

499 **A Architecture, Training, and Evaluation Details**

500 **A.1 Video Diffusion Training Details**

501 We use the same base architecture and training setup as [45] which utilizes Video U-Net architecture
502 with 3 residual blocks of 512 base channels and channel multiplier [1, 2, 4], attention resolutions [6,
503 12, 24], attention head dimension 64, and conditioning embedding dimension 1024. We use noise
504 schedule log SNR with range [-20, 20]. We make modifications Video U-Net to support first-frame
505 conditioning during training. Specifically, we replicate the first frame to be conditioned on at all
506 future frame indices, and apply temporal super resolution model condition on the replicated first
507 frame by concatenating the first frame channel-wise to the noisy data similar to [46]. We use temporal
508 convolutions as opposed to temporal attention to mix frames across time, to maintain local temporal
509 consistency across time, which has also been previously noted in [19]. We train each of our video
510 diffusion models for 2M steps using batch size 2048 with learning rate 1e-4 and 10k linear warmup
511 steps. We use 256 TPU-v4 chips for our first-frame conditioned generation model and temporal super
512 resolution model.

513 We use T5-XXL [22] to process input prompts which consists of 4.6 billion parameters. For
514 combinatorial and multi-task generalization experiments on simulated robotic manipulation, we train
515 a first-frame conditioned video diffusion models on 10x48x64 videos (skipping every 8 frames) with
516 1.7B parameters and a temporal super resolution of 20x48x64 (skipping every 4 frames) with 1.7B
517 parameters. The resolution of the videos are chosen so that the objects being manipulated (e.g.,
518 blocks being moved around) are clearly visible in the video. For the real world video results, we
519 finetune the 16x40x24 (1.7B), 32x40x24 (1.7B), 32x80x48 (1.4B), and 32x320x192 (1.2B) temporal
520 super resolution models pretrained on the data used by [19].

521 **A.2 Inverse Dynamics Training Details**

522 UniPi’s inverse dynamics model is trained to directly predict the 7-dimensional controls of the
523 simulated robot arm from an image observation mean squared error. The inverse dynamics model
524 consists of a 3x3 convolutional layer, 3 layers of 3x3 convolutions with residual connection, a
525 mean-pooling layer across all pixel locations, and an MLP layer of (128, 7) channels to predict the
526 final controls. The inverse dynamics model is trained using the Adam optimizer with gradient norm
527 clipped at 1 and learning rate 1e-4 for a total of 2M steps where linear warmup is applied to the first
528 10k steps.

529 **A.3 Baselines Training Details**

530 We describe the architecture details of various baselines below. The training details (e.g., learning
531 rate, warm up, gradient clip) of each baseline follow those of the inverse dynamics model detailed
532 above.

533 **Transformer BC [6, 26].** We employ the same transformer architecture as the 10M model of [26]
534 with 4 attention layers of 8 heads each and hidden size 512. We apply 4 layers of 3x3 convolution
535 with residual connection to extract image features, which, together with T5 text embeddings, are used
536 as inputs to the transformer. We additionally experimented with vision transformer style linearization
537 of the image patches similar to [26], but found the performance to be similar. We use a context length
538 of 4 and skip every 4 frames similar to UniPi’s inverse dynamics. We tried increasing the context
539 length of the transformer to 8 but it did not help improve performance.

540 **Transformer TT [25].** We use a similar transformer architecture as the Transformer BC baseline
541 detailed above. Instead of predicting the immediate next control in the sequence as in Transformer
542 BC, we predict the next 8 controls (skipping every 4 controls similar to other baselines) at the output
543 layer. We have also tried autoregressively predicting the next 8 controls, but found the errors to
544 accumulate quickly without additional discretization.

545 **State-Based Diffusion [21].** For the state-based diffusion baseline, we use a similar architecture
546 as UniPi’s first-frame conditioned video diffusion, where instead of diffusing and generating future
547 image frames, we replicate future controls across different pixel locations and apply the same U-Net
548 structure as UniPi to learn state-based diffusion models.

549 **A.4 Details of the Combinatorial Planning Task**

550 In the combinatorial planning tasks, we sample random 6 DOF poses for blocks, colored bowls, the
551 final placement box. Blocks start off uncolored (white) and must be placed in a bowl to obtain a
552 color. The robot then must manipulate and move the colored block to have the desired geometric

553 relation in the placement box. The underlying action space of the agent corresponds to 6 joint
554 values of robot plus a discrete contact action. When the contact action is active, the nearest block
555 on the table is attached to the robot gripper (where for methods that predict continuous actions, we
556 thresholded action prediction > 0.5 to correspond to contact). Given individual action predictions
557 for different models, we simulate the next state of the environment by running the joint controller in
558 Pybullet to try reach the predicted joint state (with a timeout of 2 seconds due to certain actions being
559 physically infeasible). As only a subset of the video dataset contained action annotations, we trained
560 the inverse-dynamics model on action annotations from 20k generated videos.

561 **A.5 Details of the CLIPort Multi-Environment Task**

562 In the CLIPort environment, we use the same action space as the combinatorial planning tasks and
563 execute actions similarly using the built in joint controller in Pybullet. As our training data, we use a
564 scripted agent on `put-block-in-bowl-unseen-colors`, `packing-unseen-google-objects-`
565 `seq`, `assembling-kits-seq-unseen-colors`, `stack-block-pyramid-seq-seen-colors`,
566 `tower-of-hanoi-seq-seen-colors`, `assembling-kits-seq-seen-colors`, `tower-of-`
567 `hanoi-seq-unseen-colors`, `stack-block-pyramid-seq-unseen-colors`, `packing-seen-`
568 `google-objects-seq`, `packing-boxes-pairs-seen-colors`, `packing-seen-google-`
569 `objects-group`. As our test data, we used the environments `put-block-in-bowl-seen-colors`,
570 `packing-unseen-google-objects-group`, `packing-boxes-pairs-unseen-colors`. We
571 trained the inverse dynamics on action annotation across the 200k generated videos.

572 **B Additional Results**

573 **B.1 Additional Results on Combinatorial Generalization**

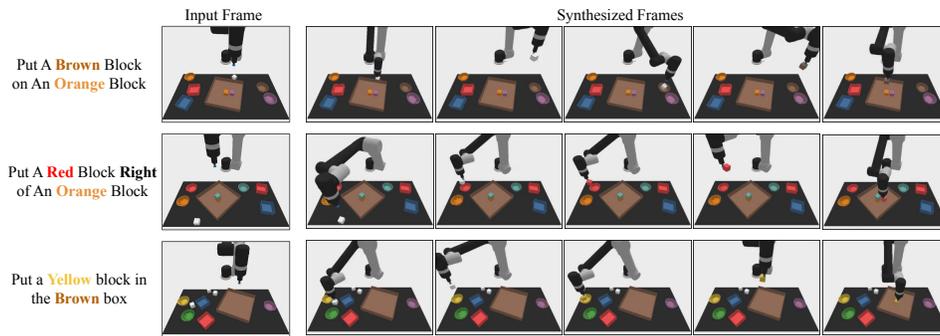


Figure 10: **Combinatorial Video Generation.** Additional results on UniPi’s generated videos for unseen language goals at test time.

574 **B.2 Additional Results on Multi-Environment Transfer**

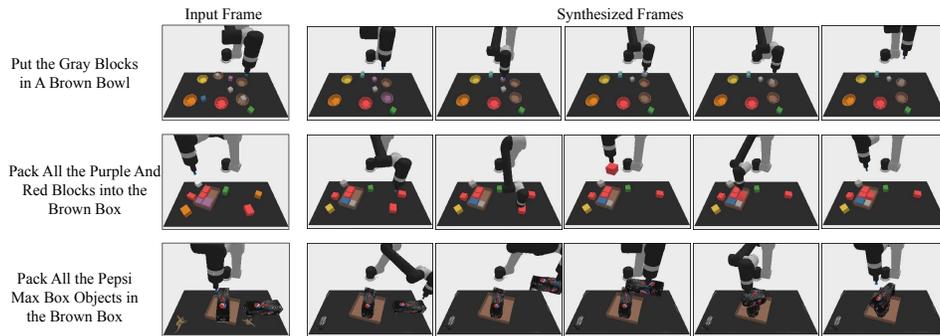


Figure 11: **Multitask Video Generation.** Additional results on UniPi’s generated video plans on different new tasks in the multitask setting.

575 **B.3 Additional Results on Real-World Transfer**

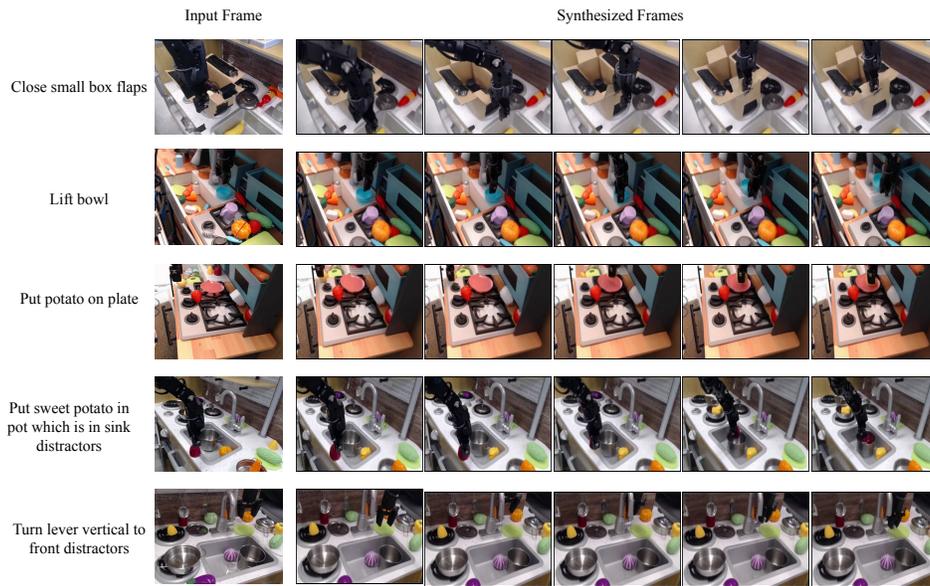


Figure 12: **High Fidelity Plan Generation.** Additional results on UniPi’s high resolution video plans across different language prompts.