Panoptic Image Annotation with a Collaborative Assistant

Jasper R. R. Uijlings*
jrru@google.com
Google Research

Mykhaylo Andriluka*
mykhayloa@google.com
Google Research

Vittorio Ferrari
vittoferrari@google.com
Google Research

Figure 1: Example of the collaborative annotation process. Given an input image, the human annotator and the automated assistant carry out actions in turn. Based on the current annotation, the annotator performs one action (highlighted in red). Then the assistant reacts with by changing labels (in blue) and adding segments (in green). For comparison we show the ground-truth panoptic segmentation.

ABSTRACT

This paper aims to reduce the time to annotate images for panoptic segmentation, which requires annotating segmentation masks and class labels for all object instances and stuff regions. We formulate our approach as a collaborative process between an annotator and an automated assistant who take turns to jointly annotate an image using a predefined pool of segments. Actions performed by the annotator serve as a strong contextual signal. The assistant intelligently reacts to this signal by annotating other parts of the image on its own, which reduces the amount of work required by the annotator. We perform thorough experiments on the COCO panoptic dataset, both in simulation and with human annotators. These demonstrate that our approach is significantly faster than

the recent machine-assisted interface of [4], and 2.4× to 5× faster than manual polygon drawing. Finally, we show on ADE20k [62] that our method can be used to efficiently annotate new datasets, bootstrapping from a very small amount of annotated data.

CCS CONCEPTS

• Computing methodologies → Image segmentation; Scene understanding; • Human-centered computing → Graphical user interfaces.

KEYWORDS

Computer vision; image annotation; human-machine collaboration

ACM Reference Format:

Jasper R. R. Uijlings, Mykhaylo Andriluka, and Vittorio Ferrari. 2020. Panoptic Image Annotation with a Collaborative Assistant. In *Proceedings of the 28th ACM International Conference on Multimedia (MM '20)*, October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394171.3413812>

1 INTRODUCTION

This paper aims to reduce the time it takes to annotate images for the panoptic segmentation task. This requires annotating segmentation masks and class labels for all object instances and stuff regions.

*Both authors contributed equally to this research

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '20, October 12–16, 2020, Seattle, WA, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-7988-5/20/10...\$15.00
<https://doi.org/10.1145/3394171.3413812>

Such annotations are expensive: it took 19 minutes for a single image for COCO [13, 38] and 1.5 hours for Cityscapes [19]. In this paper we propose to reduce annotation time by learning to predict how the image should be annotated.

To this end we formulate our approach as a collaborative process between a human annotator and an automated assistant agent who take turns to jointly annotate an image. To see this process in action, Fig. 1a shows an example image and Fig. 1b the current annotation, which is partially machine generated. At this point, the annotator converts one of the *road* segments into *pavement*. The assistant reacts by changing the labels of similar looking segments elsewhere in the image to *pavement* as well. Fig. 1e and 1f show another example image and current annotation. After the annotator corrects the *cup* to be a *bowl*, the assistant automatically adds another *bowl*, as well as a *spoon*, a *bottle*, and a *wooden floor*. This significantly helps the annotator.

We build on [4], which proposed to annotate an image by composing segments out of a pre-defined pool, using an interface in which annotators can repeatedly perform one of the following actions: add a segment from the pool, change the label of a segment, and remove a segment. In this paper we introduce an automatic assistant which helps the annotator by executing some of these actions on its own. Crucially, whenever the annotator performs an action, this directly provides a ground-truth annotation for a segment. This ground-truth serves as a strong form of context, stronger than the use of *predicted context* in previous works [26, 28, 37, 42, 44, 50, 57]. We propose an assistant that reacts intelligently to annotator actions by capitalizing on these strong contextual cues to automatically improve the annotation of other parts of the image.

To summarize, we introduce a framework in which an assistant and an annotator collaboratively annotate an image. The assistant intelligently reacts to annotator input based on context by annotating other parts of the image on its own. We experimentally demonstrate: (1) in simulations on the COCO panoptic dataset [13, 32, 38] our method is 17% – 26% faster than [4]; (2) studies with human annotators on this dataset confirm this and enable to compare to manual polygon drawing, revealing that our method is 5× faster at a small loss of quality. Furthermore, after adding a manual refinement stage, our overall pipeline is 2.4× faster without any compromise on quality. (3) a cross-dataset experiment on ADE20k [62] shows that our method can be used to quickly annotate new datasets, bootstrapping from a very small amount of labeled data.

2 RELATED WORK

Interactive Segmentation. Many works address interactive object segmentation. Most classical approaches [5, 7, 10, 17, 20, 23, 45, 51] cast the problem as an energy minimization function defined on a graph which spans over pixels. Many recent methods adapt Fully Convolutional neural networks (FCNs, e.g. [15, 39]) for single object segmentation by using user scribbles or clicks as additional input signal [8, 28, 34–36, 40, 41, 61]. In [16] they use an FCN to produce pixel-wise embedding space. They combine this with user clicks and a nearest neighbour classifier to segment objects within video. Polygon-RNN [2, 14] is a recurrent neural net which predicts polygon vertices which an annotator can adjust. In [34] they predict

object boundaries using an FCN which accepts boundary clicks, which they turn into an instance using a geodesic path solver [18].

While the vast majority of works focuses on individual objects, a few recent works have started to tackle the full image annotation problem [3, 4]. In [3] they do class-agnostic segmentation of all objects in an image starting from extreme points [46] followed by corrective scribbles. The closest related work to ours is Fluid Annotation [4]. Instead of segmenting one object at a time, they propose an interface to quickly annotate a complete image by composing segments out of a pre-defined pool. The interface is designed to make it easy for the annotator to perform the action she chooses. In this work we build on top of [4] and go beyond by introducing an automated assistant that performs some annotation actions on its own.

Assign annotation tasks. A few works propose to intelligently choose what annotation task to send to the annotator [33, 53, 59]. In [33] they focus on creating a bounding box for each image-label pair. They have an agent decide whether to ask the annotator to draw a bounding box [46, 56] or verify whether a machine-generated bounding box is good enough [47]. In [53] the machine dispatches annotation tasks to optimize the trade-off between the annotation budget and the final labeling quality. Tasks include providing image labels, providing a label for a certain box, verify a box, draw boxes around other instances of the same class, etc. In [59] they estimate the informativeness and cost of having an image label, a box, or a full segmentation of an image, which they use to dispatch these tasks in an active learning framework. In these works the machine decides which tasks to send to the annotators while the tasks themselves are not interactive. In our case there is a single interface through which both the machine and the annotator interactively take turns.

Other works on interactive annotation. In [52] they created a semantic segmentation network which can consume natural language input. This enables a user to correct a machine-predicted segmentation by giving instructions in English. In [11] they propose a framework to convert bounding boxes, instance masks, and part annotations into each other using human verification and human corrections. Other works address fine-grained classification, where annotators provide feedback or correct the classifiers based on attributes [9, 12, 48, 60].

3 METHOD

Given an input image we want to produce a dense labelling of every pixel with a semantic label and object identity. This includes “thing” classes corresponding to various countable objects, and “stuff” classes corresponding to uncountable classes which typically occupy background areas. Example annotations are shown in Fig 1d and 1h.

We start from the recent Fluid Annotation interface [4] that allows to quickly annotate an image by composing segments out of a pre-defined pool (Sec. 3.1). In this paper we turn this into a collaborative environment (Sec. 3.2) and introduce an automated assistant which helps the annotator complete its task (Sec. 3.2). Crucially, every action of the annotator provides strong contextual cues which the assistant uses to predict how the image should be annotated. Then the assistant *carries out some actions on its own*.

Figure 2: The *proposal set*, *active set* and *fixed set*. The *proposal set* is the complete segment pool created by Mask-RCNN (few are shown for clarity). The *active set* defines the current annotation (segments 3, 5, 7). The *fixed set* contains segments which were modified/added by the annotator and are considered to be ground-truth (segment 7).

3.1 Fluid Annotation [4]

Fluid Annotation [4] starts by generating a *proposal set* of segments with accompanying class labels using Mask-RCNN [24]. They modified Mask R-CNN to generate about 1000 segments per image (more than usual) and to produce segments also on stuff classes (not only on things). The image is annotated by selecting an ordered subset of these proposals and by optionally correcting their labels. The proposals are ordered to ensure that a single pixel is assigned to only one segment, following the definition of panoptic segmentation [32].

From the proposal set, the work of [4] first creates an initial annotation for the whole image using an iterative greedy algorithm [4, 32]. Starting from the empty image, first the highest scored segment is selected. Then the next-highest scored segment is placed behind all other segments, if enough of its surface is visible. This algorithm creates an initial ordered *active set*. This active set is a subset of the proposal set, and defines the current annotation (Fig. 2). The annotator can modify the active set by performing four kinds of actions: *add a segment* from the proposal set into the active set, *remove a segment*, *change label* of an active segment, and *change depth order* of an active segment.

The Fluid Annotation interface facilitates the actions made by the annotator: when adding a segment, the mouse position is used to make a small, ordered selection of segments for the annotator to quickly scroll through. When changing a label, the interfaces makes a shortlist of likely labels for that segment. Hence in [4] the system facilitates the annotator to perform the one action she chose.

3.2 Collaborative fluid annotation

In this paper we want to extend the influence of each annotator action beyond the one segment that was targeted. We achieve this by introducing an *automatic assistant* to the Fluid Annotation system. We model annotation as a collaborative environment in which the annotator and the assistant alternate in taking turns, where both use the same set of actions. Conceptually, the annotator has perfect knowledge about the visual world and about the annotation it aims to achieve. However, it has only a partial view of the proposal set, and exploring this set is a costly process. Conversely, the assistant can access all the proposal segments instantly, but has limited capability to judge which proposals belong to the final segmentation the human annotator would like to achieve. The aim of the annotator

is to produce a high-quality panoptic annotation. The goal of the assistant is to reduce the overall annotation effort.

Crucially, the annotator conveys his knowledge about the world through every action she takes, essentially creating ground-truth as she goes. This freshly created knowledge provides strong contextual cues which the assistant uses to predict how the rest of the image should be annotated. This enables the assistant to react to the annotator and carry out some actions on its own.

Fixed Set. To establish communication between the annotator and the assistant we introduce a *fixed set*. It contains all segments which have been approved by the annotator and are considered to be ground-truth. It is a subset of the active set, which in turn is a subset of the proposal set (Fig. 2). Since the fixed set can be considered as ground-truth, we do not allow the assistant to make any changes to it.

To create the fixed set, we make four natural assumptions on the behavior of the annotator: (1) whenever the annotator changes the label of a segment, they set it to the correct label; (2) the annotator only changes the label of a geometrically correct segment (i.e. matching well a real object or background region); (3) the annotator only adds segments that are geometrically correct; (4) whenever an annotator adds a segment which has an incorrect label, they will correct it with their next action. Using these assumptions, we can now create the fixed set automatically without additional costs. By combining (1) and (2), if the annotator changes a label of a segment, we immediately put it into the fixed set. By combining (3) and (4), if an annotator adds a segment, we put it into the fixed set after their next action. Using the logs of our human experiments (Sec. 4.3), we found that these assumptions hold in the vast majority of the cases (90%-96%). Moreover, these assumptions only affect the fixed set and hence the behavior of the assistants, *not* the actual behavior of the annotator.

Workflow. The assistant and the annotator take turns to collaboratively annotate an image. In particular, the assistant performs as many actions as it wants to, until it decides stop. Then the annotator performs one action. If this action alters the fixed set, the assistant has more information which it can use during its next turn.

We introduce a collaborative assistant in Sec. 3.3. This assistant reacts to context provided by the annotator and can perform the *change label* and *add* actions. We also introduce an initialization agent in Sec. 3.4. It acts before the annotator and replaces the greedy initialization stage of Fluid Annotation (Sec. 3.1).

3.3 Collaborative Assistant (CA)

Our collaborative assistant automatically performs actions to accelerate the annotation process. To identify useful actions, it relies on a *context model* that captures dependencies between fixed segments and the segments in the proposal set. Here we describe the action generation process of the assistant, and then present details of the context model.

Action generation. Denote the fixed set as X_{fix} . For the *change label* action, our context model predicts $p(c_p^k | X_{\text{fix}})$ for each segment k . The collaborative assistant generates a *change label* action when the highest scored label of the context model is different from the current label. The action updates the label to m , using $\arg \max_m p(c_p^k = m | X_{\text{fix}})$.

Figure 3: (a) Example image with three fixed segments (green dots), and a proposal segment (white dot). (b) Diagram of a simple model with skip connections that does not have access to fixed segments. (c) Diagram of our context model.

Figure 4: Structure of the E_{fix} subnetwork: each block is a fully connected layer with ReLu activation, where the number specifies the number of channels. Features specifying human-provided inputs, geometry, and appearance are first separately processed (orange, blue, and green respectively) and fused in the final layer (light gray). All feature fusion is done through concatenation.

Similarly, for the *add* action our model predicts $p(d_p^k | X_{fix})$. The *add* action is generated whenever $p(d_p^k | X_{fix})$ is greater than threshold τ and segment k is not yet in the active set. We set $\tau = 0.9$ as estimated on a validation set (Sec. 4).

Context model. For both actions we use a context model with the same architecture. Let the set of fixed segments be denoted by $X_{fix} = \{x_{fix}^k | k = 1, \dots, K\}$ and the set of proposal segments by $X_p = \{x_p^i | i = 1, \dots, N\}$.

In our model the features of the fixed and proposal segments are given by $x_{fix}^k = [c_f^k, g_f^k, s_f^k]$ and $x_p^i = [g_p^i, s_p^i]$ respectively. Each of the variables g_f^k and g_p^i corresponds to a 4D vector encoding the center of the segment and the width and height of its bounding box. This captures spatial relationships between the proposals. The variables s_f^k and s_p^i each correspond to a vector of class scores (i.e. logits) assigned to the segment by Mask-RCNN. These logits capture semantic information and, since they are linear projection of the final feature vectors, they also contain information about the segments' appearance. Finally, the variable c_f^k denotes a one-hot encoding of the segment class that has been assigned to the fixed segment by the human annotator. Note that the c_f^k component of the fixed segment representation x_{fix}^k provides an additional and potentially strong cue for resolving ambiguities in the label of the proposal segments. This is a new piece of information not accessible to Mask-RCNN as usually deployed (i.e. without a collaboration with a human).

The structure of our *change label* context model is shown in Fig. 3c. Our starting point is a simple update model shown in Fig. 3b. This model takes the segment features x_p^i as input and computes incremental updates to the class probabilities using a fully connected neural network E_p , similarly to a single layer in ResNet [25]. We extend this simple model with a graph-convolutional [30] subnetwork that computes a fixed-size vector summarizing the relationship between the proposal and the unordered, arbitrarily sized, fixed set:

$$\bar{E}_{fix} = \frac{1}{K} \sum_k E_{fix}(x_p^i, x_{fix}^k). \quad (1)$$

The output of $\bar{E}_{fix}$ is then combined with the output of E_p and is passed through a single fully connected layer to compute a difference vector with respect to the original class probabilities of a proposal segment x_p^i .

Our context model for the *add* action is almost identical in structure. But instead of updating all class scores s_p^i it only updates the highest class score while the softmax loss is replaced by a binary cross-entropy loss (present/absent).

In Fig. 3 (d) we show the structure of the component E_{fix} that encodes relationship between a fixed segment and a proposal segment. We follow the late fusion strategy and first apply a series of transformations to each type of input features before finally combining them together. Prior to feeding geometric features g_f and g_p into the network, we transform them into a 10-dimensional vector of relative features defined similarly to [27]. Let us denote the offset vector between the locations of the fixed and the proposal segment as Δ_{fp} , and after normalization by the width and height of the fixed segment as $\hat{\Delta}_{fp}$. The vector of relative geometric features is then obtained by concatenating Δ_{fp} , $\log(|\hat{\Delta}_{fp}|)$, $\log(w_p/w_f)$, $\log(h_p/h_f)$, $\text{sign}(\Delta_{fp}) \log(|\Delta_{fp}|)$, and $\text{sign}(\Delta_{fp}) \log(|\hat{\Delta}_{fp}|)$.

Local score pooling. The model described above can operate with any type of segment scores s_p^i . We found that instead of directly using the scores provided by Mask-RCNN we achieve higher accuracy by performing a form of max-pooling over the scores of nearby segments. This corresponds to defining a new score vector $\hat{s}_p^i$ as

$$\hat{s}_{p,c}^i = \max_{j \in N(i,c)} s_{p,c}^j \quad (2)$$

where c is a class label, and $N(i, c)$ is the set of segments with label c that have intersection-over-union < 0.5 with segment i .

Training the context model. The ideal training data for our context model would be composed of training examples derived from the actions performed by real humans during image annotation. Since collection of such training data is impractical we use data collected in simulation. To that end we simulate the fluid annotation process for each of the images in the training set using the simulator described in [4] and store all segments and their labels contained in the final annotation. To construct training examples of the fixed set for our context model we then randomly sample correct segments out of the final annotation. We found such random sampling to be necessary to make the model robust with respect to the size of the fixed set encountered by the agent when deployed at annotation time. We train the context model using the Adam optimizer [29].

Discussion. A number of previous works attempts to exploit context based purely on image signals [26, 28, 37, 44, 50, 57]. Instead, our contextual cues are based on *human input*, which makes them much stronger. In terms of technical realization, our context model is related to models used for visual question answering [55] and for modeling relationships between scene objects [27].

3.4 Initialization Assistant (IA)

In the original Fluid Annotation, the initialization was done greedily through non-maximum suppression of Mask-RCNN segments based on their scores [4]. In this section we use an initialization assistant to do this instead. The initialization produces a panoptic segmentation by composing segments from the proposal set, without any human annotator involved. This method therefore can also be used for classical image segmentation prediction [15, 32, 39, 43].

We represent each segment as a feature vector consisting of: (a) an 8-bit encoding of the class label predicted by Mask-RCNN; (b) the predicted score of the segment; (c) the percentage of the segment surface that does not overlap with any segments in the current active set. The assistant uses these features to iteratively perform actions, starting from an empty active set. At each time step, it scores all the segments which could be added and select the one with the highest score. If this score is above a certain threshold, it adds the segment to the active set. Otherwise, it stops.

To train our initialization assistant, we collect examples by simulating the full fluid annotation process for each image (as in Sec. 3.3). This results in a set of positive *add* actions with their features. We obtain negative *add* actions with a form of hard negative mining. In particular, at several points during the fluid annotation simulation we ask the initialization assistant to make predictions. All predicted *add* actions which are inconsistent with the original ground truth form the negative example set. As model we use a simple 4-layer fully connected neural net. We use a quadratic hinge loss and train using the Adam optimizer [29].

Our initialization assistant is conceptually related to the search-based structured prediction algorithm [21], and to other approaches that iteratively generate structured outputs [6, 22].

4 RESULTS

In this section we first evaluate the context model (Sec. 4.1) and the full collaborative annotation process in simulation (Sec. 4.2). Then we report a study with human annotators (Sec. 4.3) and report a cross-dataset experiment (Sec. 4.4).

Figure 5: Class label accuracy of our context model (red, blue) w.r.t. the number of fixed segments given to it. The Mask-RCNN baseline (dashed black) represents the starting point without context model. Note that for $K_{fix} = 0$ the accuracy of the ensemble model is higher than the baseline due to local score pooling. (Eq. (2)).

Dataset. In most experiments we use the COCO panoptic dataset[1], which combines the original COCO dataset [38] with COCO-stuff [13], while merging some stuff classes based on [32]. It contains 118K training and 5K validation images densely labeled with 80 thing and 53 stuff classes. We divide the training set into two equal sized splits, one to train Mask-RCNN and another for our assistants (including the context models). We use the validation set for evaluation.

In Sec. 4.4 we use ADE20k to demonstrate that our method can be used to annotate a new dataset (detailed protocol in Sec. 4.4).

4.1 Context model

In this section we evaluate the performance of our context model in isolation. We measure how well it can correct labels predicted by Mask-RCNN, given a set of fixed segments X_{fix} .

Protocol. Given an image and the ground-truth segmentation, we first select the set of Mask-RCNN proposals E which best match the ground-truth segments geometrically. We do this by greedily selecting proposal segments that have maximal IoU with a ground-truth segment. We then measure the initial label accuracy provided by Mask R-CNN on these selected proposals E , and compare it to the accuracy after re-labeling them with the context model. We do this on 4500 COCO validation images (out of 5000) and measure accuracy averaged over classes.

When applying our context model, for each segment $s \in E$ we randomly select a fixed set of other segments in $X_{fix} \in E \setminus s$. We evaluate over a range of different number of fixed segments $|X_{fix}|$ and each time measure accuracy over all other segments. For this experiment we train a single *generic context model* which works for any size of the fixed set. Since this may be sub-optimal, we also train a separate context model for each specific fixed set size, and combine them into an *ensemble model*.

Results. As Fig. 5 shows, our context model helps improve label prediction accuracy, and its effect grows with the number of fixed segments given to it. For example, when given 8 fixed segments the accuracy increases from 69% for the original Mask-RCNN to 78%

Figure 6: Trade-off between effort and quality for our annotation system on the COCO panoptic dataset [13, 32, 38] (in simulation). The black line shows the baseline system [4]. The green line is our full system.

with our context model. This demonstrates that the fixed segments provide a strong context signal. We note that this effect is greater than typically demonstrated in classical context works, which do not benefit from conditioning on human input for parts of image [26, 28, 37, 42, 44, 50, 57].

We also observe that our ensemble model always performs better or equal to our generic model, where the difference is especially large for $K_{\text{fix}} \leq 3$. The ensemble model takes into account the size of K_{fix} and exploits the fact that, intuitively, the more fixed segments are given, the stronger the context signal and thus the more it can be relied upon to alter the labels of other segments.

4.2 Collaborative annotation process

We now evaluate our assistant in the full collaborative annotation environment. We simulate an annotator that tries to reproduce the original ground-truth of the COCO panoptic challenge [13, 38]. To keep the experiment fully rigorous, we do not report results on the 4500 validation images used to evaluate the context model, but on the remaining 500 images instead.

Results. Fig. 6 measures quality (panoptic score [32]) as a function of annotation effort (number of human actions). We compare the original Fluid Annotation system [4] (i.e. greedy initialization and no assistants) with our system which includes all assistants. As can be seen, our system is significantly better than [4] across the whole trade-off curve between effort and quality. In particular, to reach 0.60 panoptic quality, our system needs 4.2 actions compared to 5.7 actions of [4], a speed-up of 26%. Our system reaches the final quality of 0.71 after 14.9 annotator actions vs. 18.0 required by [4], a speed-up of 17%. From an alternative perspective, if we

Figure 7: Ablation study visualizing the number of actions required to reach a certain target quality.

fix the annotation budget to 6 annotator actions, our system yields 4% higher panoptic quality. Finally, we note that our approach is several times more efficient than the traditional annotation strategy of manually drawing polygons [19, 38, 54, 62], as we demonstrate through human studies in Sec. 4.3.

Ablation. To gain insights into the contributions of each assistant, we combine each assistant individually with the baseline [4]. We show in Fig. 7 the number of human actions required to reach a desired target quality (the fewer, the better). Note that all systems eventually reach the same final panoptic quality of 0.71. We conclude the following: (1) Each individual assistant brings visible improvements over [4] alone. (2) The assistant which performs the *change label* action is consistently better than the one performing *add segment*. Intuitively, this makes sense since changing the label of an existing segment is easier than adding an entirely new segment. (3) While the initialization assistant gives the most improvements after few human actions (see target quality = 0.6), the assistant which performs *change label* gives the most improvements when more human actions are performed (see target quality = 0.7). This makes sense since the initialization assistant does not do anything during interaction, while the *change label* assistant receives an increasing amount of context signal which it can act upon. (4) Our full system consistently improves over using each assistant individually, demonstrating that our assistants are complementary and all contribute to the final speed improvements.

We can also measure the influence of the assistants in terms of how much they reduce specific types of human actions (comparing the baseline [4] with our full system). The change label assistant

Figure 8: Two examples of the annotation process. When the annotator changes a label (highlighted in red), the assistant reacts by changing labels of existing segments (in blue) and adding new segments (in green).

reduces the number of *human* change label actions by 29%. The initialization assistant and the add segment assistant together reduce the combined number of *human* add and remove segment actions by 12% (they affect both add and remove, since wrongly added segments need to be removed).

Qualitative Results. To better understand how the assistants react to human input, Fig. 8 shows two examples of the collaborative annotation process in our full system. Notice the contextual response of the assistant: In the top row, once the annotator has annotated the small *baseball glove*, our assistant changes many regions in the background to *playingfield*. In the bottom row, after the annotator changes the label of the ground from *sand* to *snow*, our assistant corrects the label of the *rocks* in the background and also adds a *snowboard* as well as other snow regions.

Annotation quality in context. To get a better understanding of what 0.71 panoptic annotation quality means, we can put this into context by comparing to MCG [49], one of the best region proposal methods. To enable this comparison, we re-evaluate quality in terms of Intersection-over-Union averaged over ground-truth regions (mIoU) [32, 49, 58, 61]. Our 0.71 panoptic quality translates to 77.1% mIoU. In contrast, MCG [49] reports a 60% mIoU on the same dataset, when selecting the best available proposal for each ground-truth region (out of 4000). Moreover, we also compare the annotation quality of our system to manual polygon drawing through human experiments in Sec. 4.3.

4.3 Study with human annotators

We now verify whether the efficiency gains observed in simulation also transfer to human annotators.

Protocol. We randomly select 100 COCO validation images and have four humans annotate them (two experts and two non-experts). We compare polygon drawing, fluid annotation [4], and our system with all assistants. Before starting the real experiment, all subjects practice for 30+ minutes on both the system with assistants and the system without. They also practice about 10 minutes with the polygon drawing interface (which is simpler).

Importantly, we want to avoid measuring confounding effects, such as annotators being inherently faster than others, or being faster because they already saw the image before. We therefore make sure that each annotator annotates an equal number of images with each interface, while never annotating an image more than once. Additionally, we want to avoid measuring human label disagreement as a confounding factor. The influence of such disagreement is serious, as the COCO group demonstrated by annotating 5000 COCO images twice by independent annotators [31]. To rule out this factor, we show annotators the original COCO ground truth during annotation on one monitor, and ask them to reproduce it using the interface on another monitor. Thanks to all the precautions above, our study can measure time differences due to the interface only.

Results. Fig. 9 shows quality (panoptic score) as a function of human annotation time (in seconds). Overall, our method emerges

Figure 9: Results of our user study on 100 images from the COCO panoptic dataset [13, 32, 38].

as more efficient than [4] and polygon drawing. For example, after 40 seconds per image, the panoptic quality is 0.08 when doing polygon annotation, 0.48 for Fluid Annotation [4], and 0.53 for our full system. Thus, in this low-time-budget regime, our approach makes a further improvement over [4], which was already delivering much better annotations than polygon drawing. These reported improvements match what we observed in the simulation experiments (Sec. 4.2).

Polygon drawing requires 338s seconds on average to reach the end of the annotation process, compared to 67s for our method. Thus our method is 5× faster. In terms of quality, it is important to note that even when the annotator draws polygons while looking at the original ground truth, they are unable to perfectly replicate the original annotation. In practice, they reach 0.74 panoptic quality, which can be seen as an upper-bound for this task. In comparison, when using our method, annotators reach 0.59 panoptic quality.

Boundary Refinement Stage. So far, our system does not allow the annotator to refine the boundaries of the segments in the pool generated by Mask R-CNN. But if we want to reach the highest annotation quality possible, we can add a post-processing refinement stage. We did this by using a free-form drawing tool with adjustable brush size (similar to MS-paint) to refine the boundaries in the 100 images in this study. Our resulting overall pipeline took a total of 139s per image to reach a panoptic annotation quality of 0.73 on average (67s for the main stage, plus 72s for the boundary refinement stage). Hence, our overall pipeline is 2.4× faster than manual polygon drawing while achieving nearly the same panoptic quality of 0.74.

4.4 Annotating a new dataset

So far we applied our system to annotate new images from the same dataset it was trained on. In practice though, often one wants

Figure 10: Simulated annotation results on the ADE20k validation set.

to annotate an entirely new dataset, with different appearance distribution and even new object and stuff classes. Here we explore this scenario by annotating ADE20k [62] (in simulation).

We fine-tune our system on a very small amount (5%, 1k images) of the ADE20k training set, so that it still brings a strong gain in overall annotation time. We use this data to train the underlying Mask-RCNN model for generating proposal sets and our assistants. We fine-tune models starting from those pre-trained on COCO, and use a 4-fold leave-one-out strategy where the assistants are trained on different images than those used to train Mask R-CNN. This trains better assistants while making efficient use of the limited training data.

We evaluate our system on the validation set of ADE20k (2k images). As Fig. 10 shows, our results in Sec. 4.2 generalize to annotating ADE20k: at 0.6 panoptic score, our system is 19% faster than [4]. Furthermore, our system reaches the final 0.63 panoptic quality 12% faster than [4].

5 CONCLUSIONS

We introduced a framework in which a human annotator and an automated assistant collaboratively annotate an image. The assistant intelligently reacts to human input based on context to annotate other parts of the image by itself. On the COCO panoptic dataset [13, 32, 38]: (1) in simulation we demonstrated that our system is 17%-26% faster than the recent interface of [4]; (2) a human experiment confirms these improvements. At the same time, it shows that our method is 5× faster than polygon drawing at a small loss of annotation quality. Furthermore, after adding a manual refinement stage, our system is 2.4× faster without any compromise on quality. In a cross-dataset experiment on ADE20k [62], (3) we show that starting from a very small amount of annotated data, our approach can be used to quickly annotate new datasets.

REFERENCES

- [1] [n.d.]. COCO panoptic dataset. <http://cocodataset.org/index.htm#panoptic-2018>.
- [2] D. Acuna, H. Ling, A. Kar, and S. Fidler. 2018. Efficient Interactive Annotation of Segmentation Datasets with Polygon-RNN++. In *CVPR*.
- [3] Eirikur Agustsson, Jasper RR Uijlings, and Vittorio Ferrari. 2019. Interactive Full Image Segmentation by Considering All Regions Jointly. In *CVPR*.
- [4] Mykhaylo Andriluka, Jasper R. R. Uijlings, and Vittorio Ferrari. 2018. Fluid Annotation: A Human-Machine Collaboration Interface for Full Image Annotation. In *ACM Multimedia*.
- [5] Xue Bai and Guillermo Sapiro. 2009. Geodesic matting: A framework for fast interactive image and video segmentation and matting. *IJCV* (2009).
- [6] Dan Banica and Cristian Sminchisescu. 2015. Second-Order Constrained Parametric Proposals and Sequential Search-Based Structured Prediction for Semantic Segmentation in RGB-D Images. (2015).
- [7] D. Batra, A. Kowdle, D. Parikh, J. Luo, and T. Chen. 2011. Interactively Co-segmenting Topically Related Images with Intelligent Scribble Guidance. *IJCV* (2011).
- [8] Rodrigo Benenson, Stefan Popov, and Vittorio Ferrari. 2019. Large-scale interactive object segmentation with human annotators. In *CVPR*.
- [9] Arijit Biswas and Devi Parikh. 2013. Simultaneous active learning of classifiers & attributes via relative feedback. In *CVPR*.
- [10] Y. Boykov and M. P. Jolly. 2001. Interactive Graph Cuts for Optimal Boundary and Region Segmentation of Objects in N-D Images. In *ICCV*.
- [11] S. Branson, K.E. Hjørleifsson, and P. Perona. 2014. Active annotation translation. In *CVPR*.
- [12] Steve Branson, Catherine Wah, Florian Schroff, Boris Babenko, Peter Welinder, Pietro Perona, and Serge Belongie. 2010. Visual recognition with humans in the loop. In *ECCV*.
- [13] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. 2018. COCO-Stuff: Thing and Stuff Classes in Context. In *CVPR*.
- [14] L. Castrejon, K. Kundu, R. Urtasun, and S. Fidler. 2017. Annotating Object Instances with a Polygon-RNN. In *CVPR*.
- [15] L-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A.L. Yuille. 2018. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Trans. on PAMI* (2018).
- [16] Yuhua Chen, Jordi Pont-Tuset, Alberto Montes, and Luc Van Gool. 2018. Blazingly Fast Video Object Segmentation with Pixel-Wise Metric Learning. In *CVPR*.
- [17] Ming-Ming Cheng, V A Prisacariu, Shuai Zheng, Philip H. S. Torr, and Carsten Rother. 2015. DenseCut: Densely Connected CRFs for Realtime GrabCut. *Computer Graphics Forum* (2015).
- [18] L. Cohen and R. Kimmel. 1996. Global Minimum for Active Contour Models: A Minimal Path Approach. In *CVPR*.
- [19] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. 2016. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *CVPR*.
- [20] A. Criminisi, T. Sharp, C. Rother, and P. Perez. 2010. Geodesic Image and Video Editing.
- [21] Hal Daumé III, John Langford, and Daniel Marcu. 2009. Search-based Structured Prediction. *Machine Learning* 75, 3 (2009), 297–325.
- [22] Georgia Gkioxari, Alexander Toshev, and Navdeep Jaitly. 2016. Chained predictions using convolutional neural networks. (2016).
- [23] V. Gulshan, C. Rother, A. Criminisi, A. Blake, and A. Zisserman. 2010. Geodesic star convexity for interactive image segmentation. In *CVPR*.
- [24] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask R-CNN. In *ICCV*.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*.
- [26] G. Heitz and D. Koller. 2008. Learning Spatial Context: Using Stuff to Find Things. In *ECCV*.
- [27] Ronghang Hu, Piotr Dollár, Kaiming He, Trevor Darrell, and Ross Girshick. 2018. Learning to Segment Every Thing. In *CVPR*.
- [28] Yang Hu, Andrea Soltoggio, Russell Lock, and Steve Carter. 2019. A fully convolutional two-stream fusion network for interactive image segmentation. *Neural Networks* (2019).
- [29] Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [30] Thomas N Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. (2017).
- [31] A. Kirillov. [n.d.]. Panoptic Challenge Intro. COCO+Mapillary Joining Recognition Challenge Workshop. <http://presentations.cocodataset.org/ECCV18/COCO18-Panoptic-Overview.pdf>.
- [32] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollar. 2018. Panoptic Segmentation. *CVPR* 1801.00868.
- [33] K. Konyushkova, J.R.R. Uijlings, C. Lampert, and V. Ferrari. 2018. Learning Intelligent Dialogs for Bounding Box Annotation. In *CVPR*.
- [34] Hoang Le, Long Mai, Brian Price, Scott Cohen, Hailin Jin, and Feng Liu. 2018. Interactive Boundary Prediction for Object Selection. In *ECCV*.
- [35] Z. Li, Q. Chen, and V. Koltun. 2018. Interactive Image Segmentation with Latent Diversity. In *CVPR*.
- [36] J.H. Liew, Y. Wei, W. Xiong, S-H. Ong, and J. Feng. 2017. Regional interactive image segmentation networks. In *ICCV*.
- [37] D. Lin, Y. Ji, D. Lischinski, D. Cohen, and H. Huang. 2019. Multi-Scale Context Intertwining for Semantic Segmentation. In *ECCV*.
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2014. Microsoft COCO: Common Objects in Context. In *ECCV*.
- [39] J. Long, E. Shelhamer, and T. Darrell. 2015. Fully Convolutional Networks for Semantic Segmentation. In *CVPR*.
- [40] S. Mahadevan, P. Voigtlaender, and B. Leibe. 2018. Iteratively Trained Interactive Segmentation. In *BMVC*.
- [41] K.-K. Maninis, S. Caelles, J. Pont-Tuset, and L. Van Gool. 2018. Deep Extreme Cut: From Extreme Points to Object Segmentation. In *CVPR*.
- [42] Davide Modolo, Alexander Vezhnevets, and Vittorio Ferrari. 2015. Context Forest for Object Class Detection. In *BMVC*.
- [43] R. Mottaghi, X. Chen, X. Liu, N.-G. Cho, S.-W. Lee, S. Fidler, R. Urtasun, and A. Yuille. 2014. The role of context for object detection and semantic segmentation in the wild. In *CVPR*.
- [44] K. Murphy, A. Torralba, and W. T. Freeman. 2003. Using the Forest to See the Trees: A Graphical Model Relating Features, Objects, and Scenes. In *NIPS*.
- [45] N. S. Nagaraja, F. R. Schmidt, and T. Brox. 2015. Video Segmentation with Just a Few Strokes. In *ICCV*.
- [46] Dim P Papadopoulos, Jasper RR Uijlings, Frank Keller, and Vittorio Ferrari. 2017. Extreme clicking for efficient object annotation. In *ICCV*.
- [47] D. P. Papadopoulos, Jasper R. R. Uijlings, F. Keller, and V. Ferrari. 2016. We don't need no bounding-boxes: Training object class detectors using only human verification. In *CVPR*.
- [48] Amar Parkash and Devi Parikh. 2012. Attributes for classifier feedback. In *ECCV*.
- [49] J. Pont-Tuset, P. Arbeláez, J. Barron, F. Marques, and J. Malik. 2017. Multiscale Combinatorial Grouping for Image Segmentation and Object Proposal Generation. *IEEE Trans. on PAMI* 39, 1 (2017), 128 – 140.
- [50] A. Rabinovich, A. Vedaldi, C. Galleguillos, E. Wiewiora, and S. Belongie. 2007. Objects in Context. In *ICCV*.
- [51] C. Rother, V. Kolmogorov, and A. Blake. 2004. GrabCut: Interactive Foreground Extraction using Iterated Graph Cuts. In *SIGGRAPH*.
- [52] Christian Ruppert, Iro Laina, Nassir Navab, Gregory D. Hager, and Federico Tombari. 2018. Guide Me: Interacting with Deep Networks. In *CVPR*.
- [53] Olga Russakovsky, Li-Jia Li, and Li Fei-Fei. 2015. Best of both worlds: human-machine collaboration for object annotation. In *CVPR*.
- [54] B. C. Russell, A. Torralba, K. P. Murphy, and W. T. Freeman. 2008. LabelMe: A database and web-based tool for image annotation. *IJCV* 77, 1-3 (2008), 157–173.
- [55] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. In *NIPS*.
- [56] H. Su, J. Deng, and L. Fei-Fei. 2012. Crowdsourcing annotations for visual object detection. In *AAAI Human Computation Workshop*.
- [57] J. Tighe and S. Lazebnik. 2011. Understanding Scenes on Many Levels. In *ICCV*.
- [58] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. 2013. Selective search for object recognition. *IJCV* (2013).
- [59] Sudheendra Vijayanarasimhan and Kristen Grauman. 2009. What's it going to cost you?: Predicting effort vs. informativeness for multi-label image annotations. In *CVPR*.
- [60] Catherine Wah, Grant Van Horn, Steve Branson, Subhrajyoti Maji, Pietro Perona, and Serge Belongie. 2014. Similarity comparisons for interactive fine-grained categorization. In *CVPR*.
- [61] N. Xu, B. Price, S. Cohen, J. Yang, and T.S. Huang. 2016. Deep interactive object selection. In *CVPR*.
- [62] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba. 2019. Semantic Understanding of Scenes through the ADE20K Dataset. *IJCV* 127, 3 (2019), 302–321.