A. Fourier Series and Approximation

Formally, the Fourier series is an expansion of a periodic function $f(x)$ of period $2L$ in terms of an infinite summation of sines and cosines. For clarity, we give the univariate case – the multivariate result can be found in literature.

$$f(x) = u_0 + \sum_{m=1}^{\infty} u_m \cos(m\omega_0 x) + \sum_{m=1}^{\infty} v_m \sin(m\omega_0 x), \tag{26}$$

where $\omega_0 \triangleq \frac{\pi}{L}$ and the coefficients for the series are:

$$u_m = \frac{1}{L}\int_{-L}^{L} f(x)\cos(m\omega_0 x)dx,$$
$$v_m = \frac{1}{L}\int_{-L}^{L} f(x)\sin(m\omega_0 x)dx.$$

By writing sine and cosine terms in their complex exponential forms, it is possible to define a complex Fourier series for real valued functions as

$$f(x) = \sum_{m=-\infty}^{m=\infty} c_m e^{im\omega_0 x}, \tag{27}$$
$$c_m = \frac{1}{2L}\int_{-L}^{L} f(x')e^{-im\omega_0 x'}dx'.$$

(26) and (27) are equivalent if we set c_m as:

$$c_m = \begin{cases} \frac{1}{2}(u_m + iv_m) & \text{for } m < 0, \\ u_0 & \text{for } m = 0, \\ \frac{1}{2}(u_m - iv_m) & \text{for } m > 0. \end{cases}$$

In reality, we cannot sum to infinity and instead use the series to approximate $f(x)$ to a finite value of m. Just as a Taylor series approximation becomes more accurate by using higher and higher order polynomials x^m, a Fourier series expansion becomes more accurate by using sinusoids of higher and higher frequencies $m\omega_0$. However, a Fourier series approximation approximates the function over its whole period, whereas the Taylor series does so only in a local neighbourhood of the given point.

Although the Fourier series is defined for periodic functions, it is still applicable to aperiodic functions. For bounded aperiodic functions, we define the period $2L$ to be the size of the domain of $f(x)$ and then integrate over this domain to obtain the Fourier coefficients. Intuitively, this is equivalent to repeating the bounded function periodically over an infinite domain. Aperiodic functions that are not bounded may be approximated by defining Fourier series over a bounded region of the function. As the size of this bounded region increases, and consequently the period $2L$ increases, the Fourier series approximation becomes more accurate and approaches a Fourier transform. Thus, for aperiodic unbounded functions, a Fourier series approximates a Fourier transform.

We now formalise the idea of taking the limit of the period going to infinity ($L \to \infty$) for a complex Fourier series representation of any general function $f(x)$. Firstly, it is convenient to rewrite (27) as:

$$f(x) = \frac{1}{2\pi}\sum_{m=-\infty}^{m=\infty} \int_{-L}^{L} f(x')e^{-im\omega_0 x'}dx' e^{im\omega_0 x}\omega_0.$$

Taking the limit as $L \to \infty$ (Stein & Shakarchi, 2003) gives

$$f(x) = \underbrace{\frac{1}{2\pi}\int_{\omega}\left(\overbrace{\int_{x'} f(x')e^{-i\omega x'}dx'}^{\mathcal{F}(f)}\right)e^{i\omega x}d\omega}_{\mathcal{F}^{-1}(\mathcal{F}(f))},$$

which is exactly equivalent to (7).

The integrals in the definition of the Fourier transform arise from taking a Fourier series representation of a function and letting the number of coefficients go to infinity.

B. n-Dimensional Fourier Transforms

Definitions Firstly, we make the definition of a n-dimensional Fourier transform precise: Consider a function $f(\cdot) : \mathbb{R}^n \to \mathbb{R}$. For $\boldsymbol{x} = (x_1, x_2, ...x_n)^\top \in \mathbb{R}^n$ and $\boldsymbol{\omega} = (\omega_1, \omega_2, ...\omega_n)^\top \in \mathbb{R}^n$, we have:

$$\mathcal{F}(f(x)) \triangleq \int_x f(\boldsymbol{x})e^{-i\boldsymbol{\omega}^\top \boldsymbol{x}}d\boldsymbol{x},$$
$$= \overbrace{\int_{x_1} ... \int_{x_n}}^{n} f(\boldsymbol{x})e^{-i\boldsymbol{\omega}^\top \boldsymbol{x}}dx_1...dx_n.$$

The corresponding n-Dimensional inverse Fourier transform is defined as:

$$\mathcal{F}^{-1}(f(x)) \triangleq \left(\frac{1}{2\pi}\right)^n \int_\omega f(\boldsymbol{x})e^{i\boldsymbol{\omega}^\top \boldsymbol{x}}d\boldsymbol{\omega},$$
$$= \left(\frac{1}{2\pi}\right)^n \overbrace{\int_{\omega_1} ... \int_{\omega_n}}^{n} f(\boldsymbol{x})e^{i\boldsymbol{\omega}^\top \boldsymbol{\omega}}d\omega_1...d\omega_n.$$

We define the Fourier transform of a vector/matrix quantity as simply the Fourier transform of individual elements of the vector/matrix. For example, the Fourier transform of matrix $\left[\boldsymbol{F}(\boldsymbol{x})\right]_{jk} = f_{jk}(\boldsymbol{x})$ is found from:

$$\left[\mathcal{F}(\boldsymbol{F}(\boldsymbol{x}))\right]_{jk} \triangleq \mathcal{F}(f_{jk}(\boldsymbol{x})). \tag{28}$$

And similarly for the inverse Fourier transform:

$$\left[\mathcal{F}^{-1}(\boldsymbol{F}(\boldsymbol{x}))\right]_{jk} \triangleq \mathcal{F}^{-1}(f_{jk}(\boldsymbol{x})).$$

Multiplication-Derivative Identities We now derive multi-dimension analogues to the single dimension multiplication-derivative property, which we state here:

$$\mathcal{F}\left(\frac{\partial}{\partial x_j} f(\boldsymbol{x})\right) = i\omega_j \mathcal{F}(f(\boldsymbol{x})). \tag{29}$$

Proofs of (29) are commonplace in Fourier Analysis references (Stein & Shakarchi, 2003). We start with a vector identity:

Lemma 1 (Multiplication-Derivative Property: Vectors)**.** *Given a function $f(\boldsymbol{x})$ with Fourier transform $\mathcal{F}\left(f(\boldsymbol{x})\right)$, multiplying $\mathcal{F}\left(f(\boldsymbol{x})\right)$ by the vector $i\boldsymbol{\omega}$ in the frequency domain is equivalent to taking the first order derivative $\nabla_{\boldsymbol{x}} f(\boldsymbol{x})$ in the action domain, that is:*

$$i\boldsymbol{\omega}\mathcal{F}\left(f(\boldsymbol{x})\right) = \mathcal{F}(\nabla_{\boldsymbol{x}} f(\boldsymbol{x})).$$

Proof. Consider the elements of the vector $i\boldsymbol{\omega}\mathcal{F}\left(f(\boldsymbol{x})\right)$:

$$\Big[i\boldsymbol{\omega}\mathcal{F}\left(f(\boldsymbol{x})\right)\Big]_j = i\omega_j \mathcal{F}\left(f(\boldsymbol{x})\right).$$

Using the single dimension multiplication-derivative property from (29) yields:

$$\Big[i\boldsymbol{\omega}\mathcal{F}\left(f(\boldsymbol{x})\right)\Big]_j = \mathcal{F}\left(\frac{\partial}{\partial x_j} f(\boldsymbol{x})\right).$$

Using the definition of the Fourier transform of a vector from (28) gives our main result:

$$i\boldsymbol{\omega}\mathcal{F}\left(f(\boldsymbol{x})\right) = \mathcal{F}(\nabla_{\boldsymbol{x}} f(\boldsymbol{x})).$$

□

We now derive a similar identity for matrices:

Lemma 2 (Multiplication-Derivative Property: Matrices)**.** *Given a function $f(\boldsymbol{x})$ with Fourier transform $\mathcal{F}\left(f(\boldsymbol{x})\right)$, multiplying $\mathcal{F}\left(f(\boldsymbol{x})\right)$ by the matrix $(i\boldsymbol{\omega})(i\boldsymbol{\omega})^\top$ in the frequency domain is equivalent to taking the second order derivative $\nabla^{(2)}_{\boldsymbol{x}} f(\boldsymbol{x})$ in the action domain, that is:*

$$(i\boldsymbol{\omega})(i\boldsymbol{\omega})^\top \mathcal{F}\left(f(\boldsymbol{x})\right) = \mathcal{F}\left(\nabla^{(2)}_{\boldsymbol{x}} f(\boldsymbol{x})\right).$$

Proof. Consider the elements of the matrix $(i\boldsymbol{\omega})(i\boldsymbol{\omega})^\top \mathcal{F}\left(f(\boldsymbol{x})\right)$:

$$\Big[(i\boldsymbol{\omega})(i\boldsymbol{\omega})^\top \mathcal{F}\left(f(\boldsymbol{x})\right)\Big]_{jk} = (i\omega_j)(i\omega_k)\mathcal{F}\left(f(\boldsymbol{x})\right).$$

Using the single dimension multiplication-derivative property from (29) twice yields:

$$\begin{aligned}\Big[(i\boldsymbol{\omega})(i\boldsymbol{\omega})^\top \mathcal{F}\left(f(\boldsymbol{x})\right)\Big]_{jk} &= (i\omega_j)\mathcal{F}\left(\frac{\partial}{\partial x_k} f(\boldsymbol{x})\right),\\ &= \mathcal{F}\left(\frac{\partial^2}{\partial x_j \partial x_k} f(\boldsymbol{x})\right).\end{aligned}$$

Using the definition of the Fourier transform of a matrix from (28) gives our main result:

$$i\boldsymbol{\omega}\mathcal{F}\left(f(\boldsymbol{x})\right) = \mathcal{F}(\nabla^{(2)}_{\boldsymbol{x}} f(\boldsymbol{x})).$$

□

C. Auxiliary Function Properties

Lemma 3 (nth Order Derivative of Auxiliary Function)**.** *Given an auxiliary function $\tilde{\beta}(\boldsymbol{\mu} - \boldsymbol{a}) = \beta(\boldsymbol{a})$ for a policy β, we may relate the m-th order derivative of $\tilde{\beta}$ w.r.t. $\boldsymbol{\mu}$ to the mth order derivative of β w.r.t. $\boldsymbol{a}$ as:*

$$\left(\nabla^{(m)}\tilde{\beta}\right)(\boldsymbol{\mu} - \boldsymbol{a}) = (-1)^n \nabla^{(m)}_{\boldsymbol{a}} \beta(\boldsymbol{a}) \ \forall \ m \geq 0.$$

Proof. For $m = 1$ From the chain rule we write:

$$\left(\nabla\tilde{\beta}\right)(\boldsymbol{\mu} - \boldsymbol{a}) = \nabla_{\boldsymbol{\mu}}\tilde{\beta}(\boldsymbol{\mu} - \boldsymbol{a}).$$

Let $\boldsymbol{\nu} = \boldsymbol{\mu} - \boldsymbol{a}$ s.t. $\tilde{\beta}(\boldsymbol{\mu} - \boldsymbol{a}) = \tilde{\beta}(\boldsymbol{\nu})$. Using the chain rule again for $\nabla_{\boldsymbol{\mu}}\tilde{\beta}(\boldsymbol{\mu} - \boldsymbol{a})$ yields:

$$\nabla_{\boldsymbol{\mu}}\tilde{\beta}(\boldsymbol{\nu}) = \nabla_{\boldsymbol{\mu}}\boldsymbol{\nu}\nabla_{\boldsymbol{\nu}}\boldsymbol{a}\nabla_{\boldsymbol{a}}\tilde{\beta}(\boldsymbol{\nu}).$$

Now, $\nabla_{\boldsymbol{\mu}}\boldsymbol{\nu} = \boldsymbol{I}$ and $\nabla_{\boldsymbol{\nu}}\boldsymbol{a} = -\boldsymbol{I}$. Substituting yields:

$$\nabla_{\boldsymbol{\mu}}\tilde{\beta}(\boldsymbol{\mu} - \boldsymbol{a}) = (-1)\nabla_{\boldsymbol{a}}\tilde{\beta}(\boldsymbol{\nu}).$$

Substituting $\tilde{\beta}(\boldsymbol{\nu}) = \tilde{\beta}(\boldsymbol{\mu} - \boldsymbol{a}) = \beta(\boldsymbol{a})$ gives our main result for $m = 1$:

$$\left(\nabla\tilde{\beta}\right)(\boldsymbol{\mu} - \boldsymbol{a}) = (-1)\nabla_{\boldsymbol{a}}\beta(\boldsymbol{a}).$$

Finally, taking $m - 1$ more derivatives will give our main result:

$$\left(\nabla^{(m)}\tilde{\beta}\right)(\boldsymbol{\mu} - \boldsymbol{a}) = (-1)^m \nabla^{(m)}_{\boldsymbol{a}} \beta(\boldsymbol{a}).$$

□

D. Turntable Experimental Setup Details

The turntable domain is a toy continuous control task. The goal is to align a disk to a desired angle by rotating it around its axis. The action is an angle in the range $a \in [-\pi, \pi]$ and the observations are the current position of the disk and the target position, both expressed as angles. The reward is set to $\sin(\alpha + \alpha_{\text{target}}) - \frac{1}{4}|a|$. For DPG, we used the OpenAI baseline implementation, where both the actor and the critic are represented using neural networks. For Fourier-EPG, we used the same setup but changed the critic to be trigonometric critic of the form $\sin(\alpha + \alpha_{\text{target}} - a) + w|a|$ with a tuneable weight w and the actor update given by Equation (23). The exploration policy was Gaussian with fixed standard deviation $\sigma = 0.05$ in both cases.

E. Gaussian Derivatives

We derive specific analytical solutions for the Gaussian policy $\beta = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ from Section 4.1. The following identities (Petersen & Pedersen, 2012) will be useful:

$$\nabla_{\boldsymbol{a}}\beta = -\boldsymbol{\Sigma}^{-1}(\boldsymbol{a}-\boldsymbol{\mu})\beta, \tag{30}$$
$$\nabla^{(2)}_{\boldsymbol{a}}\beta = \left(\boldsymbol{\Sigma}^{-1}(\boldsymbol{a}-\boldsymbol{\mu})(\boldsymbol{a}-\boldsymbol{\mu})^\top\boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}^{-1}\right)\beta. \tag{31}$$

Zeroth order $(M = 0)$ Substituting for $\nabla_{\boldsymbol{a}}\beta$ from (30) and $\nabla^{(2)}_{\boldsymbol{a}}\beta$ from (31) in (17) and (18) respectively, we obtain our analytic expression:

$$\begin{aligned}\hat{I}_{\boldsymbol{\mu}} &= \int_a \boldsymbol{\Sigma}^{-1}(\boldsymbol{a}-\boldsymbol{\mu})\hat{Q}\beta d\boldsymbol{a},\\ \hat{I}_{\boldsymbol{\Sigma}^{\frac{1}{2}}} &= \int_a \Big((\boldsymbol{\Sigma}^{\frac{1}{2}})^{-\top}(\boldsymbol{a}-\boldsymbol{\mu})(\boldsymbol{a}-\boldsymbol{\mu})^\top\boldsymbol{\Sigma}^{-1}\\ &\quad - (\boldsymbol{\Sigma}^{\frac{1}{2}})^{-\top}\Big)\hat{Q}\beta d\boldsymbol{a}.\end{aligned}$$

First order $(M = 1)$ Substituting for $\nabla_{\boldsymbol{a}}\beta$ from (30) in (20), we obtain our analytic expression:

$$\hat{I}_{\boldsymbol{\Sigma}^{\frac{1}{2}}} = \int_a (\boldsymbol{\Sigma}^{\frac{1}{2}})^{-\top}(\boldsymbol{a}-\boldsymbol{\mu})(\nabla_{\boldsymbol{a}}\hat{Q})^\top\beta d\boldsymbol{a}.$$

F. Proofs

Corollary 2.1. Let ψ be a parameter that does not depend upon $\boldsymbol{\mu}$. We can write $\hat{I}_\psi(\boldsymbol{s}_t) = \nabla_\psi \int_a \hat{Q}(\boldsymbol{s}_t, \boldsymbol{a})\beta_\theta(\boldsymbol{a}|\boldsymbol{s}_t)d\boldsymbol{a}$ as:

$$\hat{I}_\psi(\boldsymbol{s}) = \mathcal{F}^{-1}\Big(\mathcal{F}(\hat{Q})\nabla_\psi\mathcal{F}(\tilde{\beta})\Big)(\boldsymbol{\mu}).$$

Proof. Using Theorem 2, we obtain the following expression for $\hat{I}_\psi(\boldsymbol{s}_t)$:

$$\hat{I}_\psi(\boldsymbol{s}_t) = \nabla_\psi\mathcal{F}^{-1}\Big(\mathcal{F}(\hat{Q})\mathcal{F}(\tilde{\beta})\Big)(\boldsymbol{\mu}).$$

Using Leibniz's rule for integration under the integral, we move the derivative inside of the inverse Fourier transform, obtaining our result:

$$\hat{I}_\psi(\boldsymbol{s}_t) = \mathcal{F}^{-1}\Big(\mathcal{F}(\hat{Q})\nabla_\psi\mathcal{F}(\tilde{\beta})\Big)(\boldsymbol{\mu}).$$

□

G. Complete Periodic Critic Derivation

We now derive the analytic update from (22) for our periodic critic. Firstly, for ease of analysis we re-write our critic using the hyperbolic function:

$$\begin{aligned}\hat{Q}(a) &= \cos(\boldsymbol{f}^\top\boldsymbol{a} - h),\\ &= \frac{e^{i(\boldsymbol{f}^\top\boldsymbol{a}-h)} + e^{-i(\boldsymbol{f}^\top\boldsymbol{a}-h)}}{2},\\ &= \frac{e^{-ih}e^{i\boldsymbol{f}^\top\boldsymbol{a}} + e^{ih}e^{-i\boldsymbol{f}^\top\boldsymbol{a}}}{2}.\end{aligned}$$

Taking the Fourier transform yields:

$$\begin{aligned}\mathcal{F}\left(\hat{Q}\right) &= \frac{1}{2}\Bigg[e^{-ih}(2\pi)^n\prod_{j=1}^n\delta(\omega_j - f_j)\\ &\quad + e^{ih}(2\pi)^n\prod_{j=1}^n\delta(\omega_j + f_j)\Bigg],\\ &= (2\pi)^n\left[\frac{e^{-ih}\delta(\boldsymbol{\omega}-\boldsymbol{f}) + e^{ih}\delta(\boldsymbol{\omega}+\boldsymbol{f})}{2}\right].\end{aligned}$$

Recall that the characteristic function of the Gaussian auxiliary function is $\mathcal{F}\left(\tilde{\beta}\right) = e^{-\boldsymbol{\omega}^\top\boldsymbol{\Sigma}\boldsymbol{\omega}}$. Now taking inverse Fourier transforms of $\mathcal{F}(\hat{Q})\mathcal{F}(\tilde{\beta})$ yields:

$$\begin{aligned}\mathcal{F}^{-1}(\mathcal{F}(\hat{Q})\mathcal{F}(\tilde{\beta}))(\boldsymbol{a}) &= \frac{1}{(2\pi)^n}\int\mathcal{F}(\hat{Q})\mathcal{F}(\tilde{\beta})e^{i\boldsymbol{\omega}^T\boldsymbol{a}}d\boldsymbol{\omega},\\ &= \frac{1}{2}\int e^{-\boldsymbol{\omega}^\top\boldsymbol{\Sigma}\boldsymbol{\omega}}\left[e^{-ih}\delta(\boldsymbol{\omega}-\boldsymbol{f}) + e^{ih}\delta(\boldsymbol{\omega}+\boldsymbol{f})e^{i\boldsymbol{\omega}^T\boldsymbol{a}}\right]d\boldsymbol{\omega},\\ &= \frac{1}{2}\int e^{-\boldsymbol{\omega}^\top\boldsymbol{\Sigma}\boldsymbol{\omega}}\left[e^{i(\boldsymbol{\omega}^T\boldsymbol{a}-h)}\delta(\boldsymbol{\omega}-\boldsymbol{f}) + e^{i(\boldsymbol{\omega}^T\boldsymbol{a}+h)}\delta(\boldsymbol{\omega}+\boldsymbol{f})\right]d\boldsymbol{\omega},\\ &= e^{-\boldsymbol{f}^\top\boldsymbol{\Sigma}\boldsymbol{f}}\left[\frac{e^{i(\boldsymbol{f}^T\boldsymbol{a}-h)} + e^{-i(\boldsymbol{f}^T\boldsymbol{a}-h)}}{2}\right],\\ &= e^{-\boldsymbol{f}^\top\boldsymbol{\Sigma}\boldsymbol{f}}\cos(\boldsymbol{f}^T\boldsymbol{a} - h),\end{aligned}$$

where we have used the hyperbolic definition of $\cos$ to derive our desired result in the final line.