

Long-Term Image Boundary Prediction

Apratim Bhattacharyya, Mateusz Malinowski, Bernt Schiele, Mario Fritz

Max Planck Institute for Informatics
Saarland Informatics Campus, Saarbrücken, Germany
{abhattach, mmalinow, schiele, mfritz}@mpi-inf.mpg.de

Abstract

Boundary estimation in images and videos has been a very active topic of research, and organizing visual information into boundaries and segments is believed to be a corner stone of visual perception. While prior work has focused on estimating boundaries for observed frames, our work aims at predicting boundaries of future unobserved frames. This requires our model to learn about the fate of boundaries and corresponding motion patterns – including a notion of “intuitive physics”. We experiment on natural video sequences along with synthetic sequences with deterministic physics-based and agent-based motions. While not being our primary goal, we also show that fusion of RGB and boundary prediction leads to improved RGB predictions.

Introduction

Humans possess the skill to imagine future states of observed scenes in diverse scenarios. This supports various different tasks ranging from planning to object manipulation, e.g. a goalkeeper jumping to intercept the ball or reaching out for a handshake. Humans can readily perform such complex and versatile tasks because they can anticipate motions including an intuitive understanding of physical laws from the early age (Baillargeon 1994; 2004).

In this work, we propose the task of predicting future scene boundaries. Scene boundaries capture the important structure and extents of objects. Moreover, they can be accurately estimated (Khoreva et al. 2016). Prediction of future scene boundaries requires understanding of object dynamics and motion patterns including an intuitive understanding of physical laws or “intuitive physics”. In this work, we focus on two particular scenarios involving motion and local interactions. The first one, which we call physics-based motion, can fully be described by the laws of physics, e.g. dynamics of billiard balls. The second one, which we call agent-based motion, also involves understanding of intentions, e.g. dynamics of an ice-skater. Therefore, our methods have to deal with diverse situations, work on raw pixels, and should be capable of long-term predictions. Figure 1 shows example results of our method that accurately predicts future scene boundaries.

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

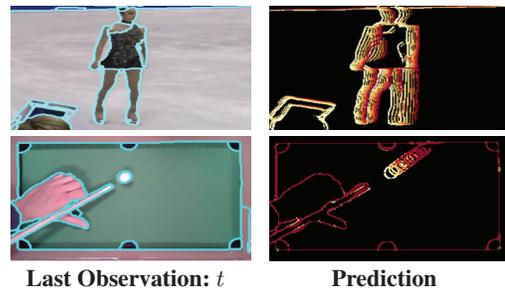


Figure 1: Predicted future boundary images, from $t + 1$ (Yellow) to $t + 8$ (Row 1), $t + 18$ (Row 2) (Red), superimposed.

Recently, full future frame prediction of observed scenes has been studied (Mathieu, Couprie, and LeCun 2016; Liu et al. 2017). But up to now, only very short range predictions of few frames have been shown, where blurriness/distortion artifacts occur in the predicted future frames – losing/incorrectly propagating high-frequency information. This high frequency information is crucial for meaningful predictions about the future, e.g. on a billiard table the location of a ball and table boundaries are necessary to infer the future state of the table. Boundaries capture this crucial high frequency information and are also known to reveal important structures of the visual scene (Wertheimer 1923; Arbelaez et al. 2011; Galasso et al. 2013). Therefore, we argue that the task of future boundary prediction is a more suitable benchmark for understanding and predicting physics or agent-based motion.

Our main contributions are as follows, 1. We propose the novel task of future boundary prediction. 2. We propose the first method that predicts future boundaries based only on the raw pixels. 3. We evaluate our model on two scenarios involving physics-based (synthetic and real billiard sequences) and agent-based motion (VSB100, (Galasso et al. 2013)). 4. Under the physics-based scenario, the method shows for the first time long-term predictions. 5. Under the agent-based scenario on VSB100 and UCF101, we show that the predicted boundaries can be used in a fusion scheme that improves RGB video prediction in the longer-term.

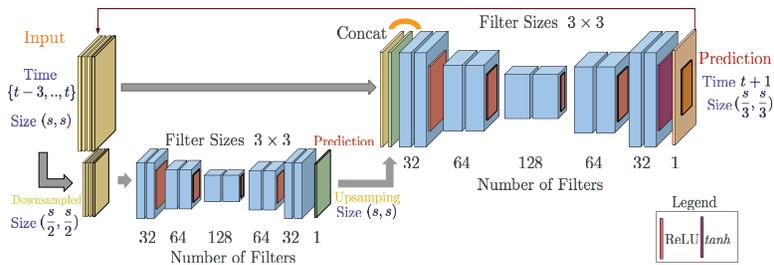


Figure 2: Convolutional Multi-scale with Context architecture (only 2 out of 4 scales illustrated).

Related work

RGB frame prediction. This problem has recently received a lot of attention. However, the predicted frames have blurriness problems. (Ranzato et al. 2014) sought to remedy this problem by discretizing the input through k-means atoms and predicting on this vocabulary instead. The work of (Mathieu, Couprie, and LeCun 2016) proposes using adversarial loss, which leads to improved results over (Ranzato et al. 2014). (Liang et al. 2017; Liu et al. 2017; Patraucean, Handa, and Cipolla 2015) shows some further improvement through the use of optical flow information. However, these approaches produce sharper short term predictions but still suffer from blurriness problems starting as soon as 3 frames into the future. (Kalchbrenner et al. 2017) focus on moving MNIST digits and like (Finn, Goodfellow, and Levine 2016) on action conditioned video prediction. (Villegas et al. 2017) proposes a hierarchical approach for making long-term frame predictions, by first estimating the high-level structure in the input frames and predicting how that structure evolves in the future. They show promising results on videos where pose is an easily identifiable and appropriate high level structure to exploit. However, such high-level structures are video domain dependent. Other works (Sutskever, Hinton, and Taylor 2009; Michalski, Memisevic, and Konda 2014) focus on deterministic bouncing ball sequences, but their dataset is limited in size and resolution and generalization with respect to the number of balls and their velocities is not considered.

Intuitive physics. Developing an intuitive understanding of physics from raw visual input has been explored recently. (Fragkiadaki et al. 2016) predict future states of balls moving on a billiard table and (Lerer, Gross, and Fergus 2016; Li, Leonardis, and Fritz 2017) predict the stability of towers made out of blocks. However, both (Fragkiadaki et al. 2016) and (Lerer, Gross, and Fergus 2016) have an “object notion”, meaning that the architecture knows a priori the location or type of the objects that it is supposed to infer. Although some recent approaches such as (Battaglia et al. 2016; Watters et al. 2017) are capable of long-term predictions, they are modeling either state-to-state or images-to-state transitions. Moreover, in the latter case, the input is visually simplified, and the focus is only on deterministic motions. In contrast to this body of work, we focus on more diverse scenarios and are agnostic to the underlying objects and causes of change.

Video segmentation. Video segmentation as the task of finding consistent spatio-temporal boundaries in a video volume has received significant attention over the last years (Galasso et al. 2014; Ochs, Malik, and Brox 2014; Galasso et al. 2013; Chang, Wei, and Fisher 2013), as it provides an initial analysis and abstraction for further processing. In contrast, our approach aims at predicting these boundaries into the future without any video observed for future frames.

Model

We present a model that observes a sequence of boundary images, where each pixel encodes the confidence of occurrence of an image boundary at that location and then predicts the boundary image(s) at the next time-step(s). An overview of our Convolutional Multi-Scale Context (CMSC) model is shown in Figure 2.

We approach long term prediction by recursion, due to the advantage of efficiency. However, errors are potentially propagated and accumulated over time. To mitigate such effects, we need our model to be accurate and to consolidate information over time. To maximize accuracy, our model has been designed through analysis of prior work on the related task of frame prediction. Furthermore, our model has many novel aspects which are key to long term prediction.

In order to generalize across diverse sequences while maintaining a tractable number of parameters, a patch based approach is adopted. Therefore, our model observes and predict on patches rather than the complete input image. Alternatively, this can be seen as multiple replicas (“patch predictors”) of our model predicting on patches of the input sequence. We now describe our model through an analysis of its various components.

Fully Convolutional.

Our CMSC model consists of only convolutional layers. The input boundary image sequence is concatenated as channels and is read by the first convolutional layer. Convolutional layers can extract high quality location invariant features. In particular, they can extract information about the orientation and direction of motion of boundaries. Neurons at upper convolutional layers have larger receptive fields and can aggregate information. In fact, as shown by the work (Jain et al. 2007), the output layer should have a wide receptive field to preserve long range spatial and temporal dependencies and learn about interaction among boundaries in a

spatio-temporal context. We therefore use several convolutional layers in our CMSC model. We also introduce pooling in between convolutional layers. Pooling further helps in the aggregation of information and increases receptive fields. However, excessive pooling (or tight bottlenecks with fully connected layers) have been shown to be successful in classification tasks, but also have shown by (Ranzato et al. 2014) to induce image degradations for synthesis tasks. Therefore, it is crucial to use moderate pooling. Finally, we use up-sampling layers after pooling to maintain resolution.

Multiple Scale Prediction.

Multi-Scale model architectures akin to a Laplacian pyramid have shown to be advantageous for generating natural images (Denton et al. 2015) and predicting future RGB frames (Mathieu, Couprie, and LeCun 2016). Such model architectures contain multiple levels which observe the input boundary image(s) at increasing (coarse to fine) scales. Down-sampling a boundary image has the effect of smoothing and discarding details of a boundary image. It is easier to predict future boundary images at a coarser resolution. Therefore, our CMSC model uses multiple scales (or levels). The input $I(L_{2k})$ to a certain level (L_{2k}) is the input boundary image sequence scaled to the current level X_{2k} and the boundary image O predicted by the previous coarser level (L_k). The boundary image predicted by the coarser level is upsampled \hat{O} to the scale at the current level. We have,

$$I(L_{2k}) = \{X_{2k}, \hat{O}(L_k)\}$$

The coarse predicted boundary images $\hat{O}(L_k)$ act as a guide for the next higher level of the model. We use four levels, with scales increasing by a factor of two.

Details of each Level in our Model. Each level of the model consists of five sets of two convolutional layers. There are 32, 64, 128, 64 and 32 filters respectively in each set, of a constant size 3×3 . Multiple convolutional layers at each scale lead to large receptive fields at the output layer. We introduce moderate 2×2 pooling layer after the first two sets of convolutional layers, leading to further aggregation of information and increased receptive fields. We double the number of convolutional filters after pooling to aid feature extraction. We upsample the convolutional maps after the third set to maintain resolution. We use *ReLU* non-linearities between every layer except the last. We use the *tanh* non-linearity at the end to ensure output in the range $[0, 1]$. (Additional details in the Appendix)

For accurate long term prediction, it is crucial to ensure global consistency through communication between the patch predictors. Consider a video of a moving ball. The trajectory of a ball might intersect with multiple patches. To correctly predict the motion far into the future, replicas of the model predicting on neighboring patches need to be consistent especially during transition of the ball between patches. Therefore, we describe next the final component of our CMSC model, the context, which ensures global consistency.

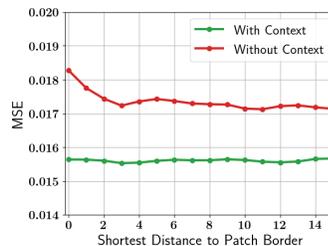


Figure 3: Our model without context has higher error near the patch boundary (red) vs. with context (green).

Context.

Our CMSC model observes a central patch along with the directly neighbouring 8 patches. This neighbourhood is called the context. However, the model only predicts on the central patch. While predicting recursively, the model observes its previous output along with the the output of the neighboring patch predictors. This enables the learning of spatially consistent predictions while keeping the same number of parameters.

The addition of a context has the added advantage that the output layer neurons now have receptive fields that are uniform in size. Without context, the neurons at the boundary of the (2D) output layer have a smaller receptive field compared to the neurons at the center. This leads to a non-uniform (training and test) error distribution at the output layer neurons. In Figure 3 we plot the average error at the output layer neurons of our CMSC model at increasing distance from the patch border, with and without context. Error increases consistently from patch center (right) to the patch border (left) without a context. Note that, the model of (Mathieu, Couprie, and LeCun 2016) is also multi-scale and fully convolutional like CMSC, but it does not have pooling or context.

Next, we evaluate our CMSC model and the effectiveness of its various components.

Experiments

We evaluate our CMSC model on natural video sequences involving agent-based motion and billiard sequences with only physics-based motion. We compare with various baselines and perform ablation studies to confirm design choices. We convert each video into 32×32 pixel patches. The CMSC model observes a central patch and eight neighbouring patches resulting in a context of size 96×96 pixels.

Training Loss. We use L2 loss (mean square error) during training, which we optimize using the ADAM optimizer.

Evaluation Metric. As we want sharp and accurate boundaries, we use the established boundary precision recall (BPR) evaluation metric from the video segmentation literature (Galasso et al. 2013). This metric is defined for a set P of predicted boundary images and G of corresponding

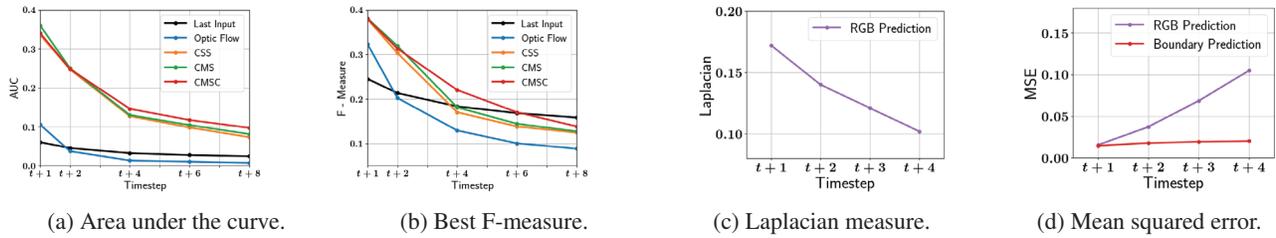


Figure 4: Left (a) and (b): Evaluation of boundary prediction on VSB100. Right (c) and (d): RGB versus boundary prediction.

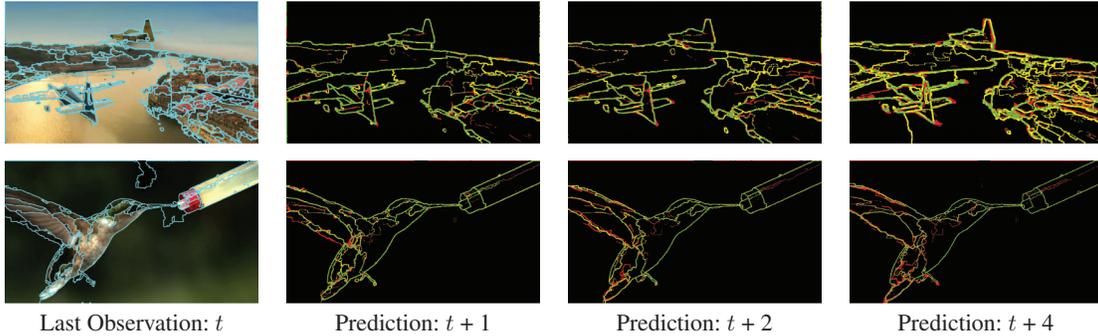


Figure 5: Rows top to bottom: Prediction on *airplane* and *hummingbird* sequences from VSB100. Correct boundaries predictions are encoded in green. Missed boundaries are encoded in yellow. Wrong boundaries are encoded in red.

ground truth boundary images as,

$$P = \frac{\sum_{B_p \in P, B_g \in G} |B_p \cap B_g|}{\sum_{B_p \in P} |B_p|}$$

$$R = \frac{\sum_{B_p \in P, B_g \in G} |B_p \cap B_g|}{\sum_{B_g \in G} |B_g|}$$

$$F = \frac{2PR}{P + R},$$

where P is boundary precision, R is boundary recall and F is the combined F-measure. As we are interested in accurate predictions, predicted boundary pixels should be at most 1 pixel away from ground-truth boundary pixels to be correct.

Evaluation on Natural Video Sequences Involving Agent-based Motion.

Dataset and Training. We use the VSB100 dataset which contains 101 videos with a maximum 121 frames each. The training set consists of 40 videos and the test set consists of 60 videos. The videos contain a wide range of objects of different sizes and shapes, including vehicles, humans and animals. The videos also have a wide variety of both object and camera motion. We use the hierarchical video segmentation algorithm in (Khoreva et al. 2016) to segment these videos. The output is an ultra-metric contour map (ucm). Boundaries higher in the hierarchy typically correspond to semantically coherent entities like animals, vehicles etc and therefore their motion corresponds to object/camera motion. We discard boundaries belonging to the lowest level of the hierarchy (corresponding to an over-segmentation), as they

are temporally very unstable. We use the ucm hierarchy as a confidence measure on boundary location at a pixel.

Experimental Settings and Baselines. The models are trained to predict boundaries of segmented VSB100 videos. Recall that, the ground-truth boundaries (ucm) in VSB100 have different confidence values. Thus, we threshold the predictions before comparison to the groundtruth. We vary the threshold to obtain a precision-recall curve and report the area under the curve (AUC) along with the best F-measure across all thresholds. We include a ‘‘Last Input’’ baseline by using the last input frame as constant prediction and a ‘‘Optical flow’’ baseline. As many boundaries do not change between frames in the videos of VSB100, the last input is a strong baseline especially when we are predicting one step into the future. In case of the optic flow baseline, the optic flow is calculated between the last two input frames (at $t - 1$ and t) using the Epic flow method of (Revaud et al. 2015). The boundary pixels at time t are propagated using the calculated flow to generate predictions at $t + 1$ to $t + 8$.

Results on VSB100. We perform an ablation study of our CMSC model and we compare to, 1. A convolutional single scale model (CSS) 2. A convolutional multi-scale model (CMS), in addition to the baselines. Both models do not have a context. We report the quantitative results in Figure 4a and Figure 4b and the qualitative results in Figure 5.

Quantitative evaluation: In the short term the CMS model (green lines) performs well. However, our CMSC (red lines) performs best in the longer term (both having the same number of parameters). This demonstrates the importance of the context for long-term prediction. The good performance of both of the multi-scale models (CMS and CMSC) versus the

Step	Last Input	CMS	CMSC-BL	CMSC
$t + 1$	0.141	0.282	0.957	0.987
$t + 5$	0.038	0.101	0.841	0.900
$t + 20$	0.002	0.066	0.347	0.632

Table 1: Evaluation on single ball billiard table worlds.

single scale CSS model, shows that multiple scales lead to more accurate predictions. The performance advantage of our CMSC model over the last input baseline shows that the model learns to predict boundaries of moving objects while keeping static boundaries intact. The recall of the CMSC model declines with time as the future becomes increasingly uncertain. The poor performance of the ‘‘Optic flow’’ baseline is due to inaccurate flow information at object boundaries.

Qualitative evaluation: The boundaries produced by our CMSC model are sharp whenever the motion is smooth, e.g. the predictions in Figure 5. However, the models are not able to deal with high uncertainty in the long-term often due to non-deterministic motion. The models in such situations react by blurring the boundaries, as a consequence of using the L2 training loss. While predicting recursively, this leads to loss of boundary confidence and eventual vanishing boundaries. The ‘‘Optic flow’’ baseline produces discontinuous (jagged) boundaries. (See Appendix for more examples). Next we evaluate and compare RGB prediction to boundary prediction.

RGB versus Boundary Prediction. We report the sharpness of RGB frames (of VSB100) predicted by the adversarial model of (Mathieu, Couprie, and LeCun 2016) (fine-tuned on VSB100) using the Laplacian measure (Krotkov 2012) in Figure 4c. The Laplacian measure pools the gradient information of the image. We observe that the model of (Mathieu, Couprie, and LeCun 2016) makes increasingly blurry predictions into the future. We also compare the mean squared error of RGB predictions of (Mathieu, Couprie, and LeCun 2016) and predicted boundaries of our CMSC model in Figure 4d. We see a sharper increase in the error of RGB predictions compared to boundaries in the long term.

Evaluation on Physics-based Motion.

Motion in the videos in the VSB100 dataset is frequently very complex as agent’s actions quickly become non-deterministic and hence increasingly uncertain. Therefore, we also look at physics-based motion, which is still challenging yet it factors out the aforementioned issues. In this scenario, we evaluate the long-term prediction performance of the models on real and synthetic billiard ball sequences. We begin by describing our dataset.

Synthetic Data Generation. The synthetic billiard ball sequences are sampled from worlds which consists of balls moving on a frictionless surface with a border, akin to a billiard table. We used pygame to create such worlds and sample boundary images from them. The output images contain boundaries that can stem from ball(s) or the

table and have binary confidence measure (indicating a boundary at that location). During evaluation, as the target is always a binary image, we report only the best F-measure obtained by thresholding the predicting boundary images and varying the threshold parameter. We sampled synthetic billiard sequences using the following parameters. 1. *Table size:* Side length randomly sampled from $\{96,128,160,192,256\}$ pixels. 2. *Ball velocity:* Randomly sampled from $[\{-3,\dots,3\},\{-3,\dots,3\}]$ pixels. 3. *Ball size:* Constant, with a radius of 13 pixels. 4. *Initial Position:* Uniformly over the table surface.

Real Data Collection. We captured a novel data-set of real billiard sequences on a mini-billiard table. Frame rate was set to 120 per second to minimize motion blur. Each sequence consists of an actor (not visible) striking the ball with a cue stick once. The motion in the sequences of the dataset are that of the cue stick and the balls. We produce boundary images using the method of (Maninis et al. 2016).

Evaluation on synthetic single ball worlds. We generate a training set using parameters mentioned previously. However, to keep our training set as diverse as possible we prefer short sequences. We restrict each sequence to a maximum length of one or two collisions with walls and set a 50% bias of the initial position of the balls being 40 pixels from the walls. We sample 500 such sequences and train our models on these sequences. We then test the models on 30 independent test sequences. We again include the ‘‘last input’’ baseline as a constant predictor. We also include a ‘‘blind’’ Convolutional multi-scale Context model (CMSC-BL), which cannot see the table borders. This is a strong baseline as starting from 42% frames in the test set, there are no ball-wall collisions 20 steps into the future. To beat this baseline, our models need to learn the physics of ball-wall collisions. We report the results in Table 1.

Our CMSC model performs the best with accurate predictions 20 time-steps into the future – also exceeding the ‘‘blind’’ version (CMSC-BL) that cannot handle ball-wall collisions. The model without a context CMS, produces inaccurate results at patch borders and thus suffers heavily especially at larger time-steps.

Evaluation on synthetic two and three ball worlds.

Worlds with more than one ball also involve harder to model physics of ball-ball collisions. To evaluate the models on such worlds we sample 100 training sequences each with two, three and six balls respectively with a maximum length of 200 frames. We use a curriculum learning approach (Bengio et al. 2009), where we initialize the models with the weights learned on single, two and three ball worlds respectively. We test the models on 30 independent sequences containing two, three and six balls respectively. We report the results in Table 2. In each case, we also include CMSC models trained on single ball worlds (CMSC-1B), two ball worlds (CMSC-2B) and three ball worlds (CMSC-3B) respectively as baselines. To beat these strong baselines learning the physics of ball-ball collisions is necessary as in case of our two-ball and three-ball test sets, there are no ball-ball and 3-ball collisions 20 steps into the future for 92% and 98% of the starting frames (and no 6-ball collisions). Again,

Step	– Evaluation on two ball worlds –			– Evaluation on three ball worlds –			– Evaluation on six ball worlds –		
	Last Input	CMSC-1B	CMSC	Last Input	CMSC-2B	CMSC	Last Input	CMSC-3B	CMSC
$t + 1$	0.246	0.966	0.969	0.246	0.967	0.968	0.250	0.962	0.964
$t + 5$	0.114	0.848	0.896	0.118	0.890	0.892	0.130	0.875	0.866
$t + 20$	0.101	0.612	0.681	0.090	0.664	0.700	0.115	0.511	0.600

Table 2: Evaluation on complex billiard table worlds.



Figure 6: Trails produced by super-imposing predicted boundaries on synthetic sequences.

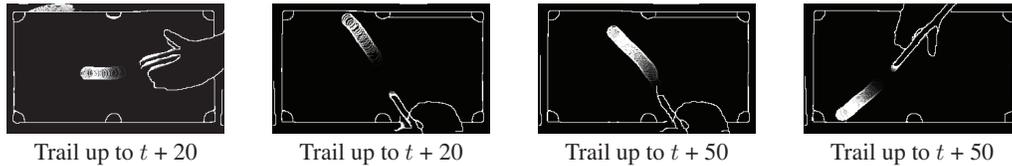


Figure 7: Trails produced by super-imposing predicted boundaries on real sequences.

Step	Last Input	CMSC	Last Input(M)	CMSC(M)
$t + 1$	0.890	0.850	0.126	0.570
$t + 5$	0.855	0.804	0.085	0.541
$t + 20$	0.844	0.746	0.087	0.497

Table 3: Evaluation on real billiard sequences (M-masked).

we see accurate prediction by the CMSC model even at 20 time-steps in the future.

Prediction over Very Long Time Scales. Although we evaluate only 20 timesteps into the future in Table 1 and Table 2, our models are stable over longer time-horizons. In Figure 6, we predict 100 timesteps and visualize the boundary images by trails obtained by superposition. We notice a few failure cases where a ball reverse direction mid table and the ball(s) get deformed or disappear.

Evaluation on Real Billiard Sequences. Prediction on real billiard table sequences is a challenging test for our models. The table fabric causes rapid deceleration of the ball (compared to the constant velocity in the synthetic sequences). Spin is sometimes inadvertently introduced and a segmentation algorithm applied on the observed frames introduces artifacts. The boundaries are not always consistent across frames of a sequence and they are jagged and change shape. We collect 350 real billiard sequences, with one ball, as our training set. To deal with deceleration, we experiment with increasing the number of input frames. We train our CMSC model with six input frames and pre-train on our synthetic one ball training set. We report the results of evaluation (F-measure as before) on 30 independent sequences in Table 3. As many boundaries (e.g table borders) remain static the last input baseline performs very well. For

fair comparison we use a mask obtained with a ball tracker, Our method is able to propagate the motion of the ball and beats the last input baseline in the masked case. We show qualitative results in Table 7 as trails, where our model predicts 20 and 50 time-steps into the future.

Sharpening RGB Predictions with Fusion

The sharp boundaries produced by our models raise the prospect of sharpening RGB predictions in a fusion scheme. We present our fusion architecture in Figure 9, which fuses RGB predictions of (Mathieu, Couprie, and LeCun 2016) with our boundaries. Note that, our approach can be used on top of any RGB frame prediction method and unlike (Villegas et al. 2017) is video domain agnostic. It is inspired by prior work (Eigen, Krishnan, and Fergus 2013; Mao, Shen, and Yang 2016) on deblurring/denoising. Like these models our fusion model is fully convolutional. Resolution is maintained by skip connections, as in (Mao, Shen, and Yang 2016). Our fusion model takes as input the predicted RGB and boundaries at each timestep and is trained with L2 loss.

Datasets and metrics. We evaluate on both VSB100 and UCF101 datasets. We randomly select 30 and 20 videos from VSB100 to train our CMSC model and our fusion model. We test on the remaining 50 videos. Similarly we randomly select 1000, 500 (training) and 1000 (test) videos from UCF101. The UCF101 train/test set was segmented using the method of (Maninis et al. 2016). We use PSNR, the sharpness loss measure from (Mathieu, Couprie, and LeCun 2016) and the Laplacian measure as evaluation metrics.

Baselines. We include a baseline de-blurring model. It has the same architecture as our fusion model, except for the

Step	PSNR			Sharpness Loss			Laplacian Measure		
	RGB prediction	De-blurring	Fusion (Ours)	RGB prediction	De-blurring	Fusion (Ours)	RGB prediction	De-blurring	Fusion (Ours)
VS100									
$t + 2$	24.4	24.5	25.1	18.5	18.5	18.6	0.142	0.139	0.155
$t + 3$	22.2	22.9	23.1	18.2	18.2	18.3	0.121	0.109	0.127
$t + 4$	20.4	21.7	22.3	18.1	18.1	18.2	0.103	0.114	0.118
UCF101									
$t + 2$	26.5	27.7	28.2	21.4	21.5	21.7	0.101	0.122	0.136
$t + 3$	23.4	25.1	25.2	20.5	20.8	20.9	0.095	0.093	0.102
$t + 4$	21.4	23.4	23.8	20.4	20.5	20.6	0.089	0.101	0.112

Table 4: Evaluation of our Fusion scheme. PSNR, Sharpness Loss and Laplacian measure: Higher is better.



Figure 8: Sharpening RGB predictions using our Fusion scheme on VSB100 (top two rows) and on UCF101 (bottom two rows).

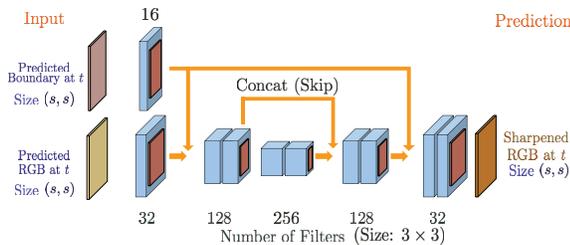


Figure 9: Our fusion model architecture.

top block. This baseline aims to de-blur RGB predictions without observing our predicted boundaries.

Evaluation. We observe improved and sharper RGB predictions (see Table 4)¹. Our fusion model learns to reintroduce lost high frequency information.

Conclusion

We propose the novel task of boundary prediction and demonstrate accurate results with our CMSC model. We argue for the key design choices, 1. A wide receptive field

¹Corresponding results in Table 5 in (Mathieu, Couprie, and LeCun 2016). We do not use motion masking as we would like our model to keep still boundaries intact.

allowing the model to learn complex spatio-temporal dependencies. 2. Accurate prediction at each time-step with a fully convolutional setup without any bottleneck layers. 3. The context which allows for information sharing thus leading to global consistency. We obtain sharp predictions using L2 loss (in contrast to RGB prediction, which leads to very blurry results with L2 loss). Predictions by our CMSC model on diverse scenarios shows that it developed a data-driven model of future boundary motions over long time horizons. This includes dynamics of moving agents and billiard balls. Moreover, while not being our primary goal, our predicted boundaries lead to sharper RGB video predictions via a fusion-based approach.

References

- Arbelaez, P.; Maire, M.; Fowlkes, C.; and Malik, J. 2011. Contour detection and hierarchical image segmentation. *TPAMI*.
- Baillargeon, R. 1994. How do infants learn about the physical world? *Current Directions in Psychological Science*.
- Baillargeon, R. 2004. Infants' physical world. *Current directions in psychological science* 13(3):89–94.
- Battaglia, P.; Pascanu, R.; Lai, M.; Rezende, D. J.; et al. 2016. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*, 4502–4510.

- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *ICML*.
- Chang, J.; Wei, D.; and Fisher, J. 2013. A video representation using temporal superpixels. In *CVPR*.
- Denton, E. L.; Chintala, S.; Szlam, A.; and Fergus, R. 2015. Deep generative image models using a laplacian pyramid of adversarial networks. In *NIPS*.
- Eigen, D.; Krishnan, D.; and Fergus, R. 2013. Restoring an image taken through a window covered with dirt or rain. In *CVPR*.
- Finn, C.; Goodfellow, I.; and Levine, S. 2016. Unsupervised learning for physical interaction through video prediction. In *NIPS*, 64–72.
- Fragkiadaki, K.; Agrawal, P.; Levine, S.; and Malik, J. 2016. Learning visual predictive models of physics for playing billiards. *ICLR*.
- Galasso, F.; Nagaraja, N.; Cardenas, T.; Brox, T.; and Schiele, B. 2013. A unified video segmentation benchmark: Annotation, metrics and analysis. In *CVPR*.
- Galasso, F.; Keuper, M.; Brox, T.; and Schiele, B. 2014. Spectral graph reduction for efficient image and streaming video segmentation. In *CVPR*.
- Jain, V.; Murray, J. F.; Roth, F.; Turaga, S.; Zhigulin, V.; Briggman, K. L.; Helmstaedter, M. N.; Denk, W.; and Seung, H. S. 2007. Supervised learning of image restoration with convolutional networks. In *CVPR*.
- Kalchbrenner, N.; Oord, A. v. d.; Simonyan, K.; Danihelka, I.; Vinyals, O.; Graves, A.; and Kavukcuoglu, K. 2017. Video pixel networks. *ICML*.
- Khoreva, A.; Benenson, R.; Galasso, F.; Hein, M.; and Schiele, B. 2016. Improved image boundaries for better video segmentation. In *ECCV Workshop*.
- Krotkov, E. P. 2012. *Active computer vision by cooperative focus and stereo*. Springer Science & Business Media.
- Lerer, A.; Gross, S.; and Fergus, R. 2016. Learning physical intuition of block towers by example. In *ICML*.
- Li, W.; Leonardis, A.; and Fritz, M. 2017. Visual stability prediction for robotic manipulation. In *IEEE International Conference on Robotics and Automation (ICRA)*. to appear.
- Liang, X.; Lee, L.; Dai, W.; and Xing, E. P. 2017. Dual motion gan for future-flow embedded video prediction. *ICCV*.
- Liu, Z.; Yeh, R.; Tang, X.; Liu, Y.; and Agarwala, A. 2017. Video frame synthesis using deep voxel flow. *ICCV*.
- Maninis, K.; Pont-Tuset, J.; Arbeláez, P.; and Gool, L. V. 2016. Convolutional oriented boundaries. In *ECCV*.
- Mao, X.-J.; Shen, C.; and Yang, Y.-B. 2016. Image restoration using convolutional auto-encoders with symmetric skip connections. *arXiv preprint arXiv:1606.08921*.
- Mathieu, M.; Couprie, C.; and LeCun, Y. 2016. Deep multi-scale video prediction beyond mean square error. *ICLR*.
- Michalski, V.; Memisevic, R.; and Konda, K. 2014. Modeling deep temporal dependencies with recurrent grammar cells. In *NIPS*.
- Ochs, P.; Malik, J.; and Brox, T. 2014. Segmentation of moving objects by long term video analysis. *TPAMI*.
- Patraucean, V.; Handa, A.; and Cipolla, R. 2015. Spatio-temporal video autoencoder with differentiable memory. *arXiv preprint arXiv:1511.06309*.
- Ranzato, M.; Szlam, A.; Bruna, J.; Mathieu, M.; Collobert, R.; and Chopra, S. 2014. Video (language) modeling: a baseline for generative models of natural videos. *arXiv:1412.6604*.
- Revaud, J.; Weinzaepfel, P.; Harchaoui, Z.; and Schmid, C. 2015. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *CVPR*.
- Srivastava, N.; Mansimov, E.; and Salakhutdinov, R. 2015. Unsupervised learning of video representations using lstms. In *ICML*.
- Sutskever, I.; Hinton, G. E.; and Taylor, G. W. 2009. The recurrent temporal restricted boltzmann machine. In *NIPS*.
- Villegas, R.; Yang, J.; Zou, Y.; Sohn, S.; Lin, X.; and Lee, H. 2017. Learning to generate long-term future via hierarchical prediction. *ICML, 2017*.
- Watters, N.; Tacchetti, A.; Weber, T.; Pascanu, R.; Battaglia, P.; and Zoran, D. 2017. Visual interaction networks. *arXiv preprint arXiv:1706.01433*.
- Wertheimer, M. 1923. Laws of organization in perceptual forms. *A source book of Gestalt Psychology*.

Appendix

We include here additional details of our model and results.

Further Details of the CMSC Model.

We include the internal details of each level L_k of our CMSC model in Table 5. We include the details of the type of layer, type specific details including the number and size of convolutional filters and pooling/upsampling layers, the non-linearity (activation) used after every layer, input and output of each layer.

We include details of each of the four levels of our CMSC model in Table 6. We include details of the scale (resolution) at which each level operates and the input to each level. We use the same notation as in the main paper, X_k denotes the input boundary image at scale $k \times k$ and $\hat{O}(L_k)$ is the up-sampled (by factor 2×2) output of level L_k . The final output is produced by the L_{96} level.

The central 32×32 patch of the output, produced by the L_{96} level, is considered valid and used to generate the boundary image at the next time-step.

Results on Moving MNIST.

We include results on the moving MNIST dataset (Srivastava, Mansimov, and Salakhutdinov 2015) to help compare our Convolutional Multi-Scale architecture against other frame prediction architectures. This dataset is suitable for this task because like boundary images the images in the moving MNIST dataset are also in the same domain $[0,1]$. However, we do not refer to this dataset in the main article as moving MNIST digits behave very differently compared to object boundaries. As the moving MNIST digits

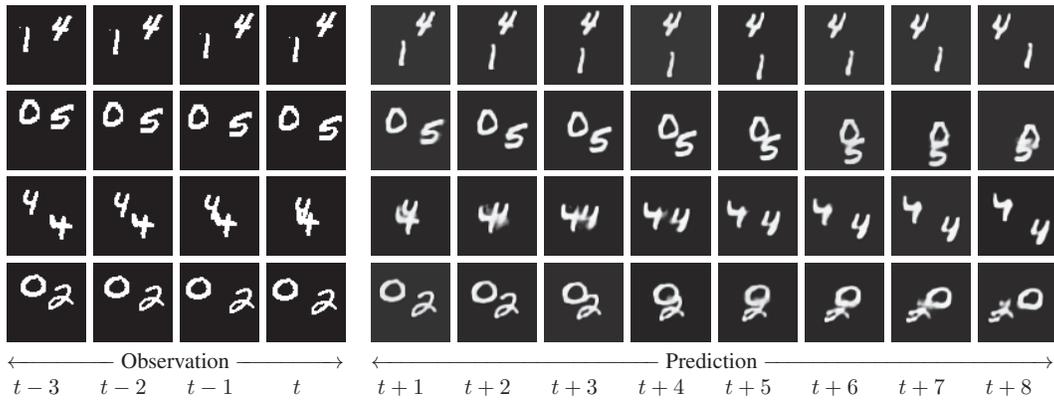


Figure 10: Example predictions on the moving MNIST dataset by our CMS model.

Layer	Type	Filters	Size	Activation	Input	Output
In ₁	Input					C ₁
C ₁	Conv	32	3×3	ReLU	In ₁	C ₂
C ₂	Conv	32	3×3	ReLU	C ₁	P ₁
P ₁	MaxPool		2×2		C ₂	C ₃
C ₃	Conv	64	3×3	ReLU	P ₁	C ₂
C ₄	Conv	64	3×3	ReLU	C ₃	P ₂
P ₂	MaxPool		2×2		C ₄	C ₅
C ₅	Conv	128	3×3	ReLU	P ₂	C ₆
C ₆	Conv	128	3×3	ReLU	C ₅	U ₁
U ₁	UpSample		2×2		C ₆	C ₇
C ₇	Conv	64	3×3	ReLU	U ₁	C ₈
C ₈	Conv	64	3×3	ReLU	C ₇	U ₂
U ₂	UpSample		2×2		C ₈	C ₉
C ₉	Conv	32	3×3	ReLU	U ₂	C ₁₀
C ₁₀	Conv	1	3×3	tanh	C ₉	

Table 5: Internal details of each level L_k of our CMSC model. Conv stands for 2D convolution, MaxPool stands for 2D max pooling and UpSample stands for 2D upsampling operations.

are of fixed size of 64x64 pixels, we do not use a context and thus use our CMS model for evaluation (which observes the full input image). This allows fair comparison against (Srivastava, Mansimov, and Salakhutdinov 2015; Patraucean, Handa, and Cipolla 2015; Kalchbrenner et al. 2017) which do not have a context and observes the full input image. We report quantitative results for prediction one time-step into the future (as in (Srivastava, Mansimov, and Salakhutdinov 2015; Patraucean, Handa, and Cipolla 2015)) in Table 7 using the Cross Entropy Loss (Srivastava, Mansimov, and Salakhutdinov 2015). Our CMS model outperforms (Srivastava, Mansimov, and Salakhutdinov 2015; Patraucean, Handa, and Cipolla 2015). Moreover, qualitative results in Figure 10 shows that our CMS model predicts accurately eight time-steps into the future. The highly complex model from (Kalchbrenner et al. 2017) performs better. However, this comparison shows that our CMS model compares favorably against other frame prediction models, while beating models with comparable number of parameters.

Level	Scale	Input
L ₁₂	12×12	X ₁₂
L ₂₄	24×24	{X ₂₄ , $\hat{O}(L_{12})$ }
L ₄₈	48×48	{X ₄₈ , $\hat{O}(L_{24})$ }
L ₉₆	96×96	{X ₉₆ , $\hat{O}(L_{48})$ }

Table 6: Details of the levels in our CMSC model.

Model	Cross Entropy Loss
(Srivastava, Mansimov, and Salakhutdinov 2015)	341.2
(Patraucean, Handa, and Cipolla 2015)	179.8
Our CMS	165.0
(Kalchbrenner et al. 2017)	87.6

Table 7: Evaluation on moving MNIST

Additional Results on VSB100.

Here, we show predictions at one, two and four steps ($t + 1$, $t + 2$, $t + 4$) in the future from a fixed time point on the airplane sequences of VSB100. We show the predictions in Figure 11. We use the same color coding as in the main article. That is, correct boundaries predictions are encoded in green, missed boundaries are encoded in yellow and wrong boundaries are encoded in red.

As expected from the quantitative performance in Figure 5 of the main article, the “Optic flow” baseline does not perform well. This method incorrectly translates the boundaries which lead to many boundaries being missed especially at $t + 4$. Compared to our CMSC model, the CSS and CMS models are unable to propagate motion in the long-term, leading to the disappearance of boundaries at $t + 4$. This highlights the importance of the context.

Running Time. Running time is GPU model and video resolution dependent. On the Nvidia Titan X GPU, our CMSC model takes approximately 16 hours to train on the VSB100 and real billiards datasets and 10 hours on synthetic

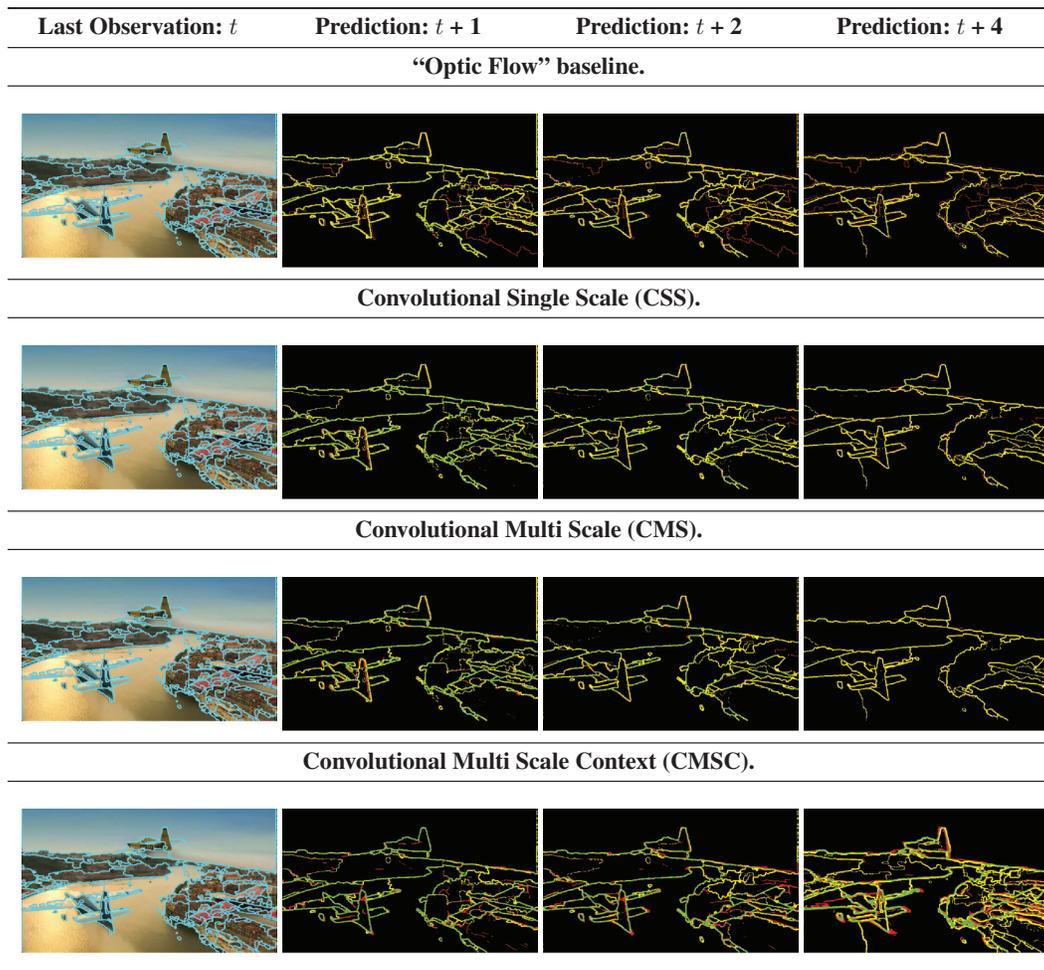


Figure 11: Prediction on *airplane* sequence from VSB100. Correct boundaries predictions are encoded in green. Missed boundaries are encoded in yellow. Wrong boundaries are encoded in red.

billiards (1 ball) dataset. During the test phase, prediction of one future frame of VSB100 (640×480 pixels) takes 1.03 seconds, synthetic billiards (256×256) 136 milliseconds and real billiards (320×240) 168 milliseconds on average.