

483 **A Selecting ϵ for MNIST**

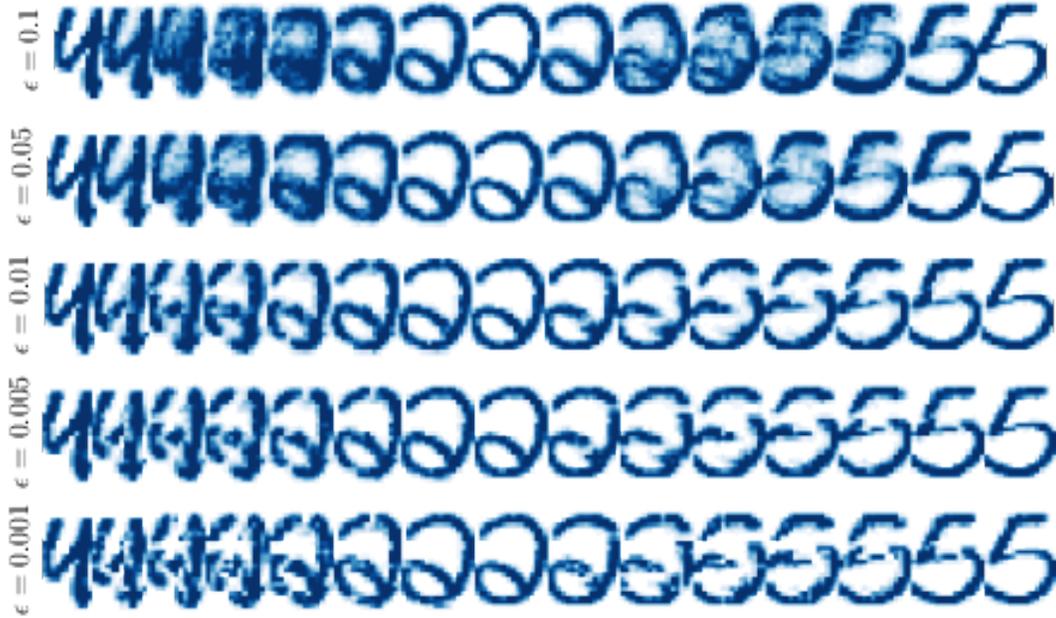


Figure 8: We selected $\epsilon = 10^{-2}$ for our MNIST coupling experiments as it results in transport maps that are not too blurry or sharp.

484 **B Other models for continuous OT**

485 While developing the hyper-network or Meta ICNN in [sect. 3.2](#) for predicting couplings between
 486 continuous measures, we considered alternative modeling formulations briefly documented in this
 487 section. We finalized only the hyper-network model because it is conceptually the most similar to
 488 predicting the optimal dual variables in the continuous setting and results in rapid predictions.

489 **B.1 Optimization-based meta-learning (MAML-inspired)**

490 The model-agnostic meta-learning setup proposed in MAML [[Finn et al., 2017](#)] could also be ap-
 491 plied in the Meta OT setting to learn an adaptable initial parameterization. In the continuous setting,
 492 one initial version would take a parameterized dual potential model $\psi_\varphi(x)$ and seek to learn an initial
 493 parameterization φ_0 so that optimizing a loss such as the W2GN loss \mathcal{L} from [eq. \(12\)](#) results in
 494 a minimal $\mathcal{L}(\varphi_K)$ after adapting the model for K steps. Formally, this would optimize:

$$\arg \min_{\varphi_0} \mathcal{L}(\varphi_K) \quad \text{where} \quad \varphi_{t+1} = \varphi_t - \nabla_{\varphi} \mathcal{L}(\varphi_t) \quad (18)$$

495 [Tancik et al. \[2021\]](#) explores similar learned initializations for coordinate-based neural implicit rep-
 496 resentations for 2D images, CT scan reconstruction, and 3d shape and scene recovery from 2D
 497 observations.

498 **Challenges for Meta OT.** The transport maps given by $T = \nabla \psi$ can significantly vary depending on
 499 the input measures α, β . We found it difficult to learn an initialization that can be rapidly adapted,
 500 and optimizing [eq. \(18\)](#) is more computationally expensive than [eq. \(17\)](#) as it requires unrolling
 501 through many evaluations of the transport loss \mathcal{L} . And, we found that *only* learning to predict
 502 the optimal parameters with [eq. \(17\)](#), conditional on the input measures, and then fine-tuning with
 503 W2GN to be stable.

504 **Advantages for Meta OT.** Exploring MAML-inspired methods could further incorporate the knowl-
 505 edge that the model’s prediction is going to be fine-tuned into the learning process. One promising

506 direction we did not try could be to integrate some of the ideas from LEO [Rusu et al., 2018] and
 507 CAVIA [Zintgraf et al., 2019], which propose learn a latent space for the parameters where the
 508 initialization is also conditional on the input.

509 B.2 Neural process

510 The (conditional) neural process models considered in Garnelo et al. [2018b,a] can also be adapted
 511 for the Meta OT setting. In the continuous setting, this would result in a dual potential that is also
 512 conditioned on a representation of the input measures, e.g. $\psi_\varphi(x; z)$ where $z := f_\varphi^{\text{emb}}(\alpha, \beta)$ is a
 513 learned embedding of the input measures that is learned with the parameters of ψ . This could be
 514 formulated as

$$\arg \min_{\varphi} \mathbb{E}_{(\alpha, \beta) \sim \mathcal{D}} \mathcal{L}(\varphi, f_\varphi^{\text{emb}}(\alpha, \beta)), \quad (19)$$

515 where \mathcal{L} modifies the model used in the loss eq. (12) to also be conditioned on the context extracted
 516 from the measures.

517 **Challenges for Meta OT.** This raises the issue on best-formulating the model to be conditional on
 518 the context. One way could be to append z to the input point x in the domain, but if ψ is an input-
 519 convex neural network, then the model would only need to be convex with respect to x and not z .

520 **Advantages for Meta OT.** A large advantage is that the representation z of the measures α, β would
 521 be significantly lower-dimensional than the parameters φ that our Meta OT models are predicting.

522 C Additional experimental and implementation details

523 We have attached the Jax source code necessary to run and reproduce all of the experiments in our
 524 paper and will open-source all of it. Here is a basic overview of the files:

```

525 |---meta_ot  Meta OT Python library code
    |   |---conjugate.py  Exact conjugate solver for the continuous setting
    |   |---data.py
    |   |---models.py
    |   |---utils.py
    |---config  Hydra configuration for the experiments (containing hyper-parameters)
    |---train_discrete.py  Train Meta OT models for discrete OT
    |---train_color_single.py  Train a single ICNN with W2GN between 2 images (for debugging)
    |---train_color_meta.py  Train a Meta ICNN with W2GN
    |---plot_mnist.py  Visualize the MNIST couplings
    |---plot_world_pair.py  Visualize the spherical couplings
    |---eval_color.py  Evaluate the Meta ICNN in the continuous setting
    |---eval_discrete.py  Evaluate the Meta ICNN for the discrete tasks
  
```

526 Connecting to the data is one difficulty in running the experiments. The easiest experiment to re-run
527 is the MNIST one, which will automatically download the dataset:

```
528  
529 1 ./train_discrete.py # Train the model, outputting to <exp_dir>  
530 2 ./eval_discrete.py <exp_dir> # Evaluate the learned models  
531 3 ./plot_mnist.py <exp_dir> # Produce further visualizations
```

533 We lastly summarize the hyper-parameters we used:

534 C.1 Hyper-parameters

535 Here we briefly summarize the hyper-parameters we used for training, which we did not extensively
536 tune. In the discrete setting, we use the same hyper-parameters for the MNIST and spherical settings.

Table 3: Discrete OT hyper-parameters.

Name	Value
Batch size	128
Number of training iterations	50000
MLP Hidden Sizes	[1024, 1024, 1024]
Adam learning rate	1e-3

537

Table 4: Continuous OT hyper-parameters.

Name	Value
Meta batch size (for α, β)	8
Inner batch size (to estimate \mathcal{L})	1024
Cycle loss weight (γ)	3.
Adam learning rate	1e-3
ℓ_2 weight penalty	1e-6
Max grad norm (for clipping)	1.
Number of training iterations	200000
Meta ICNN Encoder	ResNet18
Encoder output size (both measures)	256×2
Meta ICNN Decoder Hidden Sizes	[512]

538 **D Additional color transfer results**



Figure 9: Meta ICNN (initial prediction)

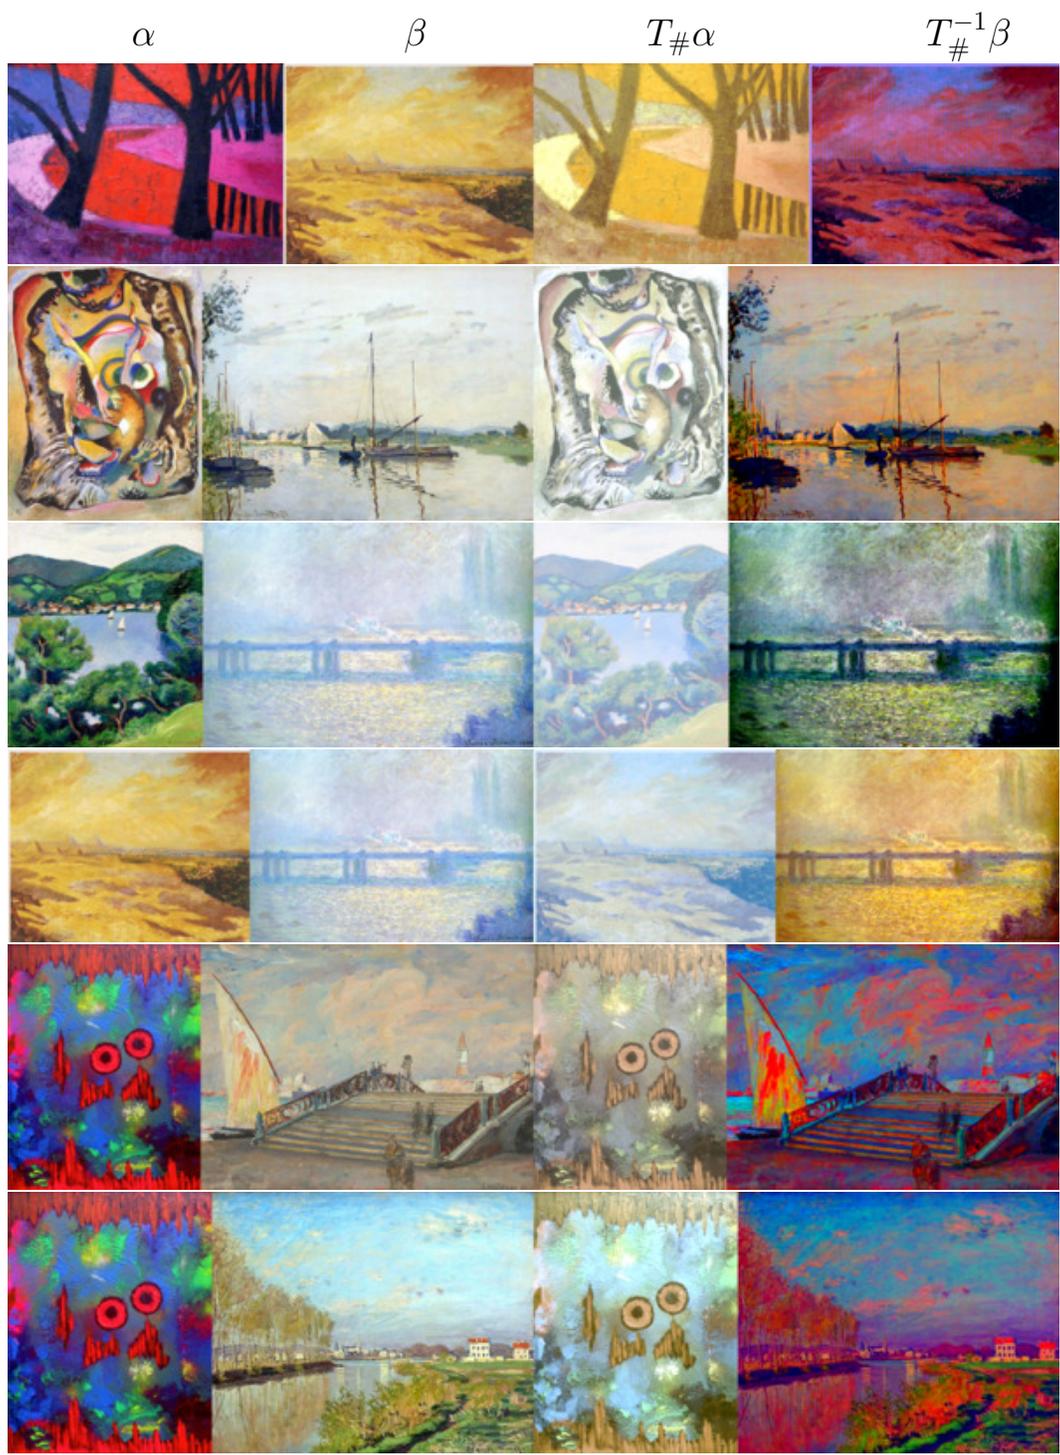


Figure 10: Meta ICNN + W2GN fine-tuning

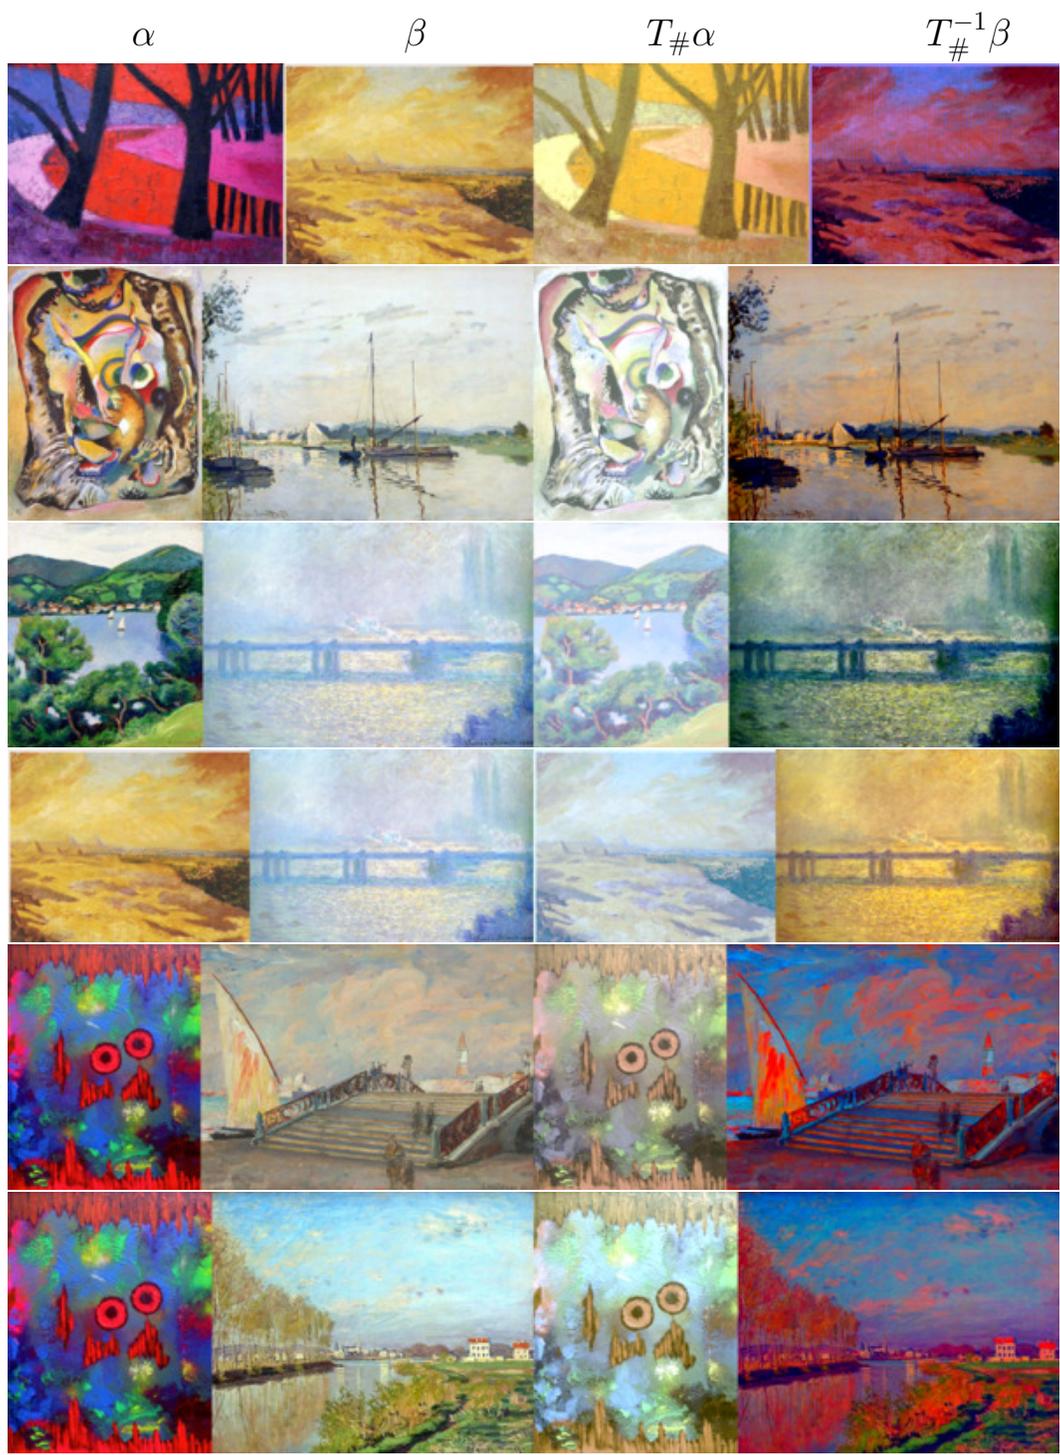


Figure 11: W2GN (final)