depth, normal, segmentation and captioning and reaching our target. These experiments are summarized in Table 1. We control network size by altering the number of filters in the enc/dec branches (Fig. 2). For example, as seen in rows #6 and #7 we can have a bigger NN that has enc/dec depth of 6 layers and a 3$\times$ larger NN handling half of the dimensionality of input and output in each dimension. As seen in experiment #4, a 5-layer enc/dec architecture that tackles just pose and depth has excellent performance. Achieving a similar HDM metric with #4 for all involved outputs, we need a substantially larger architecture with more capacity, like exp. #13. Residual connections and the 1$\times$1 pixelwise layer, add representational capacity allowing stronger networks that tackle more modalities. Pixelwise layer width however, is limited by GPU VRAM. Adding a captioning branch, extra weights do not increase the image output branch. Caption features arriving at the bridge layer (Figure 2) however seem beneficial for the decoder branch.

4.2. Quantitative Experiments

We evaluated our model on the COCO17 validation set, using ground truth data generated by our teacher models. Our quantitative metrics, summarized in Table 2, reveal varying performance across different tasks. Depth estimation yielded the lowest accuracy, which aligns with its inherent complexity. In contrast to Joints, PAFs and class segmentations that have coarse "on/off" values and large inactive areas, depth maps return a response for all pixels. Furthermore even minor inaccuracies in depth predictions result in significantly larger error contributions, particularly when depth is assigned to an incorrect distance. Estimating object depth in in-the-wild images presents additional challenges due to an absence of reference scale, multiple viewpoints, or motion cues, as discussed in [66]. Our depth teacher, DAv2 [68, 69] is trained on over 60M images, two orders of magnitude more than our training corpus, it has a

Method	M.Size	Joints	PAFs	Depth	Normals	ClassSeg.	Text Seg.	Captions
OpenPose [5]	28M	✓	✓	-	-	-	-	-
DAv2 [68, 69]	335M	-	-	✓	✓	-	-	-
Detectron2 [64]	63M	-	-	-	-	✓	-	-
DPText [71]	44M	-	-	-	-	-	✓	-
VisionGPT2 [17]	239M	-	-	-	-	-	-	✓
Y-MAP-Net (Ours)	297M	✓	✓	✓	✓	✓	✓	✓

Table 3. We tackle a wide variety of tasks via a monolithic network using 41.88% of the weights typically needed to perform the task. We use [17, 64, 69, 71] as teachers for Y-MAP-Net training.

larger size, and is dedicated to this single task as seen in Table 3. Normals perform better, with improved fidelity in areas with edges. We take advantage of this with our iterative algorithm (Figure 4). Class segmentation performs better and Joint/PAF estimation achieves the highest accuracy. We partly attribute this to the quality of human-annotated training data, since unsupervised model-generated labels may introduce inaccuracies. In captioning tasks, we use the cosine similarity metric achieving 0.59, 0.62, 0.64, 0.64, 0.53, 0.27, 0.08 and 0.01 scores for caption tokens 1 to 8. Overall we observe labels delivering a good generic image description. Our .keras encoded model occupies 1.1 GB on disk. Although beyond the scope of this study, FP16 quantization could be beneficial due to our straightforward convolutional architecture. Benchmarking an unoptimized Python runtime on Intel i7-4790 / NVIDIA GTX 1070, it needs 7.5 GB of VRAM and reaches a 5.6 Hz refresh rate. On an AMD Ryzen 7 3800X / NVIDIA RTX 4080 SUPER, it utilizes 14 GB VRAM, achieving a real-time rate of 15.9 Hz.

Figure 6. Qualitative results from COCO17 validation data showcasing various raw outputs of our model. Joints and PAFs contain the union of joint and PAF heatmaps, while column Segment. contains the union of columns 7-12 to showcase their synthesis.

4.3. Qualitative Experiments

We test our method on a variety of input sources. We observe good generalization on the COCO17 validation set, showcased in Figure 6. Since we aim at real-time applications in production environments, we experiment with Youtube videos featuring workers in production lines (see Fig. 1). Further samples of qualitative results are shown in Fig. 6. We observe normal and depth estimation consistent with input, pose estimation results that follow the bodies even when wearing protective equipment, multi-token captioning output that describes the major visible components of the RGB image and class segmentation that, when observing a few objects / subjects, is very sharp. More qualitative results are provided in the supplementary material.

5. Conclusions

We presented the first work to jointly address depth, normal, pose, class segmentation and multi-label captioning in real-time. The proposed work offers an elegant and compact architecture that could have an important impact on the field due to its versatility, with the potential to become a staple along sparse architectures like YOLO and offering multi-modal vision capabilities in an efficient and scalable way. Future work includes using a mixture of experts supervision from multiple teacher models per modality and increased training corpus. To support research on the topic, our code will be made publicly available on github.

6. Acknowledgments

This work was co-funded by the European Union (EU - HE Magician – Grant Agreement 101120731) and the Hellenic Foundation for Research and Innovation (HFRI) under the "1st Call for HFRI Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment", project I.C.Humans, no 91. The authors also gratefully acknowledge the support of the framework of the National Recovery and Resilience Plan Greece 2.0, funded by the European Union – NextGenerationEU (project Greece4.0).

References

[1] Md Zahangir Alom, Tarek M Taha, Christopher Yakopcic, Stefan Westberg, Paheding Sidike, Mst Shamima Nasrin, Brian C Van Esesn, Abdul A S Awwal, and Vijayan K Asari. The history began from alexnet: A comprehensive survey on deep learning approaches. *arXiv preprint arXiv:1803.01164*, 2018. 2

[2] Jyoti Aneja, Aditya Deshpande, and Alexander G. Schwing. Convolutional image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 3, 4, 6

[3] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 2

[4] XB Bruce, Yan Liu, and Keith CC Chan. Multimodal fusion via teacher-student network for indoor action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3199–3207, 2021. 2

[5] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. 2, 3, 4, 7, 8

[6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 2

[7] Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S Huang, and Lei Zhang. Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5386–5395, 2020. 2

[8] Kyunghyun Cho. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 2, 4

[9] François Chollet et al. Keras documentation - model checkpoint. https://keras.io/api/callbacks/model_checkpoint/, 2024. 6

[10] François Chollet et al. Keras documentation - early stopping. https://keras.io/api/callbacks/early_stopping/, 2024. 6

[11] Yuanzheng Ci, Yizhou Wang, Meilin Chen, Shixiang Tang, Lei Bai, Feng Zhu, Rui Zhao, Fengwei Yu, Donglian Qi, and Wanli Ouyang. Unihcp: A unified model for human-centric perceptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17840–17852, 2023. 2

[12] P Kingma Diederik. Adam: A method for stochastic optimization. *(No Title)*, 2014. 6

[13] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2

[14] Abhimanyu Dubey et. al. The llama 3 herd of models, 2024. 2

[15] OpenAI et. al. Gpt-4 technical report, 2024. 2

[16] Christopher Fox. A stop list for general text. *SIGIR Forum*, 24(1–2):19–21, 1989. 4, 5

[17] Shreyas Daniel Gaddam. Visiongpt2. https://github.com/shreydan/VisionGPT2, 2023. 7, 8

[18] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010. 6

[19] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7297–7306, 2018. 2, 3

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2

[21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2

[22] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 2, 4

[23] Wanting Ji and Ruili Wang. A multi-instance multi-label dual learning approach for video captioning. *ACM Trans. Multimedia Comput. Commun. Appl.*, 17(2s), 2021. 5

[24] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2

[25] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *Proceedings of the IEEE international conference on computer vision*, pages 3334–3342, 2015. 2

[26] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*, 2016. 2

[27] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2

[28] Rawal Khirodkar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens: Foundation for human vision models. *arXiv preprint arXiv:2408.12569*, 2024. 2

[29] Rawal Khirodkar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens discrete models repository. https://huggingface.co/facebook/sapiens, 2024. 3

[30] Rawal Khirodkar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens: Foundation for human vision models. In *European Conference on Computer Vision*, pages 206–228. Springer, 2025. 2, 3

[31] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2

[32] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome:

Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017. 2

[33] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 2

[34] Jing Lin, Ailing Zeng, Shunlin Lu, Yuanhao Cai, Ruimao Zhang, Haoqian Wang, and Lei Zhang. Motion-x: A large-scale 3d expressive whole-body human motion dataset. *Advances in Neural Information Processing Systems*, 36, 2024. 2

[35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 2, 4

[36] Diogo Carbonera Luvizon, Hedi Tabia, and David Picard. Learning features combination for human action recognition from skeleton sequences. *Pattern Recognition Letters*, 99: 13–20, 2017. 2

[37] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pages 181–196, 2018. 2

[38] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 2

[39] Sachin Mehta, Ezgi Mercan, Jamen Bartlett, Donald Weaver, Joann G Elmore, and Linda Shapiro. Y-net: joint segmentation and classification for diagnosis of breast biopsy images. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11*, pages 893–901. Springer, 2018. 2

[40] Tomas Mikolov. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 3781, 2013. 2

[41] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26, 2013. 2

[42] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation, 2016. 2

[43] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G. Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2

[44] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014. 2, 3, 5

[45] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10106–10116, 2024. 2

[46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2

[47] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020. 2, 4

[48] J Redmon. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 1, 2

[49] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016. 2

[50] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3433–3441, 2017. 2

[51] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 2, 3

[52] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 2

[53] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2

[54] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852, 2017. 2

[55] Jianing Sun. Mask-yolo. https://github.com/jianing-sun/Mask-YOLO, 2018. 2

[56] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 2

[57] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 2

[58] Jonathan J Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. *Advances in neural information processing systems*, 27, 2014. 2

[59] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *Proceedings of*

the IEEE conference on computer vision and pattern recognition, pages 1653–1660, 2014. 2

[60] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 2, 4

[61] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020. 2

[62] Kaiqiang Wang, Jiazhen Dou, Qian Kemao, Jianglei Di, and Jianlin Zhao. Y-net: a one-to-two deep learning framework for digital holographic reconstruction. *Optics letters*, 44(19): 4765–4768, 2019. 2

[63] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 2

[64] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. https://github.com/facebookresearch/detectron2, 2019. 2, 3, 7, 8

[65] Bing Xu. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015. 3

[66] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 7

[67] Lumin Xu, Sheng Jin, Wentao Liu, Chen Qian, Wanli Ouyang, Ping Luo, and Xiaogang Wang. Zoomnas: searching for whole-body human pose estimation in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):5296–5313, 2022. 2

[68] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024. 2, 3, 4, 7, 8

[69] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv:2406.09414*, 2024. 2, 4, 7, 8

[70] Wei Yang, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. End-to-end learning of deformable mixture of parts and deep convolutional neural networks for human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3073–3082, 2016. 2

[71] Maoyuan Ye, Jing Zhang, Shanshan Zhao, Juhua Liu, Bo Du, and Dacheng Tao. Dptext-detr: Towards better scene text detection with dynamic points in transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3241–3249, 2023. 2, 3, 7, 8

[72] Tianzhu Ye, Li Dong, Yuqing Xia, Yutao Sun, Yi Zhu, Gao Huang, and Furu Wei. Differential transformer, 2024. 4

[73] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023. 2

[74] Shan You, Chang Xu, Chao Xu, and Dacheng Tao. Learning from multiple teacher networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1285–1294, 2017. 2

[75] Changqian Yu, Bin Xiao, Changxin Gao, Lu Yuan, Lei Zhang, Nong Sang, and Jingdong Wang. Lite-hrnet: A lightweight high-resolution network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10440–10450, 2021. 2

[76] Zifan Zheng, Yezhaohui Wang, Yuxin Huang, Shichao Song, Bo Tang, Feiyu Xiong, and Zhiyu Li. Attention heads of large language models: A survey. *arXiv preprint arXiv:2409.03752*, 2024. 4