Appendix A Fourier theory on S^1

It is easy to see that every group automorphism of S^1 is given by $e^{i\theta} \mapsto e^{in\theta}$, for some integer n. So the dual group of S^1 is the set of integers, which is the domain for the Fourier transform of a function $f \in L^2(S^1)$. The Fourier transform of f is the function $\hat{f}$ defined as

$$\hat{f}(n) = \frac{1}{2\pi}\int_0^{2\pi} f(\theta)e^{-in\theta}d\theta \tag{1}$$

The functions $\{e^{in\theta}\}_{n\in\mathbb{Z}}$ are an orthonormal basis for functions in $L^2(S^1)$. The inverse Fourier transform gives the Fourier expansion of a function on the circle,

$$f(\theta) = \sum_{n=-\infty}^{n=\infty} \hat{f}(n)e^{in\theta} \tag{2}$$

Remark 8. *We say a function $f \in L^2(S^1)$ is transformed by S^1 (or use the terminology, S^1 acts on $f \in L^2(S^1)$) by defining $\theta_0 \cdot f := f_{\theta_0}$ to be the function, $f_{\theta_0}(\theta) := f(\theta - \theta_0)$.*
Clearly $(\theta_1 + \theta_2)\cdot f = \theta_1 \cdot (\theta_2 \cdot f)$ and the identity ($\theta = 0$) fixes f.

It is easy to see from equation 2 that

$$\widehat{f_{\theta_0}}(n) = e^{-in\theta_0}\hat{f}(n) \tag{3}$$

Recall that cross-correlation of $h \in L^2(S^1)$ with $f \in L^2(S^1)$ is the function

$$[h\star f](\theta_0) = \frac{1}{2\pi}\int_0^{2\pi} h(\theta-\theta_0)f(\theta)d\theta \tag{4}$$

An easy computation shows

$$\widehat{h\star f}(n) = \overline{\widehat{h}(n)}\cdot\widehat{f}(n) \tag{5}$$

So we have

Proposition 9. *Let $f,h \in L^2(S^1)$. Under the action of $\theta \in S^1$, the Fourier coefficients of the cross-correlation 4 transform as $\widehat{h\star f}(n) \mapsto e^{-in\theta}\widehat{h\star f}(n)$.*

Proof. Using Equation 3, under the action of $\theta \in S^1$, the n-th Fourier coefficient of $[h\star f]_\theta$ is $e^{-in\theta}\widehat{h\star f}(n)$. We check that this is also equal to $\widehat{h\star f_\theta}(n)$. Using Equation 3 and Equation 5, this is equal to $e^{-in\theta}\overline{\widehat{h}(n)}\cdot\widehat{f}(n)$, as required. □

Appendix B Representations of S^1

Definition 10. *Let V be a complex vector space. A representation of S^1 on V is a group homomorphism $\rho : S^1 \to GL(V)$. We say S^1 acts on V. Given $\theta \in S^1, v \in V$ we denote by $\theta\cdot v$, or $\theta(v)$, the element $\rho(\theta)v$.*

We recall the following theorem from harmonic analysis, see Carter et al. [2] for a reference.

Theorem 11. *If W is a representation of S^1 then $W = \widehat{\bigoplus}_{k\in\mathbb{Z}} W_k$ where*

$$W_k = \{w \in W : \theta \cdot w = e^{-ik\theta} w, \ for\ all\ \theta \in S^1\}$$

We recall a few more definitions from representation theory that we will need.

Definition 12. *Let W be a representation of S^1. A (closed) subspace $W' \subseteq W$ is said to be invariant under S^1 if for all $\theta \in S^1$, and $w' \in W'$ we have $\theta \cdot w' \in W'$. We say W is irreducible if there is no non trivial closed proper subspace which is invariant under S^1.*

We only deal with finite dimensional representations W, and in this case the number of summands is finite and all subspaces are closed, so can use the more standard notation $\oplus$ in the statement of the above theorem. From the above theorem it is clear that each W_k is an S^1-invariant subspace. Cleary the only irreducible subspaces are one dimensional subspaces contained in some W_k. We say the irreducible subspace of type k occurs in W with multiplicity $\dim(W_k)$.

B.1 S^1-equivariant maps

Definition 13. *Let V, W be two representations of S^1. A complex linear map $\phi : V \to W$ is said to be S^1-equivariant (or an S^1-morphism) if for all $v \in V, \theta \in S^1$, $\phi(\theta \cdot v) = \theta \cdot (\phi(v))$.*

Now if V is an irreducible representation of type n with basis vector v and W is an irreducible representation of type m with basis vector w, and $n \neq m$, then the only equivariant map from V to W is the zero map since $\phi(\theta \cdot v) = \phi(e^{-in\theta} v) = e^{-in\theta}\phi(v)$ but the right hand side in the definition above gives $e^{-im\theta}\phi(v)$. On the other hand if they are of the same type, then sending v to any complex multiple of w gives us an S^1-morphism. Note that any linear combination of vectors of the same type, is a vector of the same type. It follows that if the multiplicities of the irreducible representation of type k in V, W are m_k, n_k respectively, the dimension of the space of S^1-equivariant maps from V to W is $\sum_k m_k \cdot n_k$. Such a map is given by a block diagonal matrix, one block for each k of size $n_k \times m_k$.

Appendix C Proof: Bases W_{28} and W_{14} are coupled

We prove that the bases obtained in the coupled-autoencoder satisfy the definition of coupling. The proof is straightforward assuming that we have zero loss.

Let W_{28} and W_{14} be the final CW-bases learned. Following the bottom half of the network describing the coupling architecture, we see that under the morphism ϕ, the image space of W_{28}^T surjects onto the image space of W_{14}^T. Surjection follows since the dotted arrow going from W_{14} on the top to W_{14} below shows that we can start with any linear combination of the basis elements of W_{14} at the top and obtain it as the image of ϕ applied to an appropriate linear combination of basis elements of W_{28}. So the image space of W_{14}^T is isomorphic to a quotient of the image space of W_{28}^T. On the other hand following the top half of the network it is clear that ψ maps the image space of $W_{14}^T \oplus (W_{14}^T \otimes W_{14}^T)$ surjectively onto the image space of W_{28}^T. Surjectivity follows by reasoning as before. Hence the bases obtained in the coupling architecture satisfy definition 7.

Appendix D Bases learned for the translation group.

The bases learned for the translation group are visualized below.

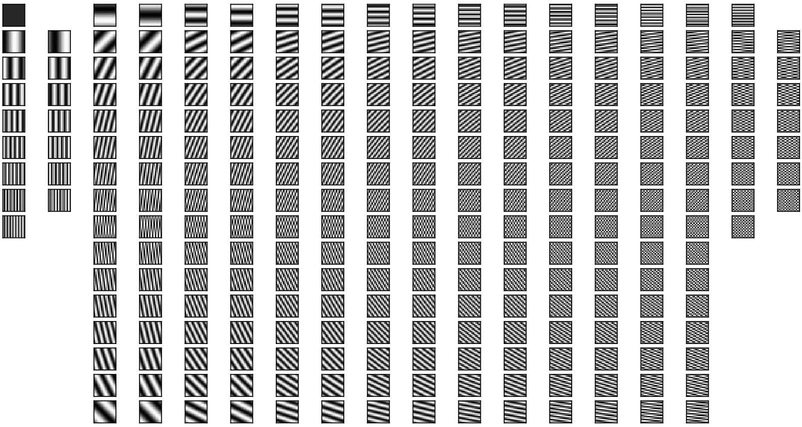

Figure 9: The W_{16} for translation group learned in AE

Appendix E Implementation details

The types and multiplicities of the Fourier coefficients in layer 0 and layer 1 of the autoencoder architecture (see Figure 2) used to discover the CW-basis of 28×28 images are given below. The first row gives the types and multiplicities of the activations in layer 0 and the second gives the types and multiplicities of activations in layer 1. Multiplicities are given in brackets. Here $\pm 1 - 4(5)$ means that types ± 1,± 2,± 3,and ± 4 were chosen to have multiplicity 5.

$$0(10), \pm 1-4(5), \pm 5-9(4), \pm 10-14(3), \pm 15-19(2), \pm 20-24(1)$$

$$0(8), \pm 1-4(4), \pm 5-9(3), \pm 10-14(2), \pm 15-19(1)$$

Appendix F Working over reals in Tensor Flow

Let S^1 act on a vector space W. It can be shown that for every irreducible subspace W_j with type $n_j \neq 0$ and basis vector w_j, there is an irreducible subspace with type $-n_j$, with basis vector the complex conjugate $\bar{w}_j$ of w_j, see for example [9][2.3.1]. In fact this can be deduced from our identification of $SO(2)$ with diagonal matrices having entries $e^{-i\theta}, e^{i\theta}$. Setting $b_{j1} = \frac{w_j + \bar{w}_j}{2}$ and $b_{j2} = \frac{w_j - \bar{w}_j}{2i}$ it is easy to see that b_{j1}, b_{j2} are real and they transform according to the columns of the matrix $R(n_j\theta)$. This two dimensional subspace (over the real numbers $\mathbb{R}$) of the image space is invariant to the real rotation group, $SO(2)$ and is irreducible for the action of $SO(2)$ on W. For our implementation purposes we work over real numbers. We call n_j the type of this irreducible representation of $SO(2)$. On the other hand subspaces W_j of type $n_j = 0$ with basis vector w_j, are invariant vectors, and satisfy $\rho(R(\theta))w_j = w_j$. These are one dimensional irreducible representations of $SO(2)$. Over reals Theorem 11 takes the form, $W = \oplus_{n \geq 0} W_n$. Here W_n is the subspace spanned by $SO(2)$-invariant subspaces of type n. Working over $\mathbb{R}$, the tensor product of two irreducible $SO(2)$-representations of type $s \geq t \geq 0$ splits into a direct sum of two irreducible $SO(2)$-representations of type $s+t, s-t$.

Appendix G Datasets used

We have used MNIST and MNIST_rot dataset in our experiments. Each sample in the dataset is a 28x28 gray scale image. MNIST dataset contains handwritten digits(upright, we call them NR). There are 60000 train samples and 10000 test samples. MNIST_rot dataset contains handwritten digits(rotated, we call them R) MNIST_rot has 12000 train samples and 50000 test samples.

We have used Fashion MNIST in our experiments. Each sample in the dataset is a 28x28 gray scale image. There are 10 different classes. Each label is one among the ten labels index by 0 to 9. The label index and the corresponding description is given in Table 3

Label	Description
0	T-shirt/top
1	Trouser
2	Pullover
3	Dress
4	Coat
6	Shirt
7	Sneaker
8	Bag
9	Ankle boot

Table 3: Fashion MNIST - Label index and Description

There are 60000 train samples and 10000 test samples. We call this set as upright (NR). For the rotated case (we call it as R), we rotate each sample in the test and the train by a random angle between 0 and 360.