

Comparatives, Quantifiers, Proportions: A Multi-Task Model for the Learning of Quantities from Vision

Sandro Pezzelle*, Ionut-Teodor Sorodoc*[†], Raffaella Bernardi*[‡]

*CIMEC - Center for Mind/Brain Sciences, University of Trento

[†]Universitat Pompeu Fabra Barcelona

[‡]DISI - Department of Information Engineering and Computer Science, University of Trento

*sandro.pezzelle@unitn.it, [†]ionutteodor.sorodoc@upf.edu

[‡]raffaella.bernardi@unitn.it

Abstract

The present work investigates whether different quantification mechanisms (set comparison, vague quantification, and proportional estimation) can be jointly learned from visual scenes by a multi-task computational model. The motivation is that, in humans, these processes underlie the same cognitive, non-symbolic ability, which allows an automatic estimation and comparison of set magnitudes. We show that when information about lower-complexity tasks is available, the higher-level proportional task becomes more accurate than when performed in isolation. Moreover, the multi-task model is able to generalize to unseen combinations of target/non-target objects. Consistently with behavioral evidence showing the interference of absolute number in the proportional task, the multi-task model no longer works when asked to provide the number of target objects in the scene.

1 Introduction

Understanding and producing sentences like ‘There are *more* cars than parking lots’, ‘*Most* of the supporters wear blue t-shirts’, ‘*Twenty percent* of the trees have been planted last year’, or ‘*Seven* students passed the exam’, is a fundamental competence which allows speakers to communicate information about quantities. Crucially, the type of information conveyed by these expressions, as well as their underlying cognitive mechanisms, are not equivalent, as suggested by evidence from linguistics, language acquisition, and cognition.

First, comparatives (‘more’, ‘less’), quantifiers (‘some’, ‘most’, ‘all’), and proportions (‘20%’, ‘two thirds’) express a comparison or relation *between sets* (e.g., between the set of cars and the set of parking lots). Such relational information is rather coarse when expressed by comparatives and vague quantifiers, more precise when denoted by proportions. In contrast, numbers (‘one’, ‘six’,

‘twenty-two’) denote the exact, absolute cardinality of the items belonging to *one set* (e.g., the set of students who passed the exam).

Second, during language acquisition, these expressions are neither learned at the same time nor governed by the same rules. Recent evidence showed that children can understand comparatives at around 3.3 years (Odic et al., 2013; Bryant, 2017), with quantifiers being learned a few months later, at around 3.4-3.6 years (Hurewitz et al., 2006; Minai, 2006; Halberda et al., 2008). Crucially, knowing the meaning of numbers, an ability that starts not before the age of 3.5 years (Le Corre and Carey, 2007), is not required to understand and use these expressions. As for proportions, they are acquired significantly later, being fully mastered only at the age of 9 or 10 (Hartnett and Gelman, 1998; Moss and Case, 1999; Sophian, 2000).

Third, converging evidence from cognition and neuroscience supports the hypothesis that some important components of these expressions of quantity are grounded on a preverbal, non-symbolic system representing magnitudes (Piazza, 2010). This system, often referred to as Approximate Number System (ANS), is invariant to the sensory modality and almost universal in the animal domain, and consists in the ability of holistically extracting and comparing approximate numerosities (Piazza and Eger, 2016). In humans, it is present since the youngest age, with 6-month-old infants being able to automatically compare sets and combine them by means of proto-arithmetical operations (Xu and Spelke, 2000; McCrink and Wynn, 2004). Since it obeys Weber’s law, according to which highly differing sets (e.g. 2:8) are easier to discriminate than highly similar sets (e.g. 7:8), ANS has been recently claimed to be a *ratio-based* mechanism (Sidney et al., 2017; Matthews et al., 2016). In support of this, behavioral findings indicate that, in non-symbolic




Figure 1: Toy representation of the quantification tasks and corresponding outputs explored in the paper. Note that quantification always refers to animals (target set).

contexts (e.g. visual scenes), proportional values are extracted holistically, i.e. without relying on the pre-computed cardinalities of the sets (Fabbri et al., 2012; Yang et al., 2015). Indeed, people are fairly accurate in providing the proportion of targets in a scene, even in high-speed settings (Healey et al., 1996; Treisman, 2006). Similarly, in briefly-presented scenes, the interpretation of quantifiers is shown to be best described by proportional information (Pezzelle et al., under review).

Altogether, this suggests that performing (1) set comparison, (2) vague quantification, and (3) proportional estimation, which all rely on information regarding relations among sets, underlies increasingly-complex steps of the same mechanism. Notably, such complexity would range from ‘more/less’ judgements to proportional estimation, as suggested by the increasing precision of ANS through years (Halberda and Feigenson, 2008), the reported boundary role of ‘half’ in early proportional reasoning (Spinillo and Bryant, 1991), and the different age of acquisition of the corresponding linguistic expressions. Finally, the ratio-based operation underlying these task would be different from (and possibly conflicting with) that of estimating the absolute numerosity of one set. Indeed, absolute numbers are found to interfere with the access to proportions (Fabbri et al., 2012).

Inspired by this converging evidence, the present work proposes a computational framework to explore various quantification tasks in the visual domain (see Figure 1). In particular, we investigate whether ratio-based quantification tasks can be modeled by a single, multi-task learning neural network. Given a synthetic scene depicting animals (in our setting, the ‘target’ objects) and artifacts (‘non-target’), our model is designed to jointly perform all the tasks by means of an architecture that reflects their increasing complex-

ity.¹ To perform proportional estimation (the most complex), the model builds on the representations learned to perform vague quantification and, in turn, set comparison (the least complex). We show that the multi-task model achieves both higher accuracy and higher generalization power compared to the one-task models. In contrast, we prove that introducing the absolute number task in the loop is not beneficial and indeed hurts the performance.

Our main contribution lies in the novel application and evaluation of a multi-task learning architecture on the task of jointly modeling 3 different quantification operations. On the one hand, our results confirm the interdependency of the mechanisms underlying the tasks of set comparison, vague quantification, and proportional estimation. On the other, we provide further evidence on the effectiveness of these computational architectures.

2 Related Work

2.1 Quantities in Language & Vision

In recent years, the task of extracting quantity information from visual scenes has been tackled via Visual Question Answering (VQA). Given a real image and a natural language question, a VQA computational model is asked to understand the image, the linguistic query, and their interaction to provide the correct answer. So-called *count* questions, i.e. ‘How many Xs have the property Y?’, are very frequent and have been shown to be particularly challenging for any model (Antol et al., 2015; Malinowski et al., 2015; Ren et al., 2015; Fukui et al., 2016). The difficulty of the task has been further confirmed by the similarly poor performance achieved even on the ‘diagnostic’ datasets, which include synthetic visual scenes depicting geometric shapes (Johnson et al., 2017; Suhr et al., 2017).

Using Convolutional Neural Networks (CNN), a number of works in Computer Vision (CV) have proposed specific architectures for counting digits (Seguí et al., 2015), people in the crowd (Zhang et al., 2015a), and penguins (Arteta et al., 2016). With a more cognitive flavor, Chattopadhyay et al. (2017) employed a ‘divide-and-conquer’ strategy to split the image into subparts and count the objects in each subpart by mimicking the ‘subitizing’ mechanism (i.e. numerosities up to 3-4 can be rapidly and accurately appreciated). Inspired by

¹The dataset and the code can be downloaded from github.com/sandropezzelle/multitask-quant

the same cognitive ability is Zhang et al. (2015b), who trained a CNN to detect and count the salient objects in the image. Except Suhr et al. (2017), who evaluated models against various types of quantity expressions (including existential quantifiers), these works were just focused on the absolute number.

More akin to our work is Stoianov and Zorzi (2012), who showed that hierarchical generative models learn ANS as a statistical property of (synthetic) images. Their networks were tested on the task of set comparison ('more/less') and obtained 93% accuracy. A few studies specifically focused on the learning of quantifiers. Sorodoc et al. (2016) proposed a model to assign the correct quantifier to synthetic scenes of colored dots, whereas Sorodoc et al. (2018) operationalized the same task in a VQA fashion, using real images and object-property queries (e.g. 'How many *dogs* are *black*?'). Overall, the results of these studies showed that vague quantification can be learned by neural networks, though the performance is much lower when using real images and complex queries. Finally, Pezzelle et al. (2017) investigated the difference between the learning of cardinals and quantifiers from visual scenes, showing that they require two distinct computational operations. To our knowledge, this is the first attempt to jointly investigate the whole range of quantification mechanisms. Moreover, we are the first to exploit a multi-task learning paradigm for exploring the interactions between set comparison, vague quantification, and proportions.

2.2 Multi-Task Learning

Multi-Task Learning (MTL) has been shown to be very effective for a wide range of applications in machine learning (for an overview, see Ruder (2017)). The core idea is that different and yet related tasks can be jointly learned by a multi-purpose model rather than by separate and highly fine-tuned models. Since they share representations between related (or 'auxiliary') tasks, multi-task models are more robust and generalize better than single-task models. Successful applications of MTL have been proposed in CV to improve object classification (Girshick, 2015), face detection and rotation (Zhang et al., 2014; Yim et al., 2015), and to jointly perform a number of tasks as object detection, semantic segmentation, etc. (Misra et al., 2016; Li and Hoiem, 2016). Though, re-

cently, a few studies applied MTL techniques to either count or estimate the number of objects in a scene (Sun et al., 2017; Sindagi and Patel, 2017), to our knowledge none of them were devoted to the learning of various quantification mechanisms.

In the field of natural language processing (NLP), MTL turned out to be beneficial for machine translation (Luong et al., 2016) and for a range of tasks such as chunking, tagging, semantic role labelling, etc. (Collobert et al., 2011; Søggaard and Goldberg, 2016; Bingel and Søggaard, 2017). In particular, Søggaard and Goldberg (2016) showed the benefits of keeping low-level tasks at the lower layers of the network, a setting which enables higher-level tasks to make a better use of the shared representations. Since this finding was also in line with previous evidence suggesting a natural order among different tasks (Shen and Sarkar, 2005), further work proposed MTL models in which several increasingly-complex tasks are hierarchically ordered (Hashimoto et al., 2017). The intuition behind this architecture, referred to as 'joint many-task model' in the source paper (Hashimoto et al., 2017), as well as its technical implementation, constitute the building blocks of the model proposed in the present study.

3 Tasks and Dataset

3.1 Tasks

Given a visual scene depicting a number of animals (targets) and artifacts (non-targets), we explore the following tasks, represented in Figure 1:

- (a) set comparison (hence, **setComp**), i.e. judging whether the targets are 'more', 'same', 'less' than non-targets;
- (b) vague quantification (hence, **vagueQ**), i.e. predicting the probability to use each of the 9 quantifiers ('none', 'almost none', 'few', 'the smaller part', 'some', 'many', 'most', 'almost all', 'all') to refer to the target set;
- (c) proportional estimation (hence, **propTarg**), i.e. predicting the proportion of targets choosing among 17 ratios, ranging from 0 to 100%.

Tasks (a) and (c) are operationalized as classification problems and evaluated through accuracy. That is, only one answer out of 3 and 17, respectively, is considered as correct. Given the vague status of quantifiers, whose meanings are 'fuzzy' and overlapping, task (b) is evaluated by means



Figure 2: Two scenes included in our dataset. The left-most one depicts a ratio 1:4 (3 animals, 12 artifacts, 15 total items), the rightmost one a ratio 2:3 (6, 9, 15).

of Pearson’s correlation (r) between the predicted and the ground-truth probability vector (cf. § 3.2), for each datapoint.² The overall r is obtained by averaging these scores. It is worth mentioning that we could either evaluate (b) in terms of a classification task or operationalize (a) and (c) in terms of a correlation with human responses. The former evaluation is straightforward and can be easily carried out by picking the quantifier with the highest probability. The latter, in contrast, implies relying on behavioral data assessing the degree of overlap between ground-truth classes and speakers’ choice. Though interesting, such evaluation is less crucial given the discrete, non-overlapping nature of the classes in tasks (a) and (c).

The tasks are explored by means of a MTL network that jointly performs the three quantification operations (see § 4.2). The intuition is that solving the lower-level tasks would be beneficial for tackling the higher-level ones. In particular, providing a proportional estimation (‘80%’) after performing vagueQ (‘most’) and setComp (‘more’) should lead to a higher accuracy in the highest-level task, which represents a further step in complexity compared to the previous ones. Moreover, lower-level tasks might be boosted in accuracy by the higher-level ones, since the latter include all the operations that are needed to carry out the former. In addition to the MTL model, we test a number of ‘one-task’ networks specifically designed to solve one task at a time (see § 4.1).

3.2 Dataset

We built a large dataset of synthetic visual scenes depicting a variable number of animals and artifacts on the top of a neutral, grey background

²We also experimented with Mean Average Error and dot product and found the same patterns of results (not reported).

	train	val	test	total
no. datapoints	11.9K	1.7K	3.4K	17K
% datapoints	70%	10%	20%	100%

Table 1: Number and partitioning of the datapoints.

(see Figure 2). In doing so, we employed the same methodology and materials used in Pezzelle et al. (under review), where the use of quantifiers in *grounded* contexts was explored by asking participants to select the most suitable quantifier for a given scene. Since the category of animals was always treated as the ‘target’, and that of artifacts as the ‘non-target’, we will henceforth use this terminology throughout the paper. The scenes were automatically generated by an in-house script using the following pipeline: (a) Two natural images, one depicting a target object (e.g. a butterfly) and one depicting a non-target (e.g. a mug) were randomly picked up from a sample of the dataset by Kiani et al. (2007). The sample was obtained by Pezzelle et al. (under review), who manually selected pictures depicting whole items (not just parts) and whose color, orientation and shape were not deceptive. In total, 100 unique instances of animals and 145 unique instances of artifacts were included; (b) The proportion of targets in the scene (e.g. 20%) was chosen by selecting one among 17 pre-defined *ratios* between targets:non-targets (e.g. 1:4, ‘four non-targets to one target’). Out of 17 ratios, 8 were positive (targets > 50%), 8 negative (targets < 50%), and 1 equal (targets = 50%); (c) The absolute number of targets/non-targets was chosen to equally represent the various combinations available for a given ratio (e.g., for ratio 1:4: 1-4, 2-8, 3-12, 4-16), with the constraint of having a number of total objects in the scene (targets+non-targets) ranging from 3 to 20. In total, 97 combinations were represented in the dataset, with an average of 5.7 combinations/ratio (min 2, max 18); (d) To inject some variability, the instances of target/non-target objects were randomly resized according to one of three possible sizes (i.e. medium, big, and small) and flipped on the vertical axis before being randomly inserted onto a 5*5-cell virtual grid. As reported in Table 1, 17K scenes balanced per ratio (1K scenes/ratio) were generated and further split into train (70%), validation (10%), and test (20%).

Ground-truth classes for the tasks of setComp and propTarg were automatically assigned to each scene while generating the data. For vagueQ,

we took the probability distributions obtained on a dataset of 340 scenes by Pezzelle et al. (under review) and we applied them to our datapoints, which were built in the exact same way. These probability distributions had been collected by asking participants to select, from a list of 9 quantifiers (reported in § 3.1), the most suitable one to describe the target objects in a visual scene presented for 1 second. In particular, they were computed against the proportion of targets in the scene, which in that study was shown to be the overall best predictor for quantifiers. To illustrate, given a scene containing 20% of targets (cf. leftmost panel in Figure 2), the probability of choosing ‘few’ (ranging from 0 to 1) is 0.38, ‘almost none’ 0.27, ‘the smaller part’ 0.25, etc. It is worth mentioning that, for scenes containing either 100% or 0% targets the probability of choosing ‘all’ and ‘none’, respectively, is around 1. In all other cases, the distribution of probabilities is fuzzier and reflects the largely overlapping use of quantifiers, as in the example above. On average, the probability of the most-chosen quantifier across ratios is 0.53. Though this number cannot be seen as a genuine inter-annotator agreement score, it suggests that, on average, there is one quantifier which is preferred over the others.

4 Models

In this section, we describe the various models implemented to perform the tasks. For each model, several settings and parameters were evaluated by means of a thorough ablation analysis. Based on a number of factors like performance, speed, and stability of the networks, we opted for using ReLU nonlinear activation at all hidden layers and the simple and effective Stochastic Gradient Descent (SGD) as optimizer ($\text{lr} = 0.01$). We run each model for 100 epochs and saved weights and parameters of the epoch with the lowest validation loss. The best model was then used to obtain the predictions in the test set. All models were implemented using Keras.³

4.1 One-Task Models

We implemented separate models to tackle one task at a time. For each task, in particular, both a network using ‘frozen’ (i.e. pretrained) visual features and one computing the visual features in an ‘end-to-end’ fashion were tested.

³<https://keras.io/>

One-Task-Frozen These models are simple, 2-layer (ReLU) Multi-Layer Perceptron (MLP) networks that take as input a 2048-d frozen representation of the scene and output a vector containing softmax probability values. The frozen representation of the scene had been previously extracted using the state-of-art Inception v3 CNN (Szegedy et al., 2016) pretrained on ImageNet (Deng et al., 2009). In particular, the network is fed with the average of the features computed by the last Convolutional layer, which has size 25×2048 .

One-Task-End2end These models are MLP networks that take as input the 203×203 -pixel image and compute the visual features by means of the embedded Inception v3 module, which outputs 25×2048 -d vectors (the grey and colored box in Figure 1). Subsequently, the 25 feature vectors are reduced twice via ReLU hidden layers, then concatenated, reduced (ReLU), and fed into a softmax layer to obtain the probability values.

4.2 Multi-Task Model

The `multi-task-prop` model performs 3 tasks at the *same time* with an architecture that reproduces in its *order* the conjectured complexity (see Figure 3 and its caption for technical details). The model has a core structure, represented by layers 1-5 in the figure, which is *shared* across tasks and trained with multiple outputs. In particular, (a) layers 1, 2, and 3 are trained using information regarding the output of all 3 tasks. That is, these layers are updated three times by as many back-propagation passes: One on the top of `setComp` output, the second on the top of `vagueQ` output, the third on the top of `propTarg` output; (b) layers 4 and 5 are affected by information regarding the output of `vagueQ` and `propTarg`, and thus updated twice; (c) layers 6 and 7 are updated once, on the top of the output of `propTarg`. Importantly, the three lower layers in Figure 3 (concatenation, ReLU, softmax) are not shared between the tasks, but specialized to output each a specific prediction. As can be noted, the order of the tasks reflects their complexity, since the last task in the pipeline has 2 more layers than the preceding one and 4 more than the first one.

5 Results

Table 2 reports the performance of each model in the various tasks (note that the lowest row and the rightmost column report results described

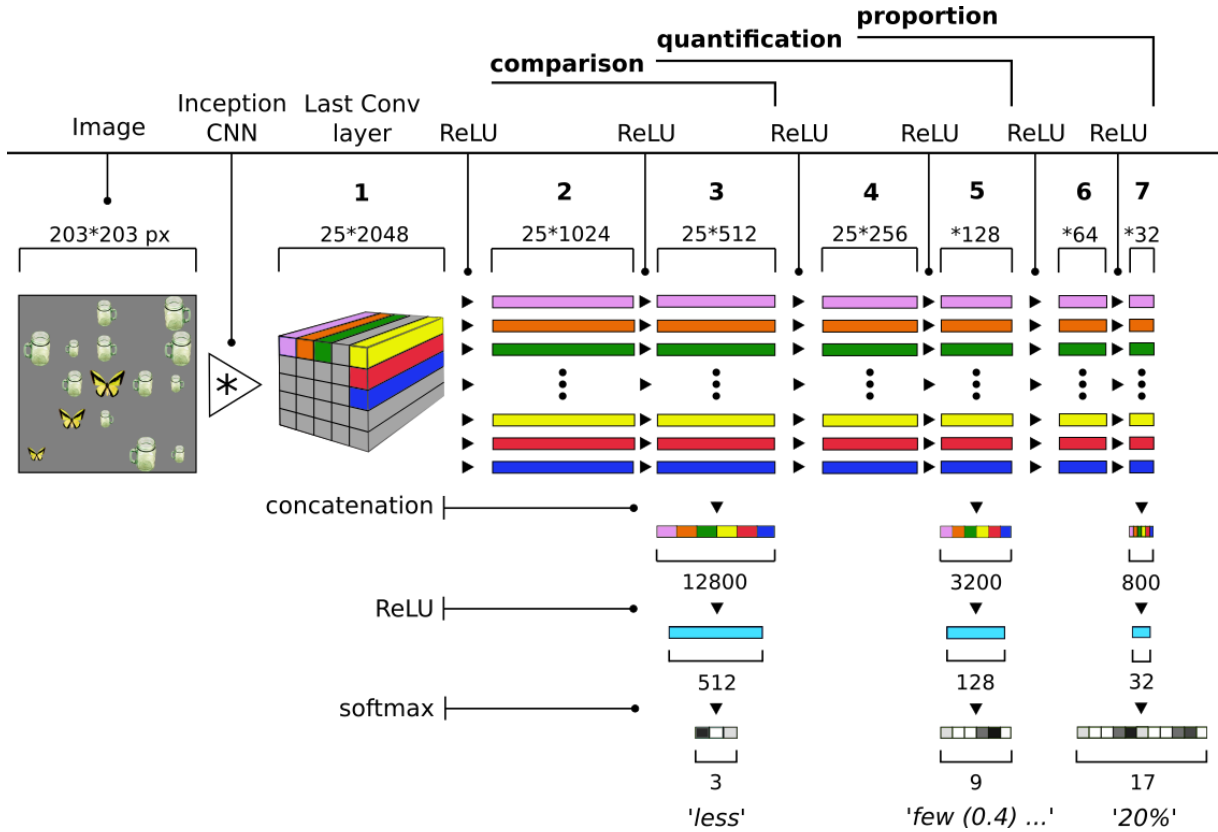


Figure 3: Architecture of the `multi-task-prop` model jointly performing (a) set comparison, (b) vague quantification, and (c) proportional estimation. Given a 203×203 -pixel image as input, the model extracts a 25×2048 representation from the last Convolutional layer of the Inception v3. Subsequently, the vectors are reduced twice via ReLU hidden layers to 1024 and 512 dimensions. The 512-d vectors are concatenated and reduced, then a softmax layer is applied to output a 3-d vector with probability distributions for task (a). The same structure (i.e., 2 hidden layers, concatenation, reduction, and softmax) is repeated for tasks (b) and (c). All the tasks are trained with cross-entropy. To evaluate tasks (a) and (c), in testing, we extract the highest-probability class and compute **accuracy**, whereas task (b) is evaluated via **Pearson’s correlation** against the 9-d ground-truth probability vector.

in § 6.1). In `setComp`, all the models are neatly above chance/majority level (0.47). In particular, the `one-task-end2end` model achieves a remarkable 0.90 acc., which is more than 10% better compared to the simple `one-task-frozen` model (0.78). The same pattern of results can be observed for `vagueQ`, where the Pearson’s correlation (r) between the ground-truth and the predicted probability vector is around 0.96, that is more than 30% over the simpler model (0.62). This gap increases even more in `propTarg`, where the accuracy of the frozen model is more than 40 points below the one achieved by the `one-task-end2end` model (0.21 against 0.66). These results firmly indicate that, on the one hand, the frozen representation of the visual scene encodes little information about the proportion of targets (likely due to the the different task for which they were pretrained,

i.e. object classification). On the other hand, computing the visual features in an end-to-end fashion leads to a significant improvement, suggesting that the network learns to pay attention to features that are helpful for specific tasks.

The most interesting results, however, are those achieved by the multi-task model, which turns out to be the best in all the tasks. As reported in Table 2, sharing the weights between the various tasks is especially beneficial for `propTarg`, where the accuracy reaches 0.92, that is, more than 25 points over the end-to-end, one-task model. An almost perfect performance of the model in this task can be observed in Figure 4, which reports the confusion matrix with the errors made by the model. As can be seen, the few errors are between ‘touching’ classes, e.g. between ratio 3:4 (43% of targets) and ratio 2:3 (40%). Since these classes

model	setComp	vagueQ	propTarg	nTarg
	<i>accuracy</i>	<i>Pearson r</i>	<i>accuracy</i>	<i>accuracy</i>
<i>chance/majority</i>	0.470	0.320	0.058	0.132
one-task-frozen	0.783	0.622	0.210	0.312
one-task-end2end	0.902	0.964	0.659	0.966
multi-task-prop	0.995	0.982	0.918	–
multi-task-number	0.854	0.807	–	0.478

Table 2: Performance of the models in the tasks of set comparison (setComp), vague quantification (vagueQ), proportional estimation (propTarg), and absolute number of targets (nTarg). Values in **bold** are the highest.

differ by a very small percentage, we gain indirect evidence that the model is learning some kind of proportional information rather than trivial associations between scenes and orthogonal classes.

To further explore this point, one way is to inspect the last layer of the proportional task (i.e. the 32-d turquoise vector in Figure 3). If the vectors contain information regarding the proportion of targets, we should expect scenes depicting the same proportion to have a similar representation. Also, scenes with similar proportions (e.g. 40% and 43%) would be closer to each other than are scenes with different proportions (e.g. 25% and 75%). Figure 5 depicts the results of a two-dimensional PCA analysis performed on the vectors of the last layer of the proportional task (the 32-d vectors).⁴ As can be noted, scenes depicting the same proportion clearly cluster together, thus indicating that using these representations in a retrieval task would lead to a very high precision. Crucially, the clusters are perfectly ordered with respect to proportion. Starting from the purple cluster on the left side (90%) and proceeding clockwise, we find 83% (green), 80% (turquoise),

75% (brown), and so on, until reaching 10% (light blue). Proportions 0% (blue) and 100% (yellow) are neatly separated from the other clusters, being at the extremes of the ‘clock’.

An improvement in the results can be also observed for setComp and vagueQ, where the model achieves 0.99 acc. and 0.98 r , respectively. Figure 6 reports, for each quantifier, the probability values predicted by the model against the ground-truth ones. As can be seen, the red lines (model) approximate very closely the green ones (humans). In the following section, we perform further experiments to provide a deeper evaluation of the results.

6 In-Depth Evaluation

6.1 Absolute Numbers in the Loop

As discussed in § 1, the cognitive operation underlying setComp, vagueQ, and propTarg is different compared to that of estimating the absolute number of objects included in one set. To investigate whether such dissociation emerges at the computational level, we tested a modified version of our proposed multi-task model where propTarg task

⁴We used <https://projector.tensorflow.org/>




Figure 4: PropTarg. Heatmap reporting the errors made by the multi-task-prop model. Note that labels refer to ratios, i.e. 14 stands for ratio 1:4 (20% targets).




Figure 5: PCA visualization of the last layer (before softmax) of the proportional task in the MTL model.




Figure 6: VagueQ. Probability values predicted by the `multi-task-prop` model against ground-truth probability distributions for each quantifier.

has been replaced with `nTarg`, namely the task of predicting the absolute number of targets. One-task models were also tested to evaluate the difficulty of the task when performed in isolation. Since the number of targets in the scenes ranges from 0 to 20, `nTarg` is evaluated as a 21-class classification task (majority class 0.13).

As reported in Table 2, the accuracy achieved by the `one-task-end2end` model is extremely high, i.e. around 0.97. This suggests that, when learned in isolation, the task is fairly easy, but only if the features are computed *within* the model. In fact, using frozen features results in a quite low accuracy, namely 0.31. This pattern of results is even more interesting if compared against the results of the `multi-task-number` model. When included in the multi-task pipeline, in fact, `nTarg` has a huge, 50-point accuracy drop (0.48). Moreover, both `setComp` and `vagueQ` turn out to be significantly *hurt* by the highest-level task, and experience a drop of around 14 and 17 points compared to the `one-task-end2end` model, respectively. These findings seem to corroborate the incompatibility of the operations needed for solving the tasks.

6.2 Reversing the Architecture

Previous work exploring MTL suggested that defining a hierarchy of increasingly-complex tasks is beneficial for jointly learning related tasks (see § 2.2). In the present work, the order of the tasks was inspired by cognitive and linguistic abilities (see § 1). Though cognitively implau-

model	setComp	vagueQ	propTarg
	<i>accuracy</i>	<i>Pearson r</i>	<i>accuracy</i>
<i>chance/majority</i>	0.470	0.320	0.058
<i>one-task-frozen</i>	0.763	0.548	0.068
<i>one-task-end2end</i>	0.793	0.922	0.059
<i>multi-task-prop</i>	0.943	0.960	0.539

Table 3: Unseen dataset. Performance of the models in each task. Values in **bold** are the highest.

sible, it might still be the case that the model is able to learn even when reversing the order of the tasks, i.e. from the conjectured highest-level to the lowest-level one. To shed light on this issue, we tested the `multi-task-prop` model after reversing its architecture. That is, `propTarg` is now the first task, followed by `vagueQ`, and `setComp`.

In contrast with the pattern of results obtained by the original pipeline, no benefits are observed for this version of MTL model compared to one-task networks. In particular, both `vagueQ` (0.32 *r*) and `propTarg` (0.08 *acc.*) performance are around chance level, with `setComp` reaching just 0.65 *acc.*, i.e. 25 point lower than the `one-task-end2end` model. The pipeline of increasing complexity motivated theoretically is thus confirmed at the computational level.

6.3 Does MTL Generalize?

As discussed in § 2.2, MTL is usually claimed to allow a higher generalization power. To investigate whether our proposed `multi-task-prop` model genuinely learns to quantify from visual scenes, and not just associations between patterns and classes, we tested it with unseen combinations of targets/non-targets. The motivation is that, even in the most challenging `propTarg` task, the model might learn to match a given combination, e.g. 3:12, to a given proportion, i.e. 20%. If this is the case, the model would solve the task by learning “just” to assign a class to each of the 97 possible combinations included in the dataset. If it learns a more abstract representation of the proportion of targets depicted in the scene, in contrast, it should be able to generalize to unseen combinations.

We built an additional dataset using the exact same pipeline described in § 3.2. This time, however, we randomly selected one combination per ratio (17 combinations in total) to be used only for validation and testing. The remaining 80 combinations were used for training. A balanced number of datapoints for each combination were generated in *val/test*, whereas datapoints in training set

were balanced with respect to ratios, by randomly selecting scenes among the remaining combinations. The *unseen* dataset included around 14K datapoints (80% train, 10% val, 10% test). Table 3 reports the results of the models on the unseen dataset. Starting from setComp, we note a similar and fairly high accuracy achieved by the two one-task models (0.76 and 0.79, respectively). In vagueQ, in contrast, the `one-task-end2end` model neatly outperforms the simpler model (0.92 vs. 0.55 r). Finally, in propTarg both models are at chance level, with an accuracy that is lower than 0.07. Overall, this pattern of results suggests that propTarg is an extremely hard task for the separate models, which are not able to generalize to unseen combinations. The `multi-task-prop` model, in contrast, shows a fairly high generalization power. In particular, it achieves 0.54 acc. in propTarg, that is, almost 10 times chance level. The overall good performance in predicting the correct proportion can be appreciated in Figure 7, where the errors are represented by means of a heatmap. The error analysis reveals that end-of-the-scale proportions (0% and 100%) are the easiest, followed by proportions 75% (3:1), 67% (2:1), 50% (1:1), and 60% (3:2). More in general, negative ratios (targets < 50%) are mispredicted to a much greater extent than are positive ones. Moreover, the model shows a bias toward some proportions, that the model seems ‘to see everywhere’. However, the fact that the errors are found among the adjacent ratios (similar proportions) seems to be a convincing evidence that the model learns representations encoding genuine proportional information. Finally, it is worth mentioning that in setComp and vagueQ the model achieves very high results, 0.94 acc. and 0.96 r , respectively.

7 Discussion

In the present study, we investigated whether *ratio-based* quantification mechanisms, expressed in language by comparatives, quantifiers, and proportions, can be computationally modeled in vision exploiting MTL. We proved that sharing a common core turned out to boost the performance in all the tasks, supporting evidence from linguistics, language acquisition, and cognition. Moreover, we showed (a) the increasing complexity of the tasks, (b) the interference of absolute number, and (c) the high generalization power of MTL. These results lead to many additional questions.




Figure 7: PropTarg. Heatmap with the errors made by the `multi-task-prop` model in the unseen dataset.

For instance, can these methods be successfully applied to datasets of real scenes? We firmly believe this to be the case, though the results might be affected by the natural biases contained in those images. Also, is this pipeline of increasing complexity specific to vision (non-symbolic level), or is it shared across modalities, *in primis* language? Since linguistic expressions of quantity are grounded on a non-symbolic system, we might expect that a model trained on one modality can be applied to another, at least to some extent. Even further, jointly learning representations from both modalities might represent an even more natural, human-like way to learn and refer to quantities. Further work is needed to explore all these issues.

Acknowledgments

We kindly acknowledge Gemma Boleda and the AMORE team (UPF), Raquel Fernández and the Dialogue Modelling Group (UvA) for the feedback, advice and support. We are also grateful to Aurélie Herbelot, Stephan Lee, Manuela Piazza, Sebastian Ruder, and the anonymous reviewers for their valuable comments. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 715154). We gratefully acknowledge the support of NVIDIA Corporation with the donation of GPUs used for this research. This paper reflects the authors’ view only, and the EU is not responsible for any use that may be made of the information it contains.



References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual Question Answering. In *Proceedings of the IEEE International Conference on Computer Vision*. pages 2425–2433.
- Carlos Arteta, Victor Lempitsky, and Andrew Zisserman. 2016. Counting in the wild. In *European Conference on Computer Vision*. Springer, pages 483–498.
- Joachim Bingel and Anders Søgaard. 2017. Identifying beneficial task relations for multi-task learning in deep neural networks. *EACL 2017* page 164.
- Peter Bryant. 2017. *Perception and understanding in young children: An experimental approach*, volume 4. Routledge.
- Prithvijit Chattopadhyay, Ramakrishna Vedantam, Ramprasaath R. Selvaraju, Dhruv Batra, and Devi Parikh. 2017. Counting everyday objects in everyday scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12(Aug):2493–2537.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, pages 248–255.
- Sara Fabbri, Sara Caviola, Joey Tang, Marco Zorzi, and Brian Butterworth. 2012. The role of numerosity in processing nonsymbolic proportions. *The Quarterly Journal of Experimental Psychology* 65(12):2435–2446.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *Conference on Empirical Methods in Natural Language Processing*. ACL, pages 457–468.
- Ross Girshick. 2015. Fast R-CNN. In *Proceedings of the IEEE international conference on computer vision*. pages 1440–1448.
- Justin Halberda and Lisa Feigenson. 2008. Developmental change in the acuity of the “Number Sense”: The Approximate Number System in 3-, 4-, 5-, and 6-year-olds and adults. *Developmental psychology* 44(5):1457.
- Justin Halberda, Len Taing, and Jeffrey Lidz. 2008. The development of ‘most’ comprehension and its potential dependence on counting ability in preschoolers. *Language Learning and Development* 4(2):99–121.
- Patrice Hartnett and Rochel Gelman. 1998. Early understandings of numbers: Paths or barriers to the construction of new understandings? *Learning and instruction* 8(4):341–374.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2017. A Joint Many-Task Model: Growing a Neural Network for Multiple NLP Tasks. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Copenhagen, Denmark, pages 446–456.
- Christopher G Healey, Kellogg S Booth, and James T Enns. 1996. High-speed visual estimation using preattentive processing. *ACM Transactions on Computer-Human Interaction (TOCHI)* 3(2):107–135.
- Felicia Hurewitz, Anna Papafragou, Lila Gleitman, and Rochel Gelman. 2006. Asymmetries in the acquisition of numbers and quantifiers. *Language learning and development* 2(2):77–96.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, pages 1988–1997.
- Roosbeh Kiani, Hossein Esteky, Koorosh Mirpour, and Keiji Tanaka. 2007. Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *Journal of neurophysiology* 97(6):4296–4309.
- Mathieu Le Corre and Susan Carey. 2007. One, two, three, four, nothing more: An investigation of the conceptual sources of the verbal counting principles. *Cognition* 105(2):395–438.
- Zhizhong Li and Derek Hoiem. 2016. Learning without forgetting. In *European Conference on Computer Vision*. Springer, pages 614–629.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-task sequence to sequence learning. In *International Conference on Learning Representations (ICLR)*. San Juan, Puerto Rico.
- Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. 2015. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE international conference on computer vision*. pages 1–9.
- Percival G Matthews, Mark Rose Lewis, and Edward M Hubbard. 2016. Individual differences in nonsymbolic ratio processing predict symbolic math performance. *Psychological science* 27(2):191–202.

- Koleen McCrink and Karen Wynn. 2004. Large-number addition and subtraction by 9-month-old infants. *Psychological Science* 15(11):776–781.
- Utako Minai. 2006. *Everyone knows, therefore every child knows: An investigation of logico-semantic competence in child language*. Ph.D. thesis, University of Maryland.
- Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. 2016. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 3994–4003.
- Joan Moss and Robbie Case. 1999. Developing children’s understanding of the rational numbers: A new model and an experimental curriculum. *Journal for research in mathematics education* pages 122–147.
- Darko Odic, Paul Pietroski, Tim Hunter, Jeffrey Lidz, and Justin Halberda. 2013. Young children’s understanding of “more” and discrimination of number and surface area. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 39(2):451.
- Sandro Pezzelle, Raffaella Bernardi, and Manuela Piazza. under review. Probing the mental scale of quantifiers. *Cognition*.
- Sandro Pezzelle, Marco Marelli, and Raffaella Bernardi. 2017. Be precise or fuzzy: Learning the meaning of cardinals and quantifiers from vision. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics, Valencia, Spain, pages 337–342.
- Manuela Piazza. 2010. Neurocognitive start-up tools for symbolic number representations. *Trends in cognitive sciences* 14(12):542–551.
- Manuela Piazza and Evelyn Eger. 2016. Neural foundations and functional specificity of number representations. *Neuropsychologia* 83:257–273.
- Mengye Ren, Ryan Kiros, and Richard Zemel. 2015. Exploring models and data for image question answering. In *Advances in neural information processing systems*. pages 2953–2961.
- Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Santi Seguí, Oriol Pujol, and Jordi Vitria. 2015. Learning to count with deep object features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. pages 90–96.
- Hong Shen and Anoop Sarkar. 2005. Voting between multiple data representations for text chunking. In *Conference of the Canadian Society for Computational Studies of Intelligence*. Springer, pages 389–400.
- Pooja G Sidney, Clarissa A Thompson, Percival G Matthews, and Edward M Hubbard. 2017. From continuous magnitudes to symbolic numbers: The centrality of ratio. *Behavioral and Brain Sciences* 40.
- Vishwanath A Sindagi and Vishal M Patel. 2017. CNN-Based cascaded multi-task learning of high-level prior and density estimation for crowd counting. In *Advanced Video and Signal Based Surveillance (AVSS), 2017 14th IEEE International Conference on*. IEEE, pages 1–6.
- Anders Søgaard and Yoav Goldberg. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. volume 2, pages 231–235.
- Catherine Sophian. 2000. Perceptions of proportionality in young children: matching spatial ratios. *Cognition* 75(2):145 – 170.
- Ionut Sorodoc, Angeliki Lazaridou, Gemma Boleda, Aurélie Herbelot, Sandro Pezzelle, and Raffaella Bernardi. 2016. “Look, some green circles!”: Learning to quantify from images. In *Proceedings of the 5th Workshop on Vision and Language*. pages 75–79.
- Ionut Sorodoc, Sandro Pezzelle, Aurélie Herbelot, Mariella Dimiccoli, and Raffaella Bernardi. 2018. Learning quantification from images: A structured neural architecture. *Natural Language Engineering* page 1–30.
- Alina G Spinillo and Peter Bryant. 1991. Children’s proportional judgments: The importance of “half”. *Child Development* 62(3):427–440.
- Ivlin Stoianov and Marco Zorzi. 2012. Emergence of a ‘visual number sense’ in hierarchical generative models. *Nature neuroscience* 15(2):194–196.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In *55th Annual Meeting of the Association for Computational Linguistics, ACL*.
- Maojin Sun, Yan Wang, Teng Li, Jing Lv, and Jun Wu. 2017. Vehicle counting in crowded scenes with multi-channel and multi-task convolutional neural networks. *Journal of Visual Communication and Image Representation* 49:412–419.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 2818–2826.
- Anne Treisman. 2006. How the deployment of attention determines what we see. *Visual Cognition* 14(4-8):411–443. PMID: 17387378.

- Fei Xu and Elizabeth S Spelke. 2000. Large number discrimination in 6-month-old infants. *Cognition* 74(1):B1–B11.
- Ying Yang, Qingfen Hu, Di Wu, and Shuqi Yang. 2015. Children’s and adults’ automatic processing of proportion in a Stroop-like task. *International Journal of Behavioral Development* 39(2):97–104.
- Junho Yim, Heechul Jung, ByungIn Yoo, Changkyu Choi, Dusik Park, and Junmo Kim. 2015. Rotating your face using multi-task deep neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 676–684.
- Cong Zhang, Hongsheng Li, Xiaogang Wang, and Xiaokang Yang. 2015a. Cross-scene crowd counting via deep convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 833–841.
- Jianming Zhang, Shugao Ma, Mehrnoosh Sameki, Stan Sclaroff, Margrit Betke, Zhe Lin, Xiaohui Shen, Brian Price, and Radomir Mech. 2015b. Salient object subitizing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 4045–4054.
- Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2014. Facial landmark detection by deep multi-task learning. In *European Conference on Computer Vision*. Springer, pages 94–108.