

Estimating Object Physical Properties from RGB-D Vision and Depth Robot Sensors Using Deep Learning

Ricardo Pedreiras Cardoso¹

Abstract—Inertial mass plays a crucial role in robotic applications such as object grasping, manipulation, and simulation, providing a strong prior for planning and control. Accurately estimating an object’s mass before interaction can significantly enhance the performance of various robotic tasks. However, mass estimation using only vision sensors is a relatively under-explored area. This paper proposes a novel approach combining sparse point-cloud data from depth images with RGB images to estimate the mass of objects. We evaluate a range of point-cloud processing architectures, alongside RGB-only methods. To overcome the limited availability of training data, we create a synthetic dataset using ShapeNetSem 3D models, simulating RGBD images via a Kinect camera. This synthetic data is used to train an image generation model for estimating dense depth maps, which we then use to augment an existing dataset of images paired with mass values. Our approach significantly outperforms existing benchmarks across all evaluated metrics.

I. INTRODUCTION

Planning how to grasp objects benefits from accurately estimating the mass before physical interaction with the object. While mass can be directly measured using a scale, this approach is impractical for robotic manipulation in dynamic or unstructured environments, where physical interaction before grasping is often not feasible. Instead, estimating the mass of an object using visual information from a camera offers significant advantages. A vision-based approach allows robots to predict mass without prior contact, enabling faster and more adaptive decision-making when grasping or lifting objects. Additionally, this method can be more flexible, as it does not require specialized hardware like force sensors or scales.

One major challenge in mass estimation from RGB data alone is the inherent ambiguity in determining object size and mass without sufficient visual context [18]. Existing datasets attempt to mitigate this by including bounding boxes, and some models predict low-resolution thickness masks to improve estimation accuracy [1]. However, these approaches are of little use in practice since they require robots first to identify the bounding box of the object they wish to interact with, a time-consuming process often requiring multiple views around the object. Depth sensors can avoid this limitation since they can disambiguate the object’s scale from the metric depth. However, previous works rarely made direct use of them, mainly due to the scarcity of datasets that integrate depth information, RGB images and mass. Consequently, no current models fully

exploit the combination of depth and image data for mass estimation across a wide range of objects.

To address these limitations, we propose augmenting an existing RGB-mass pairs dataset with depth. We achieve this by first fine-tuning a GLPDepth model [19] on a synthetic dataset made of RGB and depth pairs, and then applying the depth estimation to the whole dataset. We developed a multimodal architecture that takes as input an RGB image and a sparse point-cloud and outputs a mass prediction in Kg. We propose several variations of the architecture, and we train them on the augmented dataset. We also explore if training the model for point-cloud reconstruction simultaneously improves its mass estimation capabilities. Our results show that our method significantly outperforms previous approaches across all metrics evaluated and that performing point-cloud reconstruction simultaneously does not change the model’s predictive power in a significant way.

II. RELATED WORK

Mass estimation is often reformulated as two subproblems: (1) estimating volume and (2) estimating density. While volume can be derived through depth cameras or inferred from visual cues and background context, density estimation poses a much greater challenge [1]. Most existing research focuses on specific, well-constrained problems where mass estimation is simplified by making assumptions about the material properties, effectively reducing the problem to volume estimation. Even so, certain types of objects, such as containers, cans, or boxes, still introduce significant ambiguity. In such cases, mass estimation becomes inherently difficult, as internal contents may be occluded.

Some methods attempt to estimate mass through geometric features like object contours, size, or shape, often by correlating these features with known object properties. For example, in [2], the mass of fish was shown to be proportional to their outline, captured via a stereo camera setup. The authors used this knowledge to create a device that accurately captured the mass of swimming fish. This insight led to successful applications in livestock mass estimation, including pigs [3], chickens [4], and other animals [5], as well as fruits [6][7] and industrial products [8].

Some have integrated depth data to improve mass estimation, as demonstrated in [12] and its follow-up [13], where a Convolutional Neural Network (CNN) was used to classify and estimate the mass of scrap metals. These models employed both RGB and depth data together and achieved good results. In [12], the authors also explored the impact of using a single network for RGB and depth, treating depth as

*This work was not supported by any organization

¹Ricardo Pedreiras Cardoso is with Instituto Superior Técnico, University of Lisbon, Av. Rovisco Pais 1049-001 Lisboa, Portugal ricardoreisc@tecnico.ulisboa.pt

a fourth channel versus having two CNNs, where the network processes RGB images and depth data separately before concatenating them into a single image. Their experiments revealed that the latter approach consistently yielded better results across multiple experiments. However, similarly to the previous approaches, these models were designed for constrained environments with limited object types, which greatly restrained their generalizability to broader scenarios.

A more general approach is to use data-driven methods to infer the volume and mass of the object directly. The most notable example is the *image2mass* model [9], which addresses general mass estimation across a diverse set of objects. The authors built a large dataset of approximately 150,000 items scraped from Amazon.com, each annotated with an image, mass, and bounding box. This dataset covers a broad range of categories, including electronics, home goods, and packaged foods, making it more applicable to real-world scenarios compared to earlier, more domain-specific datasets.

The *image2mass* architecture uses two Xception CNNs [24], one to predict a coarse 19x19 thickness mask of the object alongside 14 handcrafted features (Geometry Module) and another to estimate the object’s density (Density Tower). A third module, called Volume Tower, is used to estimate the volume of the object. The Geometry Module is not trained together with the rest of the model, but rather is used as a sort of preprocessing step. The outputs of this module are then propagated to the two towers. The model makes the assumption that the object in the image has uniform density, which, while not true for most cases, does hold up in empirical experiments. Unlike the previous approaches, this model does not explicitly model volume and density. Instead, the model fuses the outputs of the Density and Volume Towers by multiplying them together. The volume tower focuses on the object’s geometric properties, using features extracted from the bounding box and other visual cues, while the density tower aims to infer material properties, effectively learning a distribution of possible densities from the training data. Subsequent research has expanded on this framework, such as in [10], which introduced material embeddings into the *image2mass* architecture, and [11], which introduced voxel grid estimation for better predictions of the thickness mask. Despite these developments, neither did improve on the original *image2mass* results.

III. DEPTH PREDICTION DATASET

Since no large-scale dataset exists that simultaneously provides RGB images, depth information, and corresponding mass values, we opted to augment the *image2mass* dataset [9] to include depth, given the substantial overlap between the tasks of mass estimation from RGB and RGB and Depth (RGBD).

To that effect, we fine-tuned a pre-trained GLPDepth model [14] on synthetic data mimicking the *image2mass* dataset. We built this using the ShapeNetSem [15] 3D models by first filtering to exclude all those that did not contain valid metric dimensions and weight. This resulted in a total of 8,948 distinct models. Then, we simulated a Kinect camera

and captured 14 RGBD images per object, 8 from equidistant rotations at a bird’s eye view and 6 for the top, bottom, front, back, left and right.

To ensure all objects appear roughly the same size in the image, we approximate each to a sphere with a diameter equal to its bounding box’s diagonal. For every sphere to take up the same amount of space on the screen, the camera’s distance must scale with the sphere’s radius. In practice, we position the camera at a distance of 2.1 times this diagonal for consistent object sizing.

By dividing the depth by the diagonal of the object’s bounding box, we normalize the depth, making it scale-agnostic. This allows the model to learn a scale-invariant depth map, which can later be converted back to metric units by multiplying by the bounding box diagonal. This approach is beneficial because the *image2mass* dataset lacks contextual size cues, and it leverages the available bounding box data. After training the GLPDepth model, we apply it to the *image2mass* dataset, creating an augmented dataset with dense depth information for training our model.

IV. MODEL ARCHITECTURE

Here, we present our model architecture for mass estimation from RGB and Depth, as shown in figure 1. It follows an encoder-decoder architecture. Similar to previous works, we make the assumption that the density of the objects is uniform. In practice, the model learns to estimate the average density of the object. We have two distinct decoders, one for density and another for volume were both take the latent vector generated by the encoders as input. The final mass is predicted by multiplying the outputs of the decoders together.

Several methodologies exist for encoding RGBD data, such as treating them as a single 4-channel image, two distinct images, an image and a point-cloud or by voxelizing the depth data. We argue that treating depth data as a point-cloud rather than an image provides significant advantages. (1) This method enables the model to leverage three-dimensional spatial information more effectively, (2) offers greater adaptability to different sensor intrinsic parameters, and (3) is adaptable to various depth sensors beyond traditional RGB-D cameras, such as LIDARs and SONARs, which produce point-clouds instead of single-image representations.

Our experiments primarily focus on evaluating how depth processing can enhance the mass estimation task. Hence, only the point-cloud encoder is swapped during the experiments. Additionally, we explore whether incorporating point-cloud reconstruction as a secondary task can further enhance the model’s performance in mass estimation. For our RGB encoder, we use a DenseNet-121 [20], mainly for its small size and good performance on various computer vision benchmarks. For the point-cloud encoders, we test three different encoder variations. The first processes point independently [16], the second considers local context and fuses multiple scales [17], and the third employs a transformer-based architecture with local context [21]. Finally, we use a FoldingNet [22] decoder to reconstruct a coarse point-cloud when performing the reconstruction task.

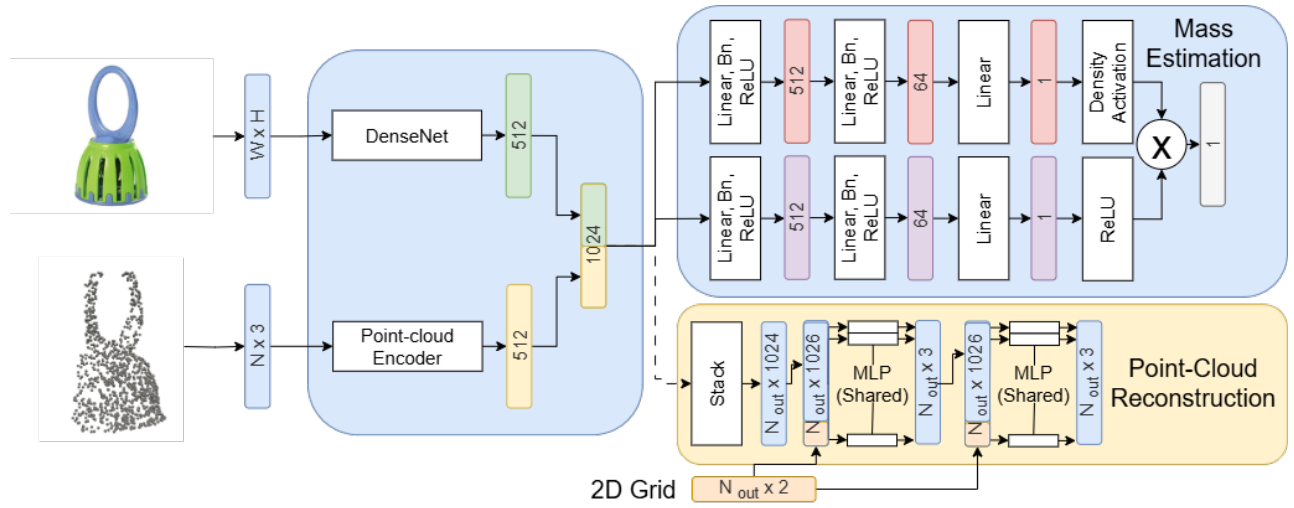


Fig. 1. Our architecture for mass estimation. The leftmost blue block works as an encoder for the point cloud and RGB, and rightmost blue block works as the mass decoder. The yellow block is a FoldingNet architecture, which can be integrated optionally.




Fig. 2. The PointNet architecture.




Fig. 3. The DGCNN architecture.

A. PointNet

Inspired by the PointNet [16] architecture, this encoder comprises a series of transform blocks that apply a linear transformation to each point separately, as shown in Figure 2. Inside the same block, the linear transformation is the same for all points; i.e. they share the same parameters. This ensures the network remains invariant to the order of the input points in the point-cloud since point-clouds are naturally unordered data structures. The model progressively transforms the input into higher-dimensional feature spaces until it reaches the maximum feature size of 512 features, at which point a max pool operation is performed to the input, resulting in a 1D vector that works as an embedding.

B. DGCNN

In our second architecture, inspired by the DGCNN (Dynamic Graph CNN) [17] paper, illustrated in Figure 3, is an evolution of the PointNet model that introduces the concept

of dynamically computing local geometric structures via k-nearest-neighbour (k-NN) graphs based on the Euclidean distance. Its feature transformation is very similar to PointNet, but each point x_i is stacked k times, and then, for each neighbour x_j , their difference is appended $[x_i - x_j, x_i]$. Then, in a similar fashion to PointNet, a shared linear layer works independently on every point; however, now there are $N \times k$ points. A MaxPool operation is used to aggregate the linear layer's output. Finally, DGCNN also makes use of residual connections from the outputs of first layers to the last, allowing the network to retain both fine-grained local details and higher-level abstract features.

C. PointTransformer

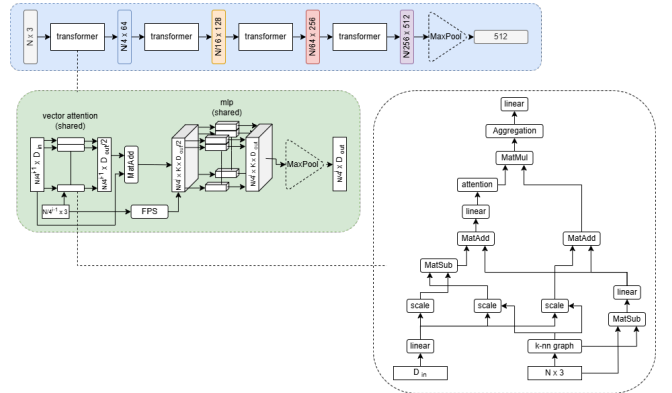


Fig. 4. The architecture of the PointTransformer. A detailed close-up of the transformer block is shown on the bottom left, and to its right, the explicit vector attention mechanism is depicted.

Finally, the PointTransformer [21], shown in Figure 4, employs a vector attention mechanism to capture descriptions for each point and its local neighbourhood. Vector attention distinguishes itself from traditional attention by considering the k-NN of each point instead of just the point itself. This

vector attention is applied individually to each point, with shared weights. In contrast to the other two previous models discussed, PointTransformer also progressively downsamples the point-cloud by a quarter at each transformation block, resulting in fewer points with more descriptive features at each step. Finally, similarly to the previous two models, a MaxPool layer is applied to aggregate the features of the local neighbour into a single feature.

D. DenseNet

DenseNet is a CNN with dense connectivity. It is divided into sets of dense blocks, in which each layer is connected to every other layer in a feed-forward manner. Additionally, DenseNet uses bottleneck layers (typically 1x1 convolutions) to reduce the number of input feature maps, significantly improving computational efficiency without sacrificing performance.

E. Density Decoder

The density decoder is a Fully Connected network acting as a regressor for density. Since typical density values are large (in the range of 500 Kg/m³ and above), we fitted a function to the density distribution in the *image2mass* dataset and used it as the activation function for the final layer [9]. This custom activation introduces an inductive bias, guiding the network toward realistic density estimates. While this bias could potentially limit generalizability to datasets with different density distributions, we argue that the large and diverse nature of the *image2mass* dataset provides a close approximation of real-world object densities when considering similar household settings. Furthermore, this approach reduces implausible predictions and stabilizes training as the network learns to fit realistic density values early on.

F. Volume Decoder

The volume decoder shares the same structure with the density decoder, except for the fact it has a ReLU as the activation function of the last layer. This results in volumes only having positive values. During training, the density decoder activation function will also guide this network since the two are multiplied to get the final output.

G. FoldingNet

The folding net serves as a point-cloud reconstruction decoder. It deforms a grid of points by sequentially applying a series of Multi-Layer Perceptrons (MLPs) to each point separately, which translates the point to a new 3D position. The MLPs use the latent vector predicted by the encoders as additional context for the deformation. Furthermore, the same MLPs are used for all points. The output of the FoldingNet is a coarse reconstruction of the input point-cloud.

H. Mass output

The model computes mass as the product of the outputs from the density and volume decoders. Since the density decoder typically outputs values several orders of magnitude larger than the volume decoder, we introduce a balancing constant b to scale the two outputs appropriately. By selecting b such that $mass = (density \cdot b) \cdot (volume/b)$, and ensuring that $(density \cdot b)$ and $(volume/b)$ are of similar magnitude, the model learns more effectively. Drawing from the approach used in [9], we found that a constant $b \approx 16.5$ yields good results.

I. Metrics

We use several scale-invariant metrics to evaluate the model’s mass predictions, accounting for the wide variation in object masses. All metrics presented are symmetrically for both over- and under-estimating.

- **Absolute Log Difference Error (ALDE):** $ALDE = |\ln(y) - \ln(\hat{y})|$. This metric measures the difference between the logarithms of the true mass (y) and the predicted mass (\hat{y}).
- **Absolute Percentage Error (APE):** $APE = \left| \frac{y - \hat{y}}{y} \right|$. APE captures the relative error between the true and predicted masses.
- **Minimum Ratio Error (MnRE):** $MnRE = \min \left(\left| \frac{\hat{y}}{y} \right|, \left| \frac{y}{\hat{y}} \right| \right)$. This metric measures how much the prediction deviates from the true value in terms of multiplicative factors, making it easy to interpret errors in practical terms. The higher the MnRE, the closer the prediction is to the true value.
- **q-metric (off by a factor of 2):** $q =$ Percentage of predictions where $MnRE < 0.5$. This metric measures the percentage of predictions that are off by a factor of 2 or more. It provides some insight into the consistency and reliability of the model.

J. Losses

We implement two distinct loss functions depending on whether we are training solely for mass estimation or also for point-cloud reconstruction.

For mass estimation, we use the ALDE loss. We believe ALDE is preferred over metrics like APE because it is less sensitive to extreme values, which can disproportionately penalize errors in smaller mass predictions.

For the task of point-cloud reconstruction, we use the Chamfer Distance, which measures the average distance between points in the predicted and ground-truth point-clouds. This metric is defined as:

$$CD = \frac{1}{|P|} \sum_{p \in P} \min_{p' \in P'} \|p - p'\|^2 + \frac{1}{|P'|} \sum_{p' \in P'} \min_{p \in P} \|p' - p\|^2,$$

where P and P' are the sets of points in the predicted and ground-truth point-clouds, respectively.

When training the model for both mass estimation and point-cloud reconstruction, we combine these two loss functions into a single objective:

Total Loss = ALDE + $\lambda \cdot$ Chamfer Distance.

Here, λ is a constant factor used to balance the relative contributions of the ALDE and Chamfer Distance components in the total loss function. From empirical experiments, we found that setting $\lambda \approx 1$ provides a good balance between mass estimation accuracy and point-cloud reconstruction quality.

V. RESULTS AND DISCUSSION

A. Depth Prediction

We train our GLPDepth model on our synthetic dataset constructed from ShapeNetSem. This dataset includes 8,948 distinct 3D models, each rendered from 14 different viewpoints. To ensure proper evaluation, we divided the dataset into training and testing sets using a 90-10 split, yielding 8,053 models for training and 895 3D models for testing. This corresponds to 112,742 training samples and 12,530 test samples, with no overlap between the 3D models in the train and test sets. This separation ensures that the objects seen during training are not present in the test set.

We use several metrics to evaluate GLPDepth dense depth map reconstructions:

- **Mean Absolute Percentage Error:**

$$mAPE = \sum_{i=1}^N \left| \frac{y - \hat{y}}{y} \right|.$$

- **Mean Square Percentage Error:**

$$mSPE = \sum_{i=1}^N \left(\frac{y - \hat{y}}{y} \right)^2.$$

- **Root Mean Square Error:**

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}.$$

- **Root Mean Square Logarithmic Error:**

$$RMSE_{log} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\log(y_i + 1) - \log(\hat{y}_i + 1))^2}.$$

- **Log10 Error:**

$$Log10 = \frac{1}{N} \sum_{i=1}^N |\log_{10}(y_i) - \log_{10}(\hat{y}_i)|.$$

- **Scale Invariant Logarithmic Error:**

$$SILog = \frac{1}{N} \sum_{i=1}^N (\log(y_i) - \log(\hat{y}_i))^2 - \frac{1}{N^2} \left(\sum_{i=1}^N \log(y_i) - \log(\hat{y}_i) \right)^2.$$

- **Absolute Log Difference Error:**

$$ALDE = |\ln(y) - \ln(\hat{y})|.$$

- **Absolute Percentage Error:**

$$APE = \left| \frac{y - \hat{y}}{y} \right|.$$

- **Minimum Ratio Error:**

$$MnRE = \min \left(\left| \frac{\hat{y}}{y} \right|, \left| \frac{y}{\hat{y}} \right| \right).$$

- **q-metric (off by a factor of 2):**

q = Percentage of predictions where $MnRE < 0.5$.

- **Chamfer Distance:**

$$CD = \frac{1}{|P|} \sum_{p \in P} \min_{p' \in P'} \|p - p'\|^2 + \frac{1}{|P'|} \sum_{p' \in P'} \min_{p \in P} \|p' - p\|^2.$$

We use the SILog loss function for training, and fine-tuned two GLPDepth NYU-v2 [23] checkpoints—one without depth normalization and another with depth normalization. Both models were fine-tuned for 5 epochs. The results of these fine-tunings are presented in Table I.

The results demonstrate that the normalization procedure significantly enhanced the model’s performance, as



Fig. 5. High definition point-clouds generated from depth predictions of the *image2mass* dataset using the Glpdepth model.

TABLE I
GLPDEPTH PERFORMANCE ON SYNTHETIC TEST DATASET (12,530 EXAMPLES)

Model	mAPE (↓)	mSPE (↓)
GLPDepth fine-tune	0.2028	0.1682
GLPDepth fine-tune + normalization	0.0793	0.0181
GLPDepth Best (20 epoch)	0.0718	0.0155
Model	RMSE (↓)	RMSE Log (↓)
GLPDepth fine-tune	0.4718	0.2119
GLPDepth fine-tune + normalization	0.1175	0.0813
GLPDepth Best (20 epoch)	0.1056	0.0741
Model	Log ₁₀ (↓)	SILog (↓)
GLPDepth fine-tune	0.0880	0.1582
GLPDepth fine-tune + normalization	0.0329	0.0629
GLPDepth Best (20 epoch)	0.0302	0.0567

evidenced by superior results across all metrics by a wide margin. This intuitively makes sense since it leveraged the 3D structure of the data to resolve the scale ambiguity of the problem. We further fine-tuned the GLPDepth model with normalization for an additional 15 epochs (GLPDepth Best), resulting in even better performance. Although the improvements were evident across all metrics, the difference between the 20-epoch model and the 5-epoch version was less pronounced.

When applied to the *image2mass* dataset, the model’s depth predictions, as shown in Figure 5, were qualitatively impressive, providing realistic 3D projections of many objects. However, we also noticed the model struggled with objects with highlights, likely due to the synthetic dataset’s lack of highlight simulation and reflective materials, which probably impacted generalization in such cases.

B. Mass estimation

We trained five variations of our original architecture and compared them against the benchmarks from the *image2mass* paper:

- **DenseNet:** A baseline setup where no depth image is provided to the model, and no point-cloud encoder is used.
- **DenseNet + PointNet:** A setup using FoldingNet as the point-cloud encoder.
- **DenseNet + DGCNN:** A setup using DGCNN as the point-cloud encoder.
- **DenseNet + PointTransformer:** A setup using PointTransformer as the point-cloud encoder.
- **DenseNet + PointNet + FoldingNet:** A setup that tests whether simultaneously learning point-cloud reconstruction helps improve the mass prediction capabilities of the model.

All models were trained on the same data split as the original *image2mass* experiments, comprising 141,950 training samples and 924 test samples. Importantly, the augmented *image2mass* dataset used in our models contains identical image and mass labels as the original but includes additional point-cloud data.

In the DenseNet setup, to compensate for the absence of a point-cloud encoder, we increase the size of the DenseNet output latent vector to 1024. In the DenseNet + PointNet + FoldingNet configuration, our synthetic *ShapeNetSem* dataset was used in conjunction with the augmented *image2mass* dataset. Since the *image2mass* dataset lacks 3D models necessary for point-cloud reconstruction, we employed a dynamic batching strategy. Before each training epoch, batches were prepared from both datasets and randomly selected during training. For batches from *ShapeNetSem*, the loss function included both Chamfer Distance and ALDE; for batches from the *image2mass* dataset, only ALDE was used.

To ensure robustness during training, several augmentation techniques were applied. Images were randomly flipped horizontally, vertically, or both, with corresponding transformations applied to the point-clouds. For models with point-cloud encoders, depth images were converted into sparse point-clouds of 1024 points. If the point-cloud had more than 1024 points, random sampling was employed to reduce the number of points to 1024. Conversely, if fewer points were available, additional points were placed at the origin to reach 1024. The point-clouds were centered by subtracting the mean coordinates, computed either after downsampling, or at the beginning of the pre-processing for smaller point-clouds.

TABLE II

MODEL PERFORMANCE ON THE *image2mass* TEST SET (924 ITEMS)

Model	ALDE (\downarrow)	APE (\downarrow)	MnRE (\uparrow)	q (\uparrow)
(i2m) Xception k-NN	0.914	2.762	0.504	0.488
(i2m) Xception k-NN (SIM)	0.677	1.14	0.570	0.584
(i2m) Pure CNN	0.550	0.711	0.634	0.696
(i2m) CNN with ALDE	0.592	0.849	0.616	0.672
(i2m) No Geometry	0.490	0.613	0.657	0.746
(i2m) Shape-aware	0.470	0.651	0.672	0.767
DenseNet	0.623	0.879	0.612	0.663
DenseNet + PointTransformer	0.477	0.624	0.671	0.757
DenseNet + DGCNN	0.448	0.607	0.682	0.783
DenseNet + PointNet	0.430	0.522	0.698	0.800
DenseNet + PointNet + FoldingNet	0.435	0.546	0.695	0.805

Table II shows the performance of our models on the *image2mass* dataset. Models starting with (i2m) refer to the *image2mass* benchmarks [9]. Our PointNet and DGCNN models outperform the original *image2mass* architecture across all metrics. Among the point-cloud encoders tested, PointNet achieved the best performance, which is somewhat surprising, given that it’s the simplest of the three models and it doesn’t consider local connectivity like the others do. Additionally, the DenseNet-only baseline performed considerably worse than all other proposed models, underscoring the critical role that depth plays in mass estimation.

We also tested PointNet with a FoldingNet, to averigate whether mass estimation and point-cloud reconstruction could aid each other. We found that both PointNet and PointNet + FoldingNet performed very similarly in mass estimation, which might suggest that point-cloud reconstruction and mass prediction are sufficiently different tasks that improvements in one do not directly translate to improvements in the other. However, this result might also be partly caused by the synthetic dataset being the only one with a complete point-cloud for reconstruction.

VI. CONCLUSION

Mass estimation using purely visual methods remains an underexplored area in robotics, with current approaches relying on images and bounding boxes to estimate object mass. However, these methods are often impractical due to their slow performance, requiring multiple views of the object. To address this challenge, we augmented an existing RGB-mass dataset with depth data, and reformulated the task as predicting mass from both depth and RGB inputs. We developed a network that integrates sparse point clouds and RGB images to estimate object mass and tested several point cloud encoders, including PointNet, DGCNN, and PointTransformer. Our model outperformed existing models on the same dataset, with PointNet yielding the best results, unexpectedly surpassing the more complex architectures.

Additionally, we explored whether learning to reconstruct a complete point cloud could improve mass estimation, though the results were inconclusive. As a next step, we aim to deploy this model in practical robotic applications, investigating how planning and control algorithms could leverage mass predictions to enhance performance in real-world tasks such as grasping and manipulation.

REFERENCES

- [1] N. Mavrakis and Rustam Stolkin, “Estimation and exploitation of objects’ inertial parameters in robotic grasping and manipulation: A survey,” *Robotics and autonomous systems*, vol. 124, pp. 103374–103374, Feb. 2020
- [2] J. A. Lines, R. D. Tillett, L. G. Ross, D. Chan, S. Hockaday, and N. J. B. McFarlane, “An automatic image-based system for estimating the mass of free-swimming fish,” *Computers and Electronics in Agriculture*, vol. 2, no. 31, pp. 151–168, 2015
- [3] Y. Yang and G. Teng, “Estimating Pig Weight From 2D Images,” *Springer eBooks*, pp. 1471–1474, Mar. 2008
- [4] S. Amraci, S. Abdanan Mehdizadeh, and S. Sallary, “Application of computer vision and support vector regression for weight prediction of live broiler chicken,” *Engineering in Agriculture, Environment and Food*, vol. 10, no. 4, pp. 266–271, Oct. 2017

- [5] M. Tschärke and T. M. Banhazi, "Review of Methods to Determine Weight and Size of Livestock from Images," *Australian Journal of Multi-Disciplinary Engineering*, vol. 10, no. 1, pp. 1–17, Jan. 2013.
- [6] M. Omid, M. Khojastehnazhand, and A. Tabatabaefar, "Estimating volume and mass of citrus fruits by image processing technique," *Journal of Food Engineering*, vol. 100, no. 2, pp. 315–321, Sep. 2010
- [7] G. Vivek Venkatesh, S. Md. Iqbal, A. Gopal, and D. Ganesan, "Estimation of Volume and Mass of Axi-Symmetric Fruits Using Image Processing Technique," *International Journal of Food Properties*, vol. 18, no. 3, pp. 608–626, Dec. 2014
- [8] D. Bailey, K. Mercer, C. Plaw, R. Ball, and H. Barraclough, "HIGH SPEED WEIGHT ESTIMATION BY IMAGE ANALYSIS," 2004
- [9] T. Standley, O. Sener, D. Chen, and S. Savarese, "image2mass: Estimating the Mass of an Object from Its Image," *PMLR*, pp. 324–333, Oct. 2017
- [10] D. Patel, A. Nath, and Rajdeep Niyogi, "Adding Material Embedding to the image2mass Problem," *Lecture notes in computer science*, pp. 77–90, Jan. 2022
- [11] J. M. L. Andrade and P. Moreno, "Improving the Estimation of Object mass from images," 2023 *IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, pp. 199–206, April 2023
- [12] Dillam Jossue Díaz-Romero, Wouter Sterkens, S. Van, Toon Goedemé, W. Dewulf, and J. Peeters, "Deep learning computer vision for the separation of Cast- and Wrought-Aluminum scrap," *Resources, Conservation and Recycling*, vol. 172, pp. 105685–105685, Sep. 2021
- [13] Dillam Jossue Díaz-Romero, "Simultaneous mass estimation and class classification of scrap metals using deep learning," *Resources, Conservation and Recycling*, vol. 181, pp. 106272–106272, Jun. 2022
- [14] Bingyuan Wu, Yongxiong Wang, "Rich global feature guided network for monocular depth estimation," *Image and Vision Computing*, Volume 125, 2022,
- [15] Manolis Savva, Angel X. Chang, Pat Hanrahan, "Semantically-Enriched 3D Models for Common-sense Knowledge," *CVPR 2015 Workshop on Functionality, Physics, Intentionality and Causality*, 2015
- [16] R. Q. Charles, H. Su, M. Kaichun and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," 2017 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, USA, 2017, pp. 77–85
- [17] Y. Chen, Y. Xu, Y. Xing and G. Liu, "DGCNN Network Architecture With Densely Connected Point Pairs in Multiscale Local Regions for ALS point-cloud Classification," in *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022
- [18] B. Liu, S. Gould and D. Koller, "Single image depth estimation from predicted semantic labels," 2010 *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, San Francisco, CA, USA, 2010, pp. 1253–1260
- [19] D. Kim, W. Ka, P. Ahn, D. Joo, S. Chun, J. Kim, "Global-Local Path Networks for Monocular Depth Estimation with Vertical CutDepth," *Computing Research Repository*, 2022
- [20] G. Huang, Z. Liu, L. Van Der Maaten and K. Q. Weinberger, "Densely Connected Convolutional Networks," 2017 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, USA, 2017, pp. 2261–2269
- [21] H. Zhao, L. Jiang, J. Jia, P. Torr and V. Koltun, "Point Transformer," 2021 *IEEE/CVF International Conference on Computer Vision (ICCV)*, Montreal, QC, Canada, 2021, pp. 16239–16248
- [22] Y. Yang, C. Feng, Y. Shen and D. Tian, "FoldingNet: point-cloud Auto-Encoder via Deep Grid Deformation," 2018 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, UT, USA, 2018, pp. 206–215
- [23] N. Silberman, D. Hoiem, P. Kohli and R. Fergus, "Indoor Segmentation and Support Inference from RGBD Images," *ECCV*, 2012
- [24] F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions," 2017 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, USA, 2017, pp. 1800–1807