

Analyzing Modular CNN Architectures for Joint Depth Prediction and Semantic Segmentation

Omid Hosseini Jafari*, Oliver Groth*, Alexander Kirillov*, Michael Ying Yang[†] and Carsten Rother*

Abstract—This paper addresses the task of designing a modular neural network architecture that jointly solves different tasks. As an example we use the tasks of depth estimation and semantic segmentation given a single RGB image. The main focus of this work is to analyze the cross-modality influence between depth and semantic prediction maps on their joint refinement. While most of the previous works solely focus on measuring improvements in accuracy, we propose a way to quantify the cross-modality influence. We show that there is a relationship between final accuracy and cross-modality influence, although not a simple linear one. Hence a larger cross-modality influence does not necessarily translate into an improved accuracy. We find that a beneficial balance between the cross-modality influences can be achieved by network architecture and conjecture that this relationship can be utilized to understand different network design choices. Towards this end we propose a Convolutional Neural Network (CNN) architecture that fuses the state-of-the-art results for depth estimation and semantic labeling. By balancing the cross-modality influences between depth and semantic prediction, we achieve improved results for both tasks using the NYU-Depth v2 benchmark.

I. INTRODUCTION

Machine perception is an important and recurrent theme throughout the Robotics and Computer Vision community. Computer Vision has contributed a broad range of tasks to the field of perception, such as estimating physical properties from an image, e.g. depth, motion, or reflectance, as well as estimating semantic properties, e.g. labeling each pixel with a semantic class. One may argue that all of these tasks contribute to one central goal, which can be broadly described as “holistic scene understanding”. In the last decade a lot of research effort has focused on solving individual tasks as good as possible. While it is certainly important to gauge the limits of individual tasks, various researchers have recently raised the question of whether the next big step forward can be achieved by focusing on improving single tasks or by considering different tasks in a joint fashion, e.g. [1].¹ This question is particularly emphasized in robotics setups where the coordination of multiple tasks and consolidation of various predictions is constitutive. In this work we focus on the question of “*How to analyze and exploit the cross-modality influence between depth and semantic predictions in order to solve tasks jointly.*”. While the idea of a “beneficial influence between different tasks” is not new, it has in our opinion not

received enough attention. This is in contrast to other fields, such as neuroscience, psychology and machine learning. In principle, there are two different ways to consider multiple tasks in a joint framework. One option is to have one big, joint model. For example this can be a neural network which has a single RGB image I as input and outputs a semantic segmentation S and a depth labeling D ; or a graphical model which represents $p(S, D|I)$. While this is a popular choice and many impressive results have been achieved, e.g. [2] (for depth, semantic segmentation and more) or [3] (for depth, surface reflectance and lighting), it has its drawbacks. Firstly, the models become rapidly complex and are hence rarely used in follow-up works. Secondly, it is very difficult to analyze whether there is indeed a beneficial influence between different tasks. For instance, as we will see later, a joint model may have no interdependency between modalities and tasks and can in fact be considered as two separate models. The second possible approach for solving multiple tasks jointly is to follow a modular design. In this work, we pursue this option. We propose a simple modular design where individual tasks are first inferred separately and then fed into our *joint refinement network* (see Fig. 1). The aim of this network is to leverage a beneficial cross-modality influence between the soft (probabilistic) input modalities in order to jointly refine both task outputs. We show experimentally that there is indeed a relation between the cross-modality influence and an improvement in accuracy for each individual task. However, the relation is not linear, i.e. a larger cross-modality influence does not necessarily mean higher accuracy.

While such a modular design is not as rich as a complex joint model, it brings many advantages: (i) New modalities can be easily integrated. For instance a module that estimates the reflectance properties can be integrated. (ii) We can quantify the cross-modality influence between different modalities, as discussed in detail later. (iii) It is easier to train all the tasks, in contrast to a full joint model. For example, in practice we often have many training images for individual modalities but fewer training images for all modalities jointly. A joint model would have to be trained in a semi-supervised fashion in order to cope with such heterogeneous data, while in a modular architecture each module is trained with the applicable training data. (iv) Since in our case each module, i.e. for the individual task and the joint refinement, is realized in the form of Convolution Neural Networks (CNNs), it is possible to conduct end-to-end training.

The advantages of modular architectures are not new and

*Computer Vision Lab Dresden, Technical University Dresden, Germany
<http://cvlab-dresden.de/people/>

[†]Scene Understanding Group, University of Twente, Netherlands
michael.yang@utwente.nl

¹See for example the recent workshop on “Recognition meets Reconstruction”, where one aim is to solve two tasks jointly.

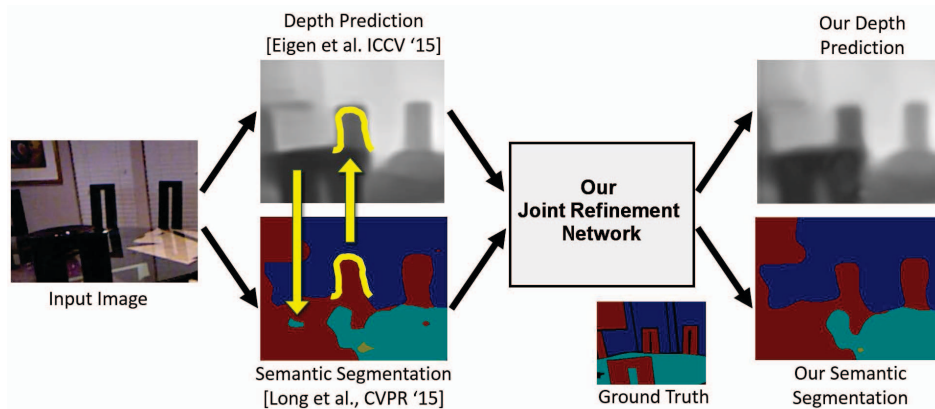


Fig. 1. **Example processing flow of our joint refinement network.** A single RGB image is first processed separately by two state-of-the-art neural networks for depth estimation and semantic segmentation. The two resulting predictions contain information which can mutually improve each other: (1) yellow arrow from depth to semantic segmentation means that a smooth depth map does not support an isolated region (cyan means furniture); (2) yellow arrow from semantic segmentation to depth map means that the exact shape of the chair can improve the depth outline of the chair. (3) In most areas the two modalities positively enforce each other (e.g. the vertical wall (dark blue) supports a smooth depth map). The *cross-modality influences* between these two modalities are exploited by our joint refinement network, which fuses the features from the input prediction maps and jointly processes both modalities for an overall prediction improvement. (*Best viewed in color.*)

indeed David Marr describes it properly in his book ([4] page 102): “This principle [of modular design] is important because if a process is not designed in this way, a small change in one place has consequences in many other places. As a result, the process as a whole is extremely difficult to debug or to improve, whether by a human designer or in the course of natural evolution.”

To summarize, our **main contributions** are threefold:

- For both tasks, semantic segmentation and depth estimation, we improve on the state-of-the-art results for the NYU-Depth v2 benchmark [5]. We achieve this by proposing a new joint refinement network which takes as input the results of the current state-of-the-art networks for the individual tasks.
- For modular architecture designs we propose an experimental setup to measure the cross-modality influence quantitatively. Such experiments are well-known in neuroscience, but have not yet been used in computer vision or robotics, to the best of our knowledge.
- We analyze different network designs with respect to their cross-modality influence and show that there is indeed a relationship between the cross-modality influences and tasks performances. Although not linear, this relationship can be used to understand different design choices in network architectures.

II. RELATED WORK

A large body of work in computer vision has focused on the two separate problems of semantic segmentation and depth estimation. In the review below, we focus on techniques that specifically address multi-modal architectures or perform semantic segmentation and depth estimation from a single monocular image.

Single tasks. Conditional Random Fields (CRFs) are popular models that have been used in both depth estimation task, e.g. [6], [7], [8], [9], [10], [11], and semantic segmentation

task, e.g. [12], [13]. Such approaches predominantly use hand-crafted features. Recently, convolutional neural networks (CNNs) are driving advances in computer vision, such as for image classification [14], object detection [15], [16], recognition [17], [18], semantic segmentation [19], [20], pose estimation [21] and depth estimation [22]. The success of CNNs is attributed to their ability to learn rich feature representations as opposed to hand-designed features. Eigen et al. [22] trained multi-scale CNNs for depth map prediction from a single image. Liu et al. [23] propose deep convolutional neural fields for depth estimation, where a CRF is used to explicitly model the relations of neighboring superpixels, and the potentials are learned in a unified CNN framework. Eigen and Fergus [24] extend their previous method [22] to predict depth, surface normals and semantic labels sequentially with a common multi-scale CNN. A number of recent approaches, including recurrent CNNs (R-CNNs) [25] and fully convolutional networks (FCN) [20] have shown a significant boost in accuracy by adapting state-of-the-art CNN-based image classifiers to the semantic segmentation problem. Pinheiro and Collobert [25] present a feed-forward approach for scene labeling based on an R-CNN. The system is trained in an end-to-end manner over raw pixels and models complex spatial dependencies with low computational cost. FCNs [20] address the coarse-graining effect of the CNN by upsampling the feature maps in deconvolution layers and combining fine-grained and coarse-grained features during prediction.

Joint models. Joint models of multiple tasks have been exploited in the computer vision literature to a certain extent, e.g. joint image segmentation and stereo reconstruction [26], [27], [28], joint object detection and semantic segmentation [29], joint instance segmentation and depth ordering [30], as well as joint intrinsic image, objects, and attributes estimation [2]. However, joint semantic segmentation and depth estimation from a single image has been rarely addressed, with a few exceptions [31], [32]. These

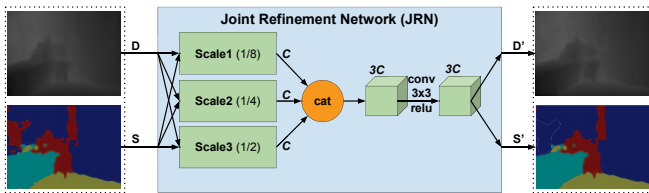


Fig. 2. **Overall Network.** JRN receives as input the predictions of two independent single modalities: Depth and semantic labeling. Inspired by [24] the inputs are considered at different scales (1/8, 1/4 and 1/2 of the total image resolution) in order to capture different levels of details. C is the number of output feature channels from each scale branch (see III-B). After processing the three scale branches, the computed features are concatenated, convolved and then mapped to the two respective output maps.

works explicitly reason about class segmentation as well as depth estimation from a single image. Ladicky *et al.* [31] jointly trained a canonical classifier considering both the loss from semantic and depth labels of the objects. However, they use local regions with hand-crafted features for prediction, which is only able to generate very coarse depth and semantic maps. Wang *et al.* [32] formulate the joint inference problem in a two-layer Hierarchical Conditional Random Field (HCRF). The unary potentials in the bottom layer are pixel-wise depth values and semantic labels, which are predicted by a CNN trained globally using the full image, while the unary potentials in the upper layer are region-wise depth and semantic maps which come from another CNN-based regressor trained on local regions. The mutual interactions between depth and semantic information are captured through the joint training of the CNNs and are further enforced in the joint inference of HCRF. They consider an alternating optimization strategy by minimizing one, fixing the other. In contrast, our model performs full joint inference.

Multi-modal learning and representation. Many different communities have addressed the problem of multi-modal learning and representation, such as machine learning [33], [34], [35], human-computer interaction [36], [37], and neuroscience [38], [39]. In [33], the authors present a series of tasks for multi-modal learning and show how to train deep networks that learn features to address these tasks. In particular, they demonstrate cross modality feature learning, where better features for downstream classification tasks are learned from a video if both audio and video signals are present during the feature learning stage. While [33] deals with an unsupervised feature learning, our approach uses supervised learning. Furthermore, unlike [33] we perform an analysis on the effect of different network architectures on the cross-modality influence. Similarly, in the neuroscience community, the authors of [38] investigated the influence of the *face-benefit* in speech and speaker recognition. Apparently, people who have heard the voice and seen the face of a speaker during training time are more likely to recognize both the speaker and the spoken words from recorded audio only during test time. Additionally, [39] revisited the *face-benefit* experiment and showed a joint audio-visual processing by the brain for the classification task, indicating a joint feature

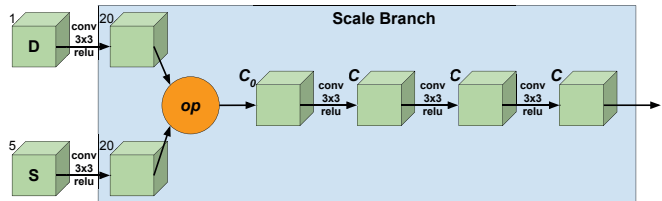


Fig. 3. **A Scale Branch.** On each scale there are first 20-dimensional feature vectors extracted by performing 3×3 convolutions on each input modality. Immediately thereafter these modality features are fused by operation op . We also consider the number of channels C_0 after the fusion operation as a network design variable, which affects the cross-modality influence. The subsequent channel number C is 60 by default (Cat 60, Sum 60). Based on op , C_0 and C we design five different network architectures and analyze their properties in Sec. IV-B.

representation is key to superior performance. Canonical correlation analysis (CCA) [40] is the de-facto approach for learning a common representation of two different modalities (so-called views) in the machine learning literature. Deep CCA, a deep learning version of CCA, is introduced in [41]. It aims at learning a complex non-linear transformation of two views such that the resulting representation is highly correlated. It can be considered as a non-linear extension of the linear CCA.

III. JOINT REFINEMENT NETWORK

In this section, we present the details of the CNN architecture which we used to predict jointly the depth map and the semantic labeling. We also discuss different architectural design decisions and their relation to the cross-modality influence between two modalities.

A. Network Architecture

Our network decomposes into two parts: (i) independent *single-modality* models that output predictions for each modality separately and (ii) our joint refinement network (JRN) that takes as input these prediction maps and outputs refined predictions of all modalities. Our model does not have any constraints on the choice of *single-modality* models. In order to capture dependencies on different scales [24], we employ a multi-scale architecture for JRN, as illustrated in Fig. 2. It has three *scale-branches* Scale1, Scale2 and Scale3 that work with different scales of the input and have the same architecture, described in Fig. 3. On each scale, 20-dimensional features are extracted by performing 3×3 convolutions on each input modality. After each convolutional layer, a ReLU non-linearity is used. In the next section, we consider different architecture designs for the combination.

B. JRN Variants

In order to analyze and understand the cross-modality influence, we design different variations of JRN. Then we measure the cross-modality influence for designed JRN architecture using the approach which we propose in section III-D.

Concatenation. In the *concatenation* architecture we have chosen concatenation operations for op in all scale branches

consistently (cf. Fig. 3). Therefore, the number of channels C_0 after concatenation has doubled to 40. The simple concatenation of the single-modality features aims at achieving cross modality features because the single-modality features are subsequently convolved together. We call this variant *Cat60*. Since concatenation is the weakest form of our possible feature interactions, we subsequently refer to this model as featuring a *loose computation*.

Summation. The *summation* architecture features an element-wise sum as fusion operation *op* on each scale. Since the channels are summed up during fusion, the number of subsequent channels C_0 is equal to 20. We call this variant *Sum60*. In contrast to the *Cat60* architecture this network forces a joint feature representation. We refer to this architecture as having a *coupled computation* later on.

Channel Squeezing. Besides summation of input features we also consider the number of channels used after fusion as an influencing factor on the cross-modality influence and the overall network performance. We hypothesize that fewer channels would generally enforce network to take into account both modalities more during training of the convolutional layers. Directly after the concatenation (*op*) of the input features, which yields a 40-channel feature, the number of channels C can be 10, 5 or 1. This gives us three more network variations *Cat10*, *Cat5* and *Cat1*. We still classify the *Cat10* network as *loose computation* whereas the *Cat5* and *Cat1* networks feature *coupled computation*.

C. JRN Training

We define D' and S' as the depth and semantic label prediction maps and D^* and S^* as the respective ground truth maps. D' and D^* are maps assigning a depth value in the range of $[0.0, 10.0]$ meters to every pixel. S' and S^* are k -channel maps, assigning a probability distribution over k semantic classes to each pixel. We restrict the loss computation to n valid pixels where we have both a depth value and a semantic label as ground truth. The loss function for JRN is a simple summation of single-task losses:

$$L_{joint}(D', S', D^*, S^*) = \underbrace{\frac{1}{n} \sum_i \frac{(D'_i - D^*_i)^2}{D^*_i}}_{L_{depth}} - \underbrace{\frac{1}{n} \sum_i S^*_i \log(S'_i)}_{L_{semantic}} \quad (1)$$

The depth loss L_{depth} is a relative quadratic distance between prediction and ground truth map. For the semantic loss $L_{semantic}$ we use the cross-entropy loss, where $S'_i = \exp[z_i] / \sum_s \exp[z_{i,s}]$ is the class prediction at pixel i given the semantic output slice z of the last convolutional layer of JRN. During training we keep the single-task networks fixed and use their predictions as inputs to JRN. In the future we plan to also perform end-to-end training of all networks, single task networks and JRN. The internal weights of JRN are initialized randomly. We train JRN jointly on both tasks with the standard NYU-Depth v2 [5] train-test split.

D. Quantifying the Cross-Modality Influence

Inspired by the ‘‘face-benefit’’ experiment in [38] we propose an evaluation proxy to measure the influence of a

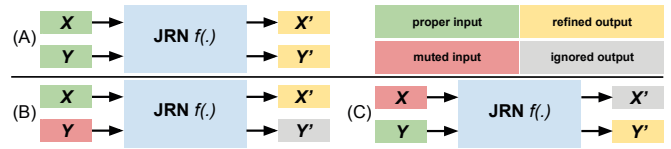


Fig. 4. **Cross-Modality Influence Test.** In order to quantify the cross-modality influence numbers we need to consider three different inference setups: (A) Predicting both outputs X' and Y' from two proper input maps X and Y : $(X', Y') = f(X, Y)$; (B) Muting the input channel Y and computing X' and Y' : $(X', Y') = f(X)$; and (C) Muting the input channel X and computing X' and Y' : $(X', Y') = f(Y)$.

modality on the final model performance and quantify it in a so-called *influence number*. During training time the JRN is trained to jointly predict two outputs X' and Y' from two input modalities X and Y . During inference time we consider three different measurement setups as explained in Fig. 4. The *performance* of the JRN predicts a particular modality X' which is measured by a function A_X (e.g. ‘‘mean IOU’’ for X being semantic labeling or ‘‘rms(linear)’’ for X being a depth map).

The cross-modality influences between the modalities are *directional* and dependent on a particular JRN architecture (represented by its transformation function $f(\cdot)$) as well as two performance functions A_X and A_Y . The cross-modality influence of input modality Y on the performance of the prediction of X' measured under A_X is defined as:

$$w_{Y \rightarrow X'} = A_X(f(X, Y)) - A_X(f(X)). \quad (2)$$

Consequently the complementary influence from X to Y' is defined as:

$$w_{X \rightarrow Y'} = A_Y(f(X, Y)) - A_Y(f(Y)). \quad (3)$$

Note that those influences are *not necessarily symmetric*. We compute the influence values in the setup of a joint semantic segmentation (X) and depth estimation (Y) in Sec. IV-C and analyze the relationship between the cross-modality influences and the model performance.

IV. EXPERIMENTS

In this section, we introduce the dataset, and then describe details of our implementation. After that we present a quantitative and qualitative comparison. We conclude with a cross-modality influence analysis.

A. Experimental setup

Dataset. We evaluate our proposed method on the NYU-Depth V2 dataset [5]. This dataset consists of 1449 RGBD images of indoor scenes, among which 795 are used for training and 654 for test (we use the standard train-test split provided with the dataset). Following [32], we also map the semantic labels into five categories conveying strong geometric properties, i.e. Ground, Vertical, Ceiling, Furniture and Object, as it is shown in Fig. 5.

Implementation details. We use two state-of-the-art single modality CNN models for providing the input depth-map and

semantic segmentation prediction map. For depth input, we use the inference code with pretrained model of Eigen *et al.* [24] which is publicly available². For semantic segmentation prediction maps, we use the FCN model of Long *et al.* [20].

We implement our network in Caffe framework [42]. We train joint refinement networks (Cat60, Cat10, Cat5, Cat1 and Sum60) with 795 training images from NYU-Depth V2 dataset [5] using SGD solver with batches of size one. The learning rate is 0.001 for all the convolutional layers and the momentum is 0.9. The global scale of the learning rate is tuned to a factor of 5. Depending on the different architectures and the number of channels in scale branches, training JRN took 5 to 6 hours using an NVidia GTX Titan X. We pass the absolute depth maps and the semantic prediction maps to JRN.

Evaluation metrics. To evaluate the semantic segmentation, we take Intersection over Union (IOU) and pixel accuracy percentage as metrics. For the depth estimation task we use several measures, which are also commonly used in prior works [9], [22]. Given the predicted depth value of a pixel d_i and the ground truth depth d_i^* , the evaluation metrics are:

- Abs relative error (rel): $\frac{1}{N} \sum_i \frac{|d_i^* - d_i|}{d_i^*}$;
- Squared relative error (rel(sqr)): $\frac{1}{N} \sum_i \frac{|d_i^* - d_i|^2}{d_i^{*2}}$;
- Average \log_{10} error (log10): $\frac{1}{N} \sum_i |\log_{10} d_i^* - \log_{10} d_i|$;
- Root mean squared error (rms(linear)): $\sqrt{\frac{1}{N} \sum_i (d_i^* - d_i)^2}$;
- Root mean squared error (rms(log)): $\sqrt{\frac{1}{N} \sum_i |\log d_i^* - \log d_i|^2}$;
- Accuracy with threshold thr : percentage (%) of d_i s.t. $\max(\frac{d_i^*}{d_i}, \frac{d_i}{d_i^*}) = \delta < thr$, where $thr \in \{1.25, 1.25^2, 1.25^3\}$;

B. Comparison of Results

We first compare our five different JRN architectures described in Sec. III-B with each other (see Table I). We observe that none of the networks consistently outperforms all others in all metrics. However, the Sum60 network is nearly the best for all metrics. Hence we chose it for comparison with other models from related work.

For *depth estimation*, we compare our results with four most recent methods, i.e. Eigen *et al.* [22], Joint HCRF [32], Eigen & Fergus [24], and Liu *et al.* [23]. Table II shows the quantitative results from all the algorithms. Our JRN network consistently outperforms all the state-of-the-art algorithms in all metrics. The main difference to the models of Eigen *et al.* [22], Eigen & Fergus [24] and Liu *et al.* [23] is that they only deal with the depth estimation task, and hence cannot exploit cross-modality influence. However, the Joint HCRF [32] also jointly predicts a depth map and semantic labeling, yet our model outperforms theirs by a large margin both in depth estimation (8.7% rel(sqr) decrease) and in semantic segmentation (10% mean IOU increase, see Table III). We evaluate the published predicted depth maps from Eigen & Fergus [24] and Eigen *et al.* [22] with our evaluation

script, and for Eigen *et al.* [22] we obtain the same reported numbers. However, we do not obtain the same numbers for [24] as reported. Our goal in this work is to improve the performance of the input predictions. Therefore, this comparison is fair since we use the same evaluation script for the input and the output of our network.

For *semantic segmentation*, we compare two recent methods: Our baseline FCN [20] and Joint HCRF [32]. Results are shown in Table III. We outperform the other methods for all five classes. Compared with the baseline FCN [20] our method is 1% better in mean IOU.

Qualitative results of both tasks are shown in Fig. 5. Even though our method does not use superpixels or any explicit CRF model, it tends to produce large homogeneously labeled regions.

C. Performance Cross-Modality Influence Analysis

We compute the cross-modality influence for all five JRN networks by looking at the relation between the cross-modality influence numbers and the performance in the respective modalities (see Fig. 7). We observe that there is no linear relationship between cross-modality influence and performance but they rather lie within an area which is, according to our observations, upper-bounded by a concave curve. This means that a larger influence between modalities does not guarantee better performance in the respective metric. Indeed a large negative effect can hamper performance (see Fig. 6).

Based on our findings about cross-modality influence, we hypothesize that the relationship between cross-modality influence and performance can be generalized into a plot which is sketched in Fig. 8. The cross-modality influence arises from certain model design decisions, as well as from modality combinations for a particular end-task. For example, we have seen in our experiments that moderately transferring shapes and class-wise depth priors from the semantic map into the depth map can help improving depth estimation (see Fig. 5 bottom). However, the cross-modality influence $\omega_{S \rightarrow D'}$ can also be too strong (large positive influence number) which causes a decrease in performance. For example, shapes from semantic segmentation can cause halos and artifacts in the depth map (see Fig. 6). We conclude that inspecting performance vs. cross-modality influence plots is a useful way to find appropriate modular architectures. Furthermore, these plots may help identifying complementary modalities to further leverage the cross-modality influence for task performance improvements.

V. CONCLUSIONS

Inspired by work in neuroscience we have introduced a systematic way to measure the cross-modality influence present in our JRN networks. By doing so, we were able to identify a network which achieves a measurable influence between modalities, has an overall good performance compared to other JRN networks, and is consistently better than the state-of-the-art input modalities.

²<http://www.cs.nyu.edu/~deigen/dn1/>

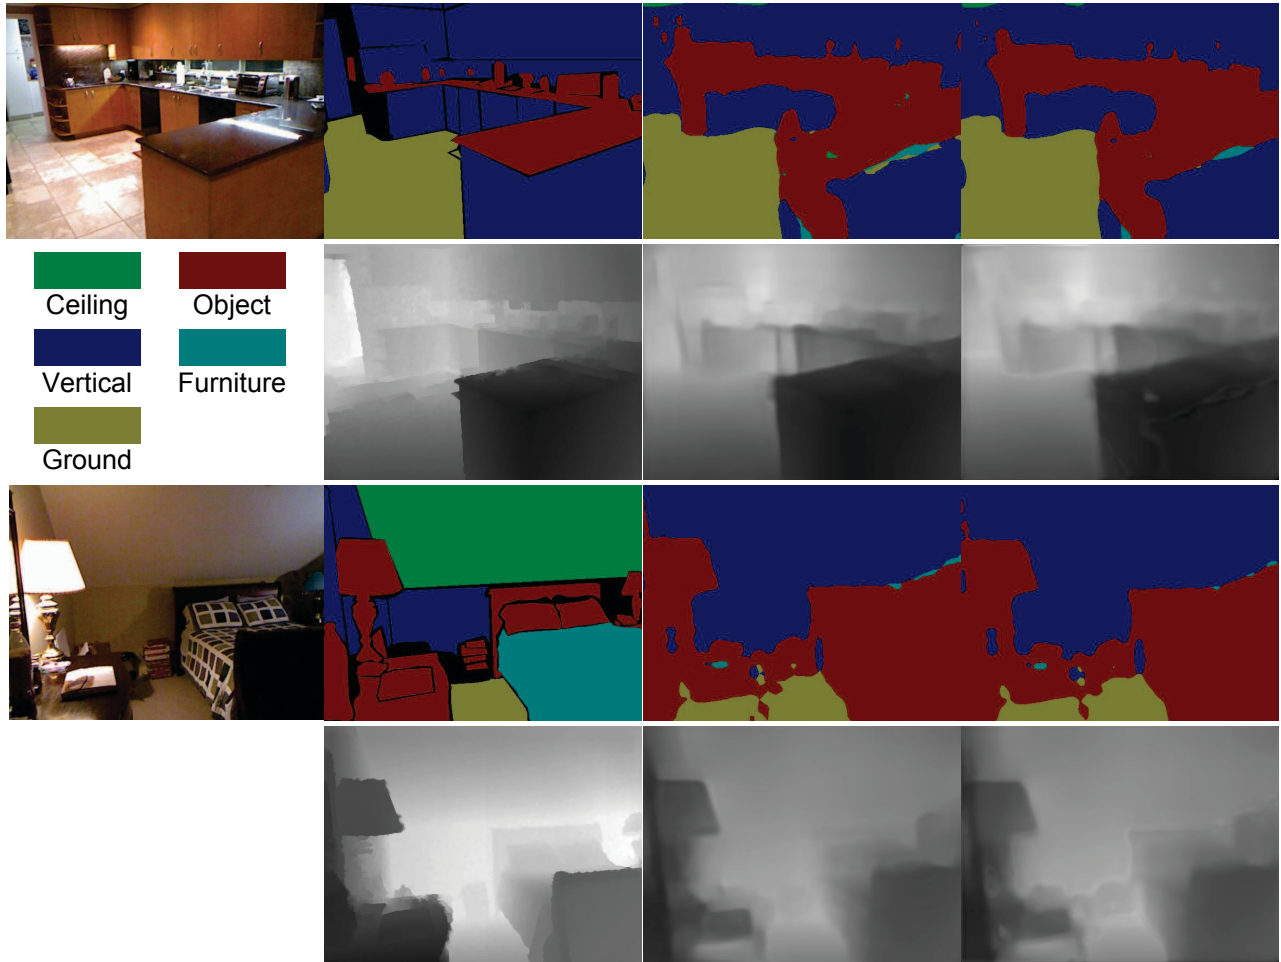


Fig. 5. Two qualitative results of our Sum60 network compared with the input (i.e state-of-the-art). Each top row (from left to right): Image, ground truth semantic labeling, result of [20] and our result. Each bottom row (from left to right): Ground truth depth map, result of [24] and our result. The first example depicts an improvement in semantic labeling, where ground label (yellow) has been removed (next to object label (red)). In the second example, the depth edges of the upper bed frame are better recovered in our result (best viewed zoomed-in). Please note that our results are smooth and follow edges in the input image, despite having no explicit CRF model.

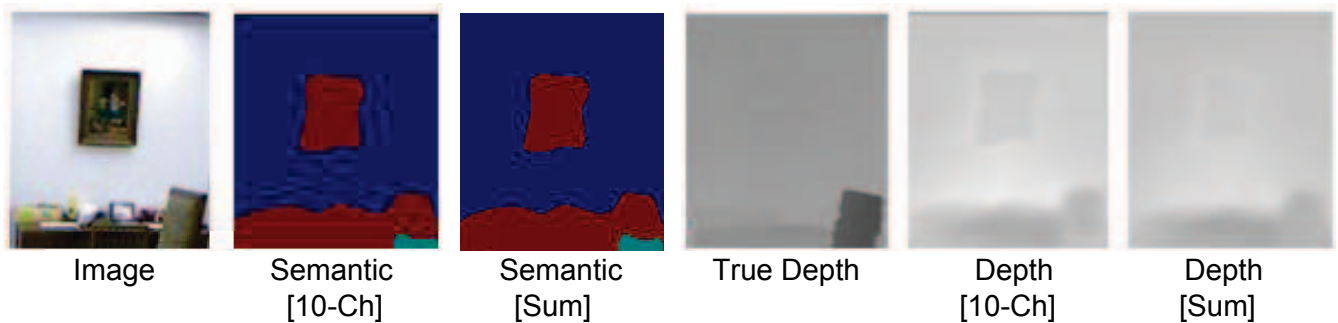


Fig. 6. **Negative influence.** From left to right: Input image, our semantic labeling of the Cat10 network, semantic labeling of the Sum60 network, ground truth depth, result of our Cat10 network, result of our Sum60 network. Despite the Cat10 network having a higher cross-modality influence number $\omega_{S \rightarrow D'}$ than the Sum60 network (cf. Fig. 7 top right), the respective depth accuracy (-rel(sqr)) of the Cat10 network is lower. This is visible in the image where the picture frame has received a wrong depth in the Cat10 network result, compared to the Sum60 network result. (Image crops shown for visualization purpose.)

TABLE I

COMPARISON OF DIFFERENT JRN ARCHITECTURES. WE COMPARE OUR DIFFERENT JRN NETWORKS WITH EACH OTHER AND ALSO WITH OUR INPUT SINGLE MODALITY NETWORKS FOR DEPTH [24] AND SEMANTIC SEGMENTATION [20]. BEST RESULTS ARE SHOWN IN BOLD.

	Error (Depth) (lower is better)					Accuracy (Depth) (higher is better)			Accuracy (Seg.) (higher is better)	
	rel	rel(sqr)	log10	rms(linear)	rms(log)	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$	Mean IOU	Pix.Acc.
Input [24]&[20]	0.158	0.125	0.070	0.687	0.221	0.751	0.946	0.987	53.284	72.268
Cat60	0.158	0.124	0.0686	0.678	0.218	0.760	0.947	0.987	54.206	72.957
Sum60	0.157	0.123	0.068	0.673	0.216	0.762	0.948	0.988	54.184	72.967
Cat10	0.158	0.125	0.069	0.681	0.219	0.756	0.946	0.987	54.080	72.953
Cat5	0.160	0.125	0.068	0.670	0.218	0.762	0.946	0.986	54.120	72.952
Cat1	0.161	0.126	0.069	0.669	0.219	0.759	0.946	0.987	53.989	72.864

TABLE II

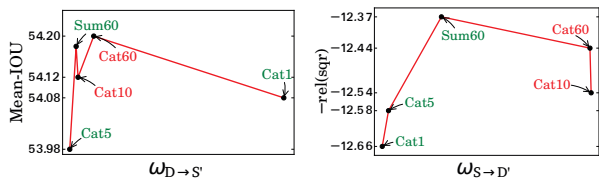
DEPTH COMPARISON. BASELINE COMPARISONS OF DEPTH ESTIMATION ON THE NYU-DEPTH V2 DATASET. OUR METHOD OUTPERFORMS STATE-OF-THE-ART METHODS.

	Error (lower is better)					Accuracy (higher is better)		
	rel	rel(sqr)	log10	rms(linear)	rms(log)	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Eigen et al. [22]	0.215	0.212	-	0.907	0.285	0.611	0.887	0.971
Joint HCRF [32]	0.220	0.210	0.094	0.745	0.262	0.605	0.890	0.970
Eigen & Fergus [24]	0.158	0.125	0.070	0.687	0.221	0.751	0.946	0.987
Liu et al. [23]	0.213	-	0.087	0.759	-	0.650	0.906	0.976
Ours (Sum60)	0.157	0.123	0.068	0.673	0.216	0.762	0.948	0.988

TABLE III

SEMANTIC SEGMENTATION COMPARISON. CLASS-WISE RESULTS REPORTING MEAN IOU AND PIXEL-WISE ACCURACY FOR SEMANTIC SEGMENTATION ON NYU-DEPTH V2 WITH FIVE CLASSES. BEST RESULTS ARE SHOWN IN BOLD.

	Ground	Vertical	Ceiling	Furniture	Object	Mean IOU	Pix.Acc.
Semantic HCRF [32]	61.84	66.344	15.977	26.291	43.121	42.715	69.351
Joint HCRF [32]	63.791	66.154	20.033	25.399	45.624	44.2	70.287
FCN16s NYU-5 [20]	66.578	67.354	46.351	35.71	50.429	53.284	72.268
Ours (Sum60)	67.87	68.707	48.166	35.82	50.770	54.267	73.035



	$\omega_{D \rightarrow S'}$	$\omega_{S \rightarrow D'}$
Cat60	0.94	4.9
Cat10	0.39	4.93
Cat5	0.10	-0.24
Sum60	0.32	1.11
Cat1	7.59	-0.40

Fig. 7. Performance vs. cross-modality influence plots for all models.

We use the mean IOU measure for semantic labels (A_S), and $-\text{rel}(\text{sqr}) \cdot 100$ for depth (A_D). The right table shows all influence numbers $\omega_{D \rightarrow S'}$ and $\omega_{S \rightarrow D'}$ for all models. The top figures are the respective performance-influence plot. Both plots exhibit a peak where the optimal trade-off between cross-modality influence and evaluated performance is achieved. We see that the Sum60 and Cat60 models are at the peaks of the respective plots. We colored models in red which feature *loose computation* and in green which feature *coupled computation* (see Sec. III-B). This supports the idea that the cross-modality influence number can facilitate the systematic exploration of network architectures.

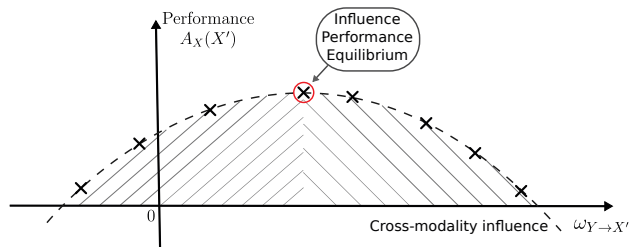


Fig. 8. The performance vs. cross-modality influence curve. For each modality pair (X, Y) and in each modality influence direction $Y \rightarrow X'$ and $X \rightarrow Y'$ the relationship between the magnitude of the cross-modality influence and the prediction performance needs to be balanced into an equilibrium. This cross-modality influence analysis will be helpful when designing models which should operate on multiple modalities and carry out joint predictions.

ACKNOWLEDGEMENTS

This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement No 647769), and German Research Foundation (DFG) YA351/2-1. The authors gratefully acknowledge the support.

REFERENCES

- [1] C. Li, A. Kowdle, A. Saxena, and T. Chen, "Toward holistic scene understanding: Feedback enabled cascaded classification models," *PAMI*, vol. 34, no. 7, pp. 1394–1408, 2012.
- [2] V. Vineet, C. Rother, and P. Torr, "Higher order priors for joint intrinsic image, objects, and attributes estimation," in *NIPS*, 2013, pp. 557–565.
- [3] J. T. Barron and J. Malik, "Shape, illumination, and reflectance from shading," *PAMI*, 2015.
- [4] D. Marr, *Vision*. Freeman, 1982.
- [5] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from rgb-d images," in *ECCV*, 2012.
- [6] A. Saxena, A. Ng, and S. Chung, "Learning Depth from Single Monocular Images," in *NIPS*, 2005.
- [7] A. Saxena, M. Sun, and A. Y. Ng, "Make3d: Learning 3d scene structure from a single still image," *PAMI*, vol. 31, no. 5, pp. 824–840, 2009.
- [8] B. Liu, S. Gould, and D. Koller, "Single image depth estimation from predicted semantic labels," in *CVPR*, 2010, pp. 1253–1260.
- [9] M. Liu, M. Salzmann, and X. He, "Discrete-continuous depth estimation from a single image," in *CVPR*, 2014, pp. 716–723.
- [10] C. Hane, L. Ladicky, and M. Pollefeys, "Direction matters: Depth estimation with a surface normal classifier," in *CVPR*, 2015.
- [11] W. Zhuo, M. Salzmann, X. He, and M. Liu, "Indoor scene structure analysis for single image depth estimation," in *CVPR*, 2015.
- [12] J. Shotton, J. Winn, C. Rother, and A. Criminisi, "Textronboost for image understanding: Multi-class object recognition and segmentation by jointly modeling texture, layout, and context," *IJCV*, vol. 81, no. 1, pp. 2–23, 2009.
- [13] L. Ladicky, C. Russell, P. Kohli, and P. H. S. Torr, "Associative hierarchical random fields," *PAMI*, vol. 36, no. 6, pp. 1056–1077, 2014.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS*, 2012, pp. 1097–1105.
- [15] N. Zhang, J. Donahue, R. B. Girshick, and T. Darrell, "Part-based r-cnns for fine-grained category detection," in *ECCV*, 2014, pp. 834–849.
- [16] S. Gupta, R. Girshick, P. Arbelaez, and J. Malik, "Learning rich features from RGB-D images for object detection and segmentation," in *ECCV*, 2014, pp. 345–360.
- [17] P. Agrawal, R. B. Girshick, and J. Malik, "Analyzing the performance of multilayer neural networks for object recognition," in *ECCV*, 2014, pp. 329–344.
- [18] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *CVPR*, 2014, pp. 1717–1724.
- [19] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *CVPR*, 2014.
- [20] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *CVPR*, 2015.
- [21] A. Toshev and C. Szegedy, "Deeppose: Human pose estimation via deep neural networks," in *CVPR*, 2014, pp. 1653–1660.
- [22] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," in *NIPS*, 2014, pp. 2366–2374.
- [23] F. Liu, C. Shen, G. Lin, and I. Reid, "Learning depth from single monocular images using deep convolutional neural fields," *PAMI*, 2016.
- [24] D. Eigen and R. Fergus, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," in *ICCV*, 2015.
- [25] P. Pinheiro and R. Collobert, "Recurrent convolutional neural networks for scene labeling," in *ICML*, 2014.
- [26] M. Bleyer, C. Rother, P. Kohli, D. Scharstein, and S. Sinha, "Object stereo - joint stereo matching and object segmentation," in *CVPR*, 2011, pp. 3081–3088.
- [27] J.-Y. Guillemaut and A. Hilton, "Joint multi-layer segmentation and reconstruction for free-viewpoint video applications," *IJCV*, vol. 93, no. 1, pp. 73–100, 2011.
- [28] L. Ladicky, P. Sturges, C. Russell, S. Sengupta, Y. Bastanlar, W. Clocksin, and P. Torr, "Joint optimization for object class segmentation and dense stereo reconstruction," *IJCV*, vol. 100, no. 2, pp. 122–133, 2012.
- [29] J. Yao, S. Fidler, and R. Urtasun, "Describing the scene as a whole: Joint object detection, scene classification and semantic segmentation," in *CVPR*, 2012, pp. 702–709.
- [30] Z. Zhang, A. G. Schwing, S. Fidler, and R. Urtasun, "Monocular object instance segmentation and depth ordering with cnns," in *ICCV*, 2015.
- [31] L. Ladicky, J. Shi, and M. Pollefeys, "Pulling things out of perspective," in *CVPR*, 2014.
- [32] P. Wang, X. Shen, Z. Lin, S. Cohen, B. Price, and A. L. Yuille, "Towards unified depth and semantic prediction from a single image," in *CVPR*, 2015.
- [33] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, "Multimodal deep learning," in *ICML*, 2011.
- [34] N. Srivastava and R. R. Salakhutdinov, "Multimodal learning with deep boltzmann machines," in *NIPS*, 2012, pp. 2222–2230.
- [35] S. Chandar, M. M. Khapra, H. Larochelle, and B. Ravindran, "Correlational neural networks," *Neural Computation*, 2015.
- [36] Z. Obrenovic and D. Starcevic, "Modeling multimodal human-computer interaction," *Computer*, vol. 37, no. 9, pp. 65–72, 2004.
- [37] A. Jaimes and N. Sebe, "Multimodal human-computer interaction: A survey," *CVIU*, vol. 108, no. 1-2, pp. 116–134, 2007.
- [38] K. von Kriegstein, D. Ozgr, M. Grter, A.-L. Giraud, C. A. Kell, T. Grter, A. Kleinschmidt, and S. J. Kiebel, "Simulation of talking faces in the human brain improves auditory speech recognition," *PNAS*, vol. 105, no. 18, pp. 6747–6752, 2008.
- [39] S. Schall, S. J. Kiebel, B. Maess, and K. von Kriegstein, "Early auditory sensory processing of voices is facilitated by visual mechanisms," *NeuroImage*, vol. 77, pp. 237 – 245, 2013.
- [40] H. Hotelling, "Relations between two sets of variates," *Biometrika*, vol. 28, pp. 321–377, 1936.
- [41] G. Andrew, R. Arora, J. Bilmes, and K. Livescu, "Deep canonical correlation analysis," in *ICML*, 2013, pp. 1247–1255.
- [42] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.