

522 A Experimental Setup

523 Here, we discuss the implementation details of the experiments in Section 5. Source code is included
524 in the supplementary materials.

525 A.1 Batch Learning from Logged Compression Data

526 In our experiments, we found that initializing the compression model such that it outputs uniform-
527 random mask probabilities, collecting a batch of 1000 tuples ($T = 1, \mathbf{x}, \hat{\mathbf{x}}, \mathbf{a}$) using this random
528 compression model, and training the models D_ϕ , D_ψ , and f_θ to convergence on this data yielded
529 a high-performing compression model f_θ . In other words, rather than alternating one step of data
530 collection with one gradient step as in Algorithm 1, we used batch learning. The initial random
531 compression model explored the structured output space of feature masks well enough to generate
532 useful training data for our models, so we did not need to learn from on-policy data generated
533 by partially-trained compression models. This approach illustrates how PICO can be practically
534 deployed in real-world applications where other compression algorithms are already in use and have
535 generated large amounts of offline data \mathcal{D} , and where online learning may be difficult to implement.

536 In the digit identification, car shopping and survey, and car racing experiments, we set the compression
537 rate hyperparameter λ (see Section 4) to 0.5 during training. In the photo verification experiments,
538 we set $\lambda = 0.25$ during training.

539 A.2 Measuring Bitrates

540 For the digit identification, car racing, and car shopping and survey experiments, we use the following
541 procedure to measure compression rates. To estimate the prior distribution (introduced in Section
542 4), we fit a multivariate Gaussian distribution to the latent embeddings of the images in our training
543 set. To measure the number of bits needed to encode a given latent embedding, we normalize the
544 latent feature values to their z-scores, discretize the z-scores into bins of width 0.1, and sum the
545 negated base-2 log-probabilities of the discretized values under the prior distribution. For the photo
546 verification experiments, we use the base-2 KL-divergence between the latent posterior and prior in
547 the NVAE model [35].

548 In the digit identification experiments in Figure 3, we sweep $\lambda \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 1\}$ and
549 measure the resulting bitrates (the hyperparameter λ is defined in Section 4). In the car shopping
550 experiments in Figure 4, we sweep $\lambda \in \{0, 0.375, 0.5, 0.625, 1\}$. In the car survey experiments in
551 Figure 4, we sweep $\lambda \in \{0, 0.25, 0.5, 0.75, 1\}$. In the photo verification experiments in Figure 5,
552 we sweep $\lambda \in \{0, 0.25, 0.375, 0.5, 0.625, 0.75, 1\}$. In the car racing experiments in Figure 6, we set
553 $\lambda = 0.5$.

554 A.3 Network Architectures and Training

555 We use stochastic gradient descent – in particular, Adam [51] – to perform the optimization steps in
556 Algorithm 1.

557 In the car racing and digit identification experiments, we use a feedforward network with 2 layers of
558 256 units to represent the discriminators; to represent the compression model, the same architecture,
559 but with 64 instead of 256 units. In the car shopping and survey experiments, we use the same
560 architecture, but with 64 instead of 256 units, for the discriminators. In the photo verification ex-
561 periment, we combine the convolutional network architecture from [https://github.com/yzwxx/
562 vae-celebA/blob/master/model_vae.py](https://github.com/yzwxx/vae-celebA/blob/master/model_vae.py) with 2 additional fully-connected layers of 256 units
563 to represent the discriminators and the compression model.

564 A.4 Compressing Images using a Generative Model

565 Following up on the discussion in Section 4 about structuring the compression model $f_\theta(\hat{\mathbf{z}}|\mathbf{z})$, let
566 $\hat{\mathbf{z}} = [\mathbf{z}_1 \mathbf{z}_2]$ denote the decomposition of $\hat{\mathbf{z}}$ into masked features \mathbf{z}_1 and transmitted features \mathbf{z}_2 . In
567 the digit identification, photo verification, and car racing experiments, we set the masked features to

568 follow the distribution $\mathbf{z}_1 \sim \mathcal{N}(\bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\Sigma}})$, where

$$\begin{aligned}\bar{\boldsymbol{\mu}} &= \boldsymbol{\mu}_1 + \boldsymbol{\Sigma}_{12}\boldsymbol{\Sigma}_{22}^{-1}(\mathbf{z}_2 - \boldsymbol{\mu}_2), \\ \bar{\boldsymbol{\Sigma}} &= \boldsymbol{\Sigma}_{11} - \boldsymbol{\Sigma}_{12}\boldsymbol{\Sigma}_{22}^{-1}\boldsymbol{\Sigma}_{21}.\end{aligned}$$

569 We estimate $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ empirically, from the data used to train the generative model. In the car racing
570 experiments, at each timestep t , we set the prior mean $\boldsymbol{\mu}$ to the transmitted feature values at the
571 previous timestep $t - 1$, and estimate the prior covariance $\boldsymbol{\Sigma}$ empirically from state transition data.
572 In the car shopping and survey experiments, we sample the masked features from the StyleGAN2
573 prior – i.e., by feeding Gaussian noise input to the StyleGAN2 mapping network, and computing the
574 intermediate latents \mathbf{w} .

575 To make exploration easier (see Section 4 for discussion), we reduce the dimensionality d of the
576 mask output space by grouping together consecutive latent features. In the car racing experiments,
577 we train a VAE with 32 latent features using prior methods [52], and reduce the dimensionality of the
578 compression model’s output space from 32 to $d = 8$ by creating 8 groups of 4 latent features each –
579 where group 1 contains latent features 1-4, group 2 contains features 5-8, etc. – and masking groups
580 instead of masking individual features. In the car shopping experiments, we use a StyleGAN2 model
581 with 16 style layers – trained on the LSUN Car dataset using prior methods [34] – and reduce the
582 dimensionality to $d = 8$ by dividing the 16 style layers into 8 groups. In the car survey experiments,
583 we reduce to $d = 4$ using the same method. To encode images into the StyleGAN2 latent space, we
584 use the optimization-based projection method described in Section 5 of [34]. In the photo verification
585 experiments, we use the NVAE model for CelebA 64x64 described in Table 6 of [35]. We always
586 sample the latents in the second and third scales from the prior. For the latents in the first scale,
587 we reduce the dimensionality of the mask output space from $d = 5 \cdot 8^2$ to $d = 8$ by applying the
588 same mask to all 5 groups, and dividing the 64 latents into groups of 8. In the digit identification
589 experiments, we use a β -VAE with 10 latent features, which we do not group together as in the other
590 experiments.

591 In the digit identification, car shopping and survey, and car racing experiments, we use the latent
592 embedding \mathbf{z} instead of the full image \mathbf{x} as input to the discriminators – i.e., we set $D_\phi(\mathbf{a}, \mathbf{x}) \leftarrow$
593 $D_\phi(\mathbf{a}, \mathbf{z})$ and $D_\psi(\mathbf{p}, \mathbf{x}) \leftarrow D_\psi(\mathbf{p}, \mathbf{z})$.

594 A.5 Positive Examples for Discriminator Training

595 In the digit identification experiments, we treat 63,000 labeled images from the MNIST training set as
596 positive examples of user behavior without compression. In the photo verification experiments, we do
597 the same with 202,397 examples from the labeled CelebA training set – in particular, the Eyeglasses
598 and Hat labels. In the car shopping experiments, we automatically label the Ferrari, Bugatti, McLaren,
599 Aston Martin, Lamborghini, Spyker, and Porsche categories as unaffordable, and the Wagon, Minivan,
600 and Van categories as affordable, discard images belonging to any other categories, and treat 1,507
601 of the remaining labeled images as positive examples. In the car survey experiments, we collect
602 positive examples by eliciting 1,507 binary labels of “dark-colored” vs. “light-colored” on Amazon
603 Mechanical Turk.

604 A.6 Prompts for Amazon Mechanical Turk Participants

605 For the photo verification experiment in Figure 5 in which users check if eyes are covered:

606 In this task, you will examine photos of people and check if their eyes are covered.
607 Photos of people wearing eyeglasses or sunglasses should be classified as covered.
608 Choose the appropriate label that best suits the image: ‘Eyes are not covered’ or
609 ‘Eyes are covered’.

610 For the photo verification experiment in Figure 5 in which users check if heads are covered:

611 In this task, you will examine photos of people and check if their head is covered.
612 Photos of people wearing hats or caps should be classified as covered. Choose
613 the appropriate label that best suits the image: ‘Head is not covered’ or ‘Head is
614 covered’.

615 For the car shopping experiments in Figure 4

616 In this task, you will examine photos of cars and guess if they are affordable for
617 someone with a budget of approximately \$20,000. Choose the appropriate label
618 that best suits the image: 'Affordable' or 'Not affordable'.

619 For the car survey experiments in Figure 4:

620 In this task, you will examine photos of cars and determine if they are dark-colored
621 (black, dark blue, dark red, etc.) or light-colored (white, silver, light red, yellow,
622 etc.). Choose the appropriate label that best suits the image: 'Dark-colored car' or
623 'Light-colored car'.

624 For the handwritten digit identification experiments in Figure 3:

625 Choose the appropriate label that best suits the image: 0, 1, 2, 3, 4, 5, 6, 7, 8, or 9

626 A.7 Subject Allocation

627 **Amazon Mechanical Turk experiments.** In the car shopping and survey tasks, we assigned 10
628 users to label each of the 100 held-out images, in order to reduce the variance introduced by our
629 intentionally-vague prompts (see previous section). For the other AMT experiments, we only assigned
630 one user to each image, since we found that behavior did not vary substantially across users.

631 **Car racing video game experiment.** We recruited 10 male and 2 female participants, with an
632 average age of 25. Each participant was provided with the rules of the game and played 5 practice
633 episodes to familiarize themselves with the controls. To generative positive and negative examples
634 for training the PICO discriminator, we had a pilot user play 10 episodes without compression and 15
635 episodes with a compression model that outputs uniform-random mask probabilities. Each of the 12
636 participants played in both experimental conditions: with the non-adaptive compression baseline, and
637 with the trained compression model from PICO. To avoid the confounding effect of users learning to
638 play the game better over time, we counterbalanced the order of the two conditions. Each condition
639 lasted 15 episodes, with 100 timesteps (10 seconds) per episode.

640 B Subjective Evaluations in Car Racing Experiment

641 After evaluating the non-adaptive compression baseline:

642 It was quite hard to understand where the car/road were when the video got hazy
643 by the end of the training, the delay felt less significant. I almost didn't notice it. I
644 had difficulty figuring out what the environment wanted me to do when it would
645 bend the road far ahead of me but not near me.

646 It was often hard to tell if the car was moving or not, and the road sometimes
647 disappeared, which also made it hard to tell when steering was needed

648 Often the task wasn't too hard, but it was most challenging when the scene geometry
649 would suddenly shift and I couldn't anticipate how to react with my controls.

650 After evaluating PICO:

651 It was a lot more predictable and the blur was very infrequent. The road did behave
652 pretty unpredictably sometimes and I could not control

653 This environment was a lot easier. It felt more consistent. I felt like we had a
654 mutual understanding of when I would turn and what it would show me to make
655 me turn.

656 After model training much easier than before model training

657 This time around the task was a lot easier – the fact that the scene geometry changed
658 more naturally, and the fact that the effects of any delayed actions were predictable,
659 made it easier to decide how to steer

Table 1: Subjective Evaluations in the Car Racing User Study

| | <i>p</i> -value | Non-Adaptive | PICO |
|---|-----------------|--------------|-------------|
| I was able to keep the car on the road | < .0001 | 3.45 | 5.64 |
| I could anticipate the consequences of my steering actions | < .01 | 3.82 | 5.36 |
| I could tell when the car was about to go off road | < .01 | 3.55 | 5.36 |
| I could tell when I needed to steer to keep the car on the road | < .01 | 4.09 | 5.73 |
| I was often able to determine the car’s current position | < .001 | 4.00 | 5.82 |

Means reported for responses on a 7-point Likert scale, where 1 = Strongly Disagree, and 7 = Strongly Agree. *p*-values from a one-way repeated measures ANOVA with the use of PICO as a factor influencing responses.

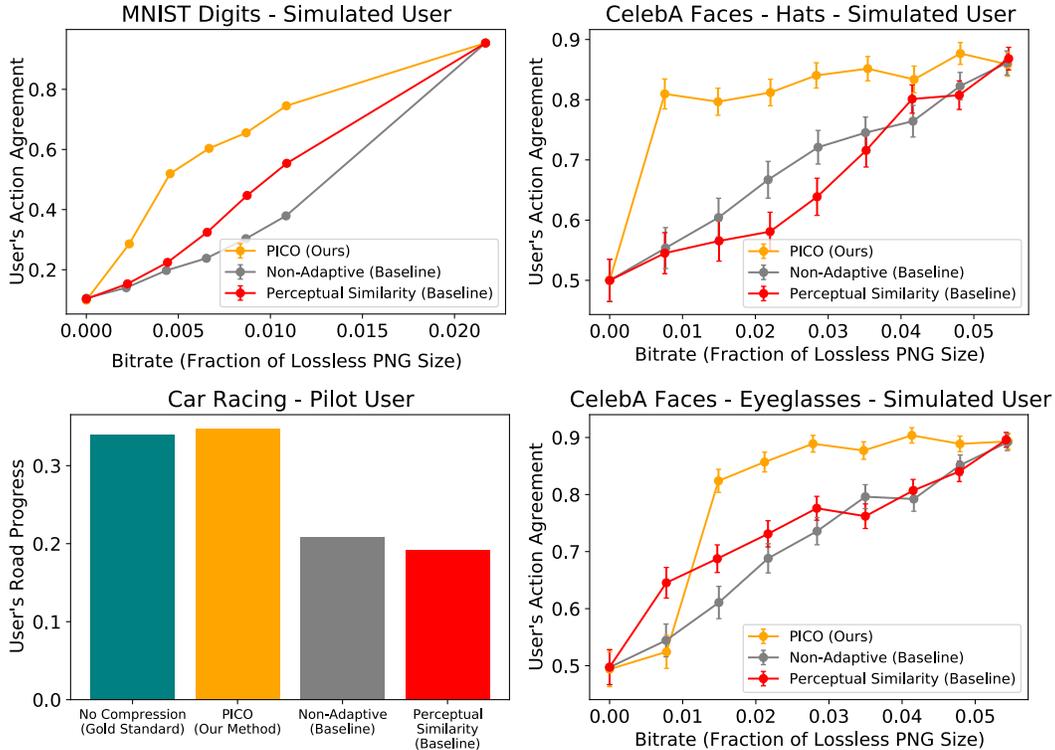


Figure 7: Experiments with simulated and pilot users show that PICO outperforms both baselines, and that the perceptual similarity baseline only performs better than the non-adaptive baseline on the digit identification task (top left).

660 C Simulation Experiments

661 To determine which baseline methods to compare with PICO in the user studies in Section 5, we ran
 662 preliminary experiments in which we simulated user behavior. In the digit identification and photo
 663 verification tasks, we simulated the user’s policy by training a classifier on labeled data (see Appendix
 664 A.5 for a description of the labeled data in each domain). In the car shopping, car surveying, and car
 665 racing tasks, we did not have enough labeled data to train a policy that qualitatively matched real user
 666 behavior. Hence, in the car racing task, we conducted a small-scale experiment with a single pilot
 667 user; and in the car shopping and surveying tasks, we perform a qualitative analysis of compressed
 668 image samples.

669 Figure 7 shows that PICO outperformed both the non-adaptive and perceptual similarity baselines
 670 in all domains. Furthermore, the perceptual similarity baseline only performed better than the non-
 671 adaptive baseline in the digit identification task; hence, our decision to omit the perceptual similarity
 672 baseline from the other user studies in Section 5. Figure 8 shows that, while PICO learns to preserve
 673 the perceived price of the car in the shopping task, and to preserve the color of the car in the survey
 674 task, the perceptual similarity baseline does not preserve either of the two features.

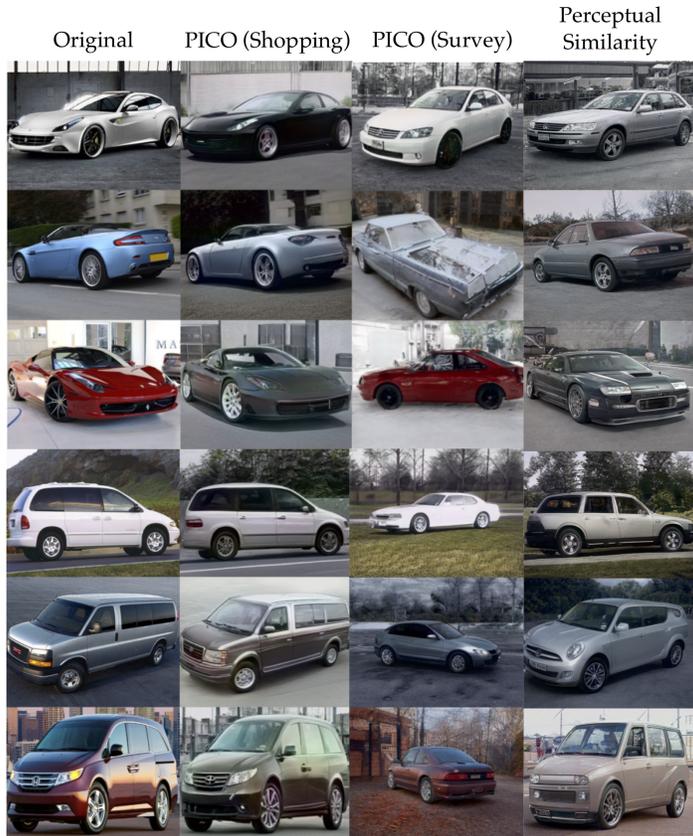


Figure 8: While PICO learns to preserve the perceived price of the car in the shopping task (second column), and to preserve the color of the car in the survey task (third column), the perceptual similarity baseline fails to preserve either of the two features (fourth column).

675 D Failure Cases

676 There are several ways in which PICO can fail to match the user’s actions with and without com-
 677 pression. For example, the latent embedding z produced by the pre-trained generative model (see
 678 Section 4) may lack the necessary features for performing the downstream task: the yellow sports
 679 car in rows 7-8 of Figure 10 gets distorted when encoded into the StyleGAN2 latent space, even
 680 without any additional compression. Another failure mode is for the latent features to be entangled,
 681 causing the structured mask output space of the compression model (see Section 4) to be insufficiently
 682 expressive for learning an effective compression policy: many of the compressed faces in Figure 9
 683 are visually distorted, most likely because the true prior distribution over latent embeddings is not
 684 modeled accurately by a Gaussian (see Appendix A.4).

685 E Examples of Compression at Different Bitrates

686 Figures 9 and 10 show that PICO tends to preserve task-relevant features like digit number, eyeglasses
 687 and hats, and the price and color of a car, more often than the non-adaptive baseline, and especially
 688 at lower bitrates. As the bitrate decreases, PICO discards task-irrelevant features before discarding
 689 task-relevant features. At extremely low bitrates (e.g., zero), PICO gracefully degrades to sampling a
 690 random image from the pre-trained generative model (see the right-most columns in Figures 9 and
 691 10), instead of, e.g., transmitting a heavily-distorted image with visual artifacts that make it difficult
 692 for the user to even attempt to perform their task.



Figure 9: Additional samples from the non-adaptive compression baseline and PICO drawn at varying bitrates. Left: digit identification experiments from Section 5.1 and Figure 3. Right: photo verification experiments from Section 5.2 and Figure 5.



Figure 10: Additional samples from the non-adaptive compression baseline and PICO drawn at varying bitrates, for the car shopping experiments in Section 5.1 and Figure 4