

Supplementary Material

A. Notation Summary

Table 6. Abbreviation.

OCT/OCF	Optimal Control Theory/Programming
MPDG	Multi-Player Dynamic Game
CG	Cooperative Game
OLNE	Open-loop Nash Equilibria
FNE	Feedback Nash Equilibria
GR	Group Rationality
IR	Individual Rationality

Table 7. Terminology mapping.

MPDG		Training generic (non-Markovian) DNNs
t	Stage order	Computation order from input to output
n	Player index	Index of parallel layers aligned at t
$f_{t,n}$	-	Layer module indexed by (t, n)
F_t	Shared dynamics	Joint propagation rule of layers $\{f_{t,n} : n \in [N]\}$
$\theta_{t,n}$	Action committed at stage t by Player n	Trainable parameter of layer $f_{t,n}$
$\mathbf{z}_{t,n}$	-	Pre-activation vector of layer $f_{t,n}$
\mathbf{x}_t	State at stage t	Augmentation of pre-activation vectors of layers $\{f_{t,n} : n \in [N]\}$
$\ell_{t,n}$	Cost incurred at stage t for Player n	Weight decay for layer $f_{t,n}$
ϕ_n	Cost incurred at final stage T for Player n	Lost w.r.t. network output (e.g. cross entropy in classification)

Table 8. Dynamic game theoretic terminology w.r.t. different optimality principles.

OLNE	$\eta_{t,n}^O$	Open-loop information structure
	$H_{t,n}$	Optimality objective (Hamiltonian) for OLNE
	$\mathbf{p}_{t,n}$	Co-state at stage t for Player n
FNE	$\eta_{t,n}^C$	Feedback information structure
	$Q_{t,n}$	Optimality objective (Isaacs-Bellman objective) for FNE
	$V_{t,n}$	Value function for FNE
	$\mathbf{k}_{t,n}$	Open gain of the locally optimal update for FNE
	$\mathbf{K}_{t,n}$	Feedback gain of the locally optimal update for FNE
GR	$\eta_{t,n}^{O-CG}$	Cooperative open-loop information structure
	$\eta_{t,n}^{C-CG}$	Cooperative feedback information structure
	P_t	Optimality objective (group Bellman objective) for GR
	W_t	Value function for GR
	$\tilde{\mathbf{k}}_t$	Open gain of the locally optimal update for GR
	$\tilde{\mathbf{K}}_t$	Feedback gain of the locally optimal update for GR

B. OCP Characterization of Training Feedforward Networks

The optimality principle to OCP (2), or equivalently the training process of feedforward networks, can be characterized by Dynamic Programming (DP) or Pontryagin Principle (PP). We synthesize the related results below.

Theorem 7 (Bellman (1954); Pontryagin et al. (1962)).

(DP) Define a value function V_t computed recursively by the Bellman equation (17), starting from $V_T(z_T) = \phi(z_T)$,

$$V_t(z_t) = \min_{\pi_t} Q_t(z_t, \theta_t), \quad \text{where} \quad Q_t(z_t, \theta_t) := \ell_t(\theta_t) + V_{t+1}(f_t(z_t, \theta_t)) \quad (17)$$

is called the Bellman objective. $\pi_t \equiv \theta_t(z_t)$ is an arbitrary mapping from z_t to θ_t . Let π_t^* be the minimizer to (17), then $\{\pi_t^* : t \in [T]\}$ is the optimal feedback policy to (2).

(PP) The optimal trajectory $\pi_t^* \equiv \theta_t^*(z_t^*)$ along (17) obeys

$$z_{t+1}^* = \nabla_{p_{t+1}} H_t(z_t^*, p_{t+1}^*, \theta_t^*), \quad z_0^* = z_0, \quad (18a)$$

$$p_t^* = \nabla_{z_t} H_t(z_t^*, p_{t+1}^*, \theta_t^*), \quad p_T^* = \nabla_{z_T} \phi(z_T^*), \quad (18b)$$

$$\theta_t^* = \arg \min_{\theta_t} H_t(z_t^*, p_{t+1}^*, \theta_t), \quad (18c)$$

where (18b) is the adjoint equation for the co-state p_t^* and

$$H_t(z_t, p_{t+1}, \theta_t) := \ell_t(\theta_t) + p_{t+1}^\top f_t(z_t, \theta_t)$$

is the discrete-time Hamiltonian.

Theorem 7 provides an OCP characterization of training feedforward networks. First, notice that the time-varying OCP objectives (Q_t, H_t) are constructed through some backward processes similar to the Back-propagation (BP). Indeed, one can verify that the adjoint equation (18b) gives the exact BP dynamics. Similarly, the dynamics of V_t in (17) also relate to BP under some conditions (Liu et al., 2021). The parameter update, $\theta_t \leftarrow \theta_t - \delta \theta_t$, for standard training methods can be seen as solving the discrete-time Hamiltonian H_t with different precondition matrices (Li et al., 2017a). On the other hand, DDPNOpt (Liu et al., 2021) minimizes the time-dependent Bellman objective Q_t with $\theta_t \leftarrow \theta_t - \delta \pi_t$. This elegant connection is, however, limited to the interpretation between feedforward networks and Markovian dynamical systems (1).

C. Missing Derivations in Section 3




Figure 12. Forward propagation (left) and Back-propagation (right) of a residual block and how each quantity connects to OLNE optimality.

Proof of Proposition 2. Expand the expression of the Hamiltonian in OLNE:

$$H_{t,n}(x_t, p_{t+1,n}, \theta_{t,1}, \dots, \theta_{t,N}) := \ell_{t,n}(\theta_{t,1}, \dots, \theta_{t,N}) + F_t(x_t, \theta_{t,1}, \dots, \theta_{t,N})^\top p_{t+1,n},$$

where $p_{t,n}$ is the co-state whose dynamics obey

$$p_{t,n} = \nabla_{x_t} H_{t,n}, \quad p_{T,n} = \nabla_{x_T} \phi_n(x_T).$$

Recall §3.1 where we demonstrate that for training generic DNNs, one shall consider $\ell_{t,n} := \ell_{t,n}(\theta_{t,n})$ and $\phi_n := \phi$. Hence, the dynamics of $p_{t,n}$ become

$$p_{t,n} = \nabla_{x_t} H_{t,n}, \quad p_{T,n} = \nabla_{x_T} \phi(x_T), \quad \text{where } H_{t,n} = \ell_{t,n}(\theta_{t,n}) + F_t(x_t, \theta_{t,1}, \dots, \theta_{t,N})^\top p_{t+1,n}. \quad (19)$$

Our goal is to show that (19) gives the exact Back-propagation dynamics. First, notice that the terminal condition of (19), *i.e.* $\mathbf{p}_{T,n} = \nabla_{\mathbf{x}_T} \phi$, is already the gradient w.r.t. the network output without any manipulation. Next, to show that $\mathbf{p}_{t,n}$ corresponds to the Back-propagation at stage t , consider, for instance, the computation graphs of the residual block in Fig. 12, where we replot Fig. 2 together with its Back-propagation dynamic and denote \mathbf{q} as the gradient w.r.t. the activation vector \mathbf{z} . Then, it can be shown by induction that $\mathbf{p}_{t,n}$ augments all “ \mathbf{q} ”s aligned at stage t . Indeed, suppose $\mathbf{p}_{t+1,n}$ is the augmentation of the Back-propagation gradients at stage $t+1$, *i.e.* $\mathbf{p}_{t+1,n} := [\mathbf{q}_{t+1}, \mathbf{q}_{\text{res}}]^\top$, then the co-state at the current stage t can be expanded as

$$\mathbf{p}_{t,n} = \nabla_{\mathbf{x}_t} H_{t,n} = \nabla_{\mathbf{x}_t} F_t^\top \mathbf{p}_{t+1,n} = \begin{bmatrix} \nabla_{\mathbf{z}_t} f_{t,1} & \nabla_{\mathbf{z}_{t-1}} f_{t,1} \\ \nabla_{\mathbf{z}_t} f_{t,2} & \nabla_{\mathbf{z}_{t-1}} f_{t,2} \end{bmatrix}^\top \begin{bmatrix} \mathbf{q}_{t+1} \\ \mathbf{q}_{\text{res}} \end{bmatrix} = \begin{bmatrix} \nabla_{\mathbf{z}_t} f_{t,1}^\top \mathbf{q}_{t+1} \\ \nabla_{\mathbf{z}_{t-1}} f_{t,2}^\top \mathbf{q}_{\text{res}} \end{bmatrix} = \begin{bmatrix} \mathbf{q}_t \\ \mathbf{q}_{t-1} \end{bmatrix},$$

which augments all “ \mathbf{q} ”s at stage t . Once we connect $\mathbf{p}_{t,n}$ to the Back-propagation dynamics, it can be verified that

$$H_\theta^{t,n} \equiv \nabla_{\theta_{t,n}} H_{t,n} = \nabla_{\theta_{t,n}} \ell_{t,n} + \nabla_{\theta_{t,n}} F_t^\top \mathbf{p}_{t+1,n}.$$

is indeed the gradient w.r.t. the parameter $\theta_{t,n}$ of each layer $f_{t,n}$. Therefore, taking the iterative update $\theta_{t,n} \leftarrow \theta_{t,n} - H_\theta^{t,n}$ is equivalent to descending along the SGD direction, up to a learning rate scaling. Similarly, setting different precondition matrices \mathbf{M} will recover other standard methods. Hence, we conclude the proof. \square

Optimality principle for $\eta_{t,n}^{\text{O-CG}}$. For the completeness, below we provide the optimality principle for the cooperative open-loop information structure $\eta_{t,n}^{\text{O-CG}}$.

Definition 8 (Cooperative optimality principle by $\eta_{t,n}^{\text{O-CG}}$). *A set of strategy, $\{\theta_{t,n}^* : \forall t, n\}$, provides an open-loop optimal solution to the joint optimization (4) if*

$$\theta_{t,1}^*, \dots, \theta_{t,N}^* = \arg \min_{\theta_{t,n} : n \in [N]} \tilde{H}_t(\mathbf{x}_t, \tilde{\mathbf{p}}_{t+1}, \theta_{t,1}, \dots, \theta_{t,N}),$$

where $\theta_{t,n}^* \equiv \theta_{t,n}^*(\eta_{t,n}^{\text{O-CG}})$ and $\tilde{H}_t := \sum_{n=1}^N \ell_{t,n} + F_t^\top \tilde{\mathbf{p}}_{t+1}$

is the “group” Hamiltonian at stage t . Similar to OLNE, the joint co-state $\tilde{\mathbf{p}}_t$ can be simulated by

$$\tilde{\mathbf{p}}_t = \nabla_{\mathbf{x}_t} \tilde{H}_t, \quad \tilde{\mathbf{p}}_T = \sum_{n=1}^N \nabla_{\mathbf{x}_T} \phi_n.$$

In this work, we focus on solving the optimality principle inherited in $\eta_{t,n}^{\text{C-CG}}$ as a representative of the CG optimality. Since $\eta_{t,n}^{\text{O-CG}} \subset \eta_{t,n}^{\text{C-CG}}$, the latter captures richer information and tends to perform better in practice, as evidenced by Fig. 3.

D. Missing Derivations in Section 4

D.1. Complete Derivation of the Iterative Updates

Derivation of FNE update. Our goal is to approximately solve the Isaacs-Bellman recursion (6) only up to second-order. Recall that the second-order expansion of $Q_{t,n}$ at some fixed point $(\mathbf{x}_t, \theta_{t,n})$ takes the form

$$Q_{t,n} \approx \frac{1}{2} \begin{bmatrix} \mathbf{1} \\ \delta \mathbf{x}_t \\ \delta \theta_{t,n} \end{bmatrix}^\top \begin{bmatrix} Q_{t,n} & Q_{\mathbf{x}}^{t,n} & Q_{\theta}^{t,n} \\ Q_{\mathbf{x}}^{t,n} & Q_{\mathbf{x}\mathbf{x}}^{t,n} & Q_{\mathbf{x}\theta}^{t,n} \\ Q_{\theta}^{t,n} & Q_{\theta\mathbf{x}}^{t,n} & Q_{\theta\theta}^{t,n} \end{bmatrix} \begin{bmatrix} \mathbf{1} \\ \delta \mathbf{x}_t \\ \delta \theta_{t,n} \end{bmatrix}, \quad \text{where} \quad \begin{aligned} Q_{\mathbf{x}}^{t,n} &= F_{\mathbf{x}}^t V_{\mathbf{x}}^{t+1,n} \\ Q_{\theta}^{t,n} &= F_{\theta}^t V_{\mathbf{x}}^{t+1,n} + \ell_{\theta}^{t,n} \\ Q_{\theta\theta}^{t,n} &= F_{\theta}^t V_{\mathbf{x}\mathbf{x}}^{t+1,n} F_{\theta}^t + \ell_{\theta\theta}^{t,n} \\ Q_{\mathbf{x}\theta}^{t,n} &= F_{\theta}^t V_{\mathbf{x}\mathbf{x}}^{t+1,n} F_{\mathbf{x}}^t \\ Q_{\mathbf{x}\mathbf{x}}^{t,n} &= F_{\mathbf{x}}^t V_{\mathbf{x}\mathbf{x}}^{t+1,n} F_{\mathbf{x}}^t \end{aligned} \quad (20)$$

follow standard chain rule (recall $Q_{t,n} := \ell_{t,n} + V_{t+1,n} \circ F_t$) with the linearized dynamics $F_\theta^t \equiv \nabla_{\theta_{t,n}} F_t$ and $F_{\mathbf{x}}^t \equiv \nabla_{\mathbf{x}_t} F_t$. The expansion (20) is a standard quadratic programming, and its analytic solution is given by

$$-\delta \pi_{t,n}^* \equiv -\delta \theta_{t,n}^*(\delta \mathbf{x}_t) = -(Q_{\theta\theta}^{t,n})^\dagger (Q_{\theta}^{t,n} + Q_{\theta\mathbf{x}}^{t,n} \delta \mathbf{x}_t) =: -(\mathbf{k}_{t,n} + \mathbf{K}_{t,n} \delta \mathbf{x}_t).$$

Substituting this solution back to the Isaacs-Bellman recursion gives us the local expression of $V_{t,n}$,

$$V_{t,n} \approx Q_{t,n} - \frac{1}{2}(Q_{\theta}^{t,n})^\top (Q_{\theta\theta}^{t,n})^\dagger Q_{\theta}^{t,n}. \quad (21)$$

Therefore, the local derivatives of $V_{t,n}$ can be computed by

$$\begin{aligned} V_x^{t,n} &= Q_x^{t,n} - Q_{x\theta}^{t,n}(Q_{\theta\theta}^{t,n})^\dagger Q_{\theta}^{t,n} = Q_x^{t,n} - Q_{x\theta}^{t,n} \mathbf{k}_{t,n} \\ V_{xx}^{t,n} &= Q_{xx}^{t,n} - Q_{x\theta}^{t,n}(Q_{\theta\theta}^{t,n})^\dagger Q_{\theta x}^{t,n} = Q_{xx}^{t,n} - Q_{x\theta}^{t,n} \mathbf{K}_{t,n}. \end{aligned}$$

Derivation of GR update. We will adopt the same terminology $\mathbf{u} \equiv \theta_{t,1}, \mathbf{v} \equiv \theta_{t,2}$. Following the procedure as in the FNE case, we can perform the second-order expansion of P_t at some fixed point $(\mathbf{x}_t, \mathbf{u}, \mathbf{v})$. The analytic solution to the corresponding quadratic programming is given by

$$-\begin{bmatrix} \delta\pi_{t,1}^* \\ \delta\pi_{t,2}^* \end{bmatrix} = -\begin{bmatrix} P_{uu}^t & P_{uv}^t \\ P_{vu}^t & P_{vv}^t \end{bmatrix}^\dagger \left(\begin{bmatrix} P_u^t \\ P_v^t \end{bmatrix} + \begin{bmatrix} P_{ux}^t \\ P_{vx}^t \end{bmatrix} \delta\mathbf{x}_t \right), \quad (22)$$

where the block-matrices inversion can be expanded using the Schur complement.

$$\begin{bmatrix} P_{uu}^t & P_{uv}^t \\ P_{vu}^t & P_{vv}^t \end{bmatrix}^\dagger = \begin{bmatrix} \overbrace{(P_{uu}^t - P_{uv}^t P_{vv}^{t\dagger} P_{vu}^t)}^{\tilde{P}_{uu}^t} & -\tilde{P}_{uu}^{t\dagger} P_{uv}^t P_{vv}^{t\dagger} \\ -\tilde{P}_{vv}^{t\dagger} P_{vu}^t P_{uu}^{t\dagger} & \underbrace{(P_{vv}^t - P_{vu}^t P_{uu}^{t\dagger} P_{uv}^t)}_{\tilde{P}_{vv}^t} \end{bmatrix}^\dagger. \quad (23)$$

Hence, (22) becomes

$$\begin{aligned} \begin{bmatrix} \delta\pi_{t,1}^* \\ \delta\pi_{t,2}^* \end{bmatrix} &= \begin{bmatrix} \tilde{P}_{uu}^{t\dagger} (P_u^t - P_{uv}^t P_{vv}^{t\dagger} P_v^t) \\ \tilde{P}_{vv}^{t\dagger} (P_v^t - P_{vu}^t P_{uu}^{t\dagger} P_u^t) \end{bmatrix} + \begin{bmatrix} \tilde{P}_{uu}^{t\dagger} (P_{ux}^t - P_{uv}^t P_{vv}^{t\dagger} P_{vx}^t) \\ \tilde{P}_{vv}^{t\dagger} (P_{vx}^t - P_{vu}^t P_{uu}^{t\dagger} P_{ux}^t) \end{bmatrix} \delta\mathbf{x}_t \\ &= \begin{bmatrix} \tilde{P}_{uu}^{t\dagger} (P_u^t - P_{uv}^t \mathbf{I}_t) \\ \tilde{P}_{vv}^{t\dagger} (P_v^t - P_{vu}^t \mathbf{K}_t) \end{bmatrix} + \begin{bmatrix} \tilde{P}_{uu}^{t\dagger} (P_{ux}^t - P_{uv}^t \mathbf{L}_t) \\ \tilde{P}_{vv}^{t\dagger} (P_{vx}^t - P_{vu}^t \mathbf{K}_t) \end{bmatrix} \delta\mathbf{x}_t \\ &=: \begin{bmatrix} \tilde{\mathbf{k}}_t \\ \tilde{\mathbf{I}}_t \end{bmatrix} + \begin{bmatrix} \tilde{\mathbf{K}}_t \\ \tilde{\mathbf{L}}_t \end{bmatrix} \delta\mathbf{x}_t, \end{aligned}$$

where we denote the non-cooperative iterative update for Player 1 and 2 respectively by

$$\begin{aligned} \delta\mathbf{u}_t(\delta\mathbf{x}_t) &= \mathbf{k}_t + \mathbf{K}_t \delta\mathbf{x}_t, \quad \text{where } \mathbf{k}_t := P_{uu}^{t\dagger} P_u^t \quad \text{and} \quad \mathbf{K}_t := P_{uu}^{t\dagger} P_{ux}^t, \\ \delta\mathbf{v}_t(\delta\mathbf{x}_t) &= \mathbf{I}_t + \mathbf{L}_t \delta\mathbf{x}_t, \quad \text{where } \mathbf{I}_t := P_{vv}^{t\dagger} P_v^t \quad \text{and} \quad \mathbf{L}_t := P_{vv}^{t\dagger} P_{vx}^t. \end{aligned}$$

Substituting this solution back to the GR Bellman equation gives the local expression of W_t ,

$$W_t \approx P_t - \frac{1}{2} \begin{bmatrix} P_u^t \\ P_v^t \end{bmatrix}^\top \begin{bmatrix} P_{uu}^t & P_{uv}^t \\ P_{vu}^t & P_{vv}^t \end{bmatrix}^\dagger \begin{bmatrix} P_u^t \\ P_v^t \end{bmatrix}. \quad (24)$$

Finally, taking the derivatives yields the formula for updating the derivatives of W_t ,

$$W_x^t = P_x^t - \frac{1}{2} \left(P_{xu}^t \tilde{\mathbf{k}}_t + P_{xv}^t \tilde{\mathbf{I}}_t + \tilde{\mathbf{K}}_t^\top P_u^t + \tilde{\mathbf{L}}_t^\top P_v^t \right) \quad \text{and} \quad W_{xx}^t = P_{xx}^{t,n} - P_{xu}^t \tilde{\mathbf{K}}_t - P_{xv}^t \tilde{\mathbf{L}}_t, \quad (25)$$

which is much complex than (10).

D.2. Kronecker Factorization and Proof of Proposition 5

We first provide some backgrounds for the Kronecker factorization (KFAC; Martens & Grosse (2015)). KFAC relies on the fact that for an affine mapping layer, *i.e.* $\mathbf{z}_{t+1} = f_t(\mathbf{z}_t, \theta_t) := \mathbf{W}_t \mathbf{z}_t + \mathbf{b}_t$, $\theta_t := \text{vec}([\mathbf{W}_t, \mathbf{b}_t])$, the gradient of the training objective L w.r.t. the parameter θ_t admits a compact factorization,

$$\nabla_{\theta_t} L = \nabla_{\theta_t} f_t^\top \nabla_{\mathbf{z}_{t+1}} L = \mathbf{z}_t \otimes \nabla_{\mathbf{z}_{t+1}} L,$$

where \otimes denotes the Kronecker product. With this, the layer-wise Fisher information matrix, or equivalently the Gauss-Newton (GN) matrix, for classification can be approximated with

$$\mathbb{E}[\nabla_{\theta_t} L \nabla_{\theta_t} L^\top] = \mathbb{E}[(z_t \otimes \nabla_{z_{t+1}} L)(z_t \otimes \nabla_{z_{t+1}} L)^\top] \approx \mathbb{E}[z_t z_t^\top] \otimes \mathbb{E}[\nabla_{z_{t+1}} L \nabla_{z_{t+1}} L^\top].$$

We can adopt this factorization to our setup by first recalling from our proof of Proposition 2 (see Appendix C) that $(\nabla_{\theta_t} L, \nabla_{z_{t+1}} L)$ are interchangeable with (H_θ^t, p_{t+1}) , or equivalently (H_θ^t, H_z^{t+1}) . Hence, the GN approximation of $\mathbb{E}[H_{\theta\theta}^t]$ can be factorized by

$$\mathbb{E}[H_\theta^t H_\theta^{t\top}] \approx \mathbb{E}[z_t z_t^\top] \otimes \mathbb{E}[p_{t+1} p_{t+1}^\top] = \mathbb{E}[z_t z_t^\top] \otimes \mathbb{E}[H_z^{t+1} H_z^{t+1\top}]. \quad (26)$$

Equation (26) suggests that KFAC factorizes the parameter curvature with two smaller matrices using the activation state z_t and the derivative of *some optimality* (in this case the Hamiltonian H) w.r.t. z_{t+1} . The main advantage of this factorization is to exploit the following formula,

$$(A \otimes B)^\dagger \text{vec}(W) = (A^\dagger \otimes B^\dagger) \text{vec}(W) = \text{vec}(B^\dagger W A^{-\top}), \quad (27)$$

which allows one to efficiently inverse the parameter curvature with two smaller matrices.

Now, let us proceed to the proof of Proposition 5. First notice that for the shared dynamics considered in Fig. 2, we have

$$F_t(x_t, u, v) := \begin{bmatrix} f_{t,1}(z_1, u) \\ f_{t,2}(z_2, v) \end{bmatrix} = \begin{bmatrix} f_{t,1}(\cdot, u) & \mathbf{0} \\ \mathbf{0} & f_{t,2}(\cdot, v) \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix},$$

which resembles the affine mapping concerned by KFAC. This motivates the following approximation,

$$\mathbb{E}[P_\theta^t P_\theta^{t\top}] \approx \mathbb{E}[x_t x_t^\top] \otimes \mathbb{E}[W_x^{t+1} W_x^{t+1\top}]. \quad (28)$$

Similar to (26), this approximation (28) factorizes the GN matrix with the MPDG state x_t and the derivative of an optimality (in this case it becomes the GR value function W_{t+1}) w.r.t. x_{t+1} .

If we denote the derivatives w.r.t. the outputs of $f_{t,1}$ and $f_{t,2}$ by g_1 and g_2 , i.e. $W_x^{t+1} := [g_1, g_2]^\top$, and rewrite $x_t := [z_1, z_2]^\top$, then (28) can be expanded by

$$\begin{aligned} \mathbb{E}[x_t x_t^\top] &= \begin{bmatrix} \mathbb{E}[z_1 z_1^\top] & \mathbb{E}[z_1 z_2^\top] \\ \mathbb{E}[z_2 z_1^\top] & \mathbb{E}[z_2 z_2^\top] \end{bmatrix} =: \begin{bmatrix} A_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix} \\ \mathbb{E}[W_x^{t+1} W_x^{t+1\top}] &= \begin{bmatrix} \mathbb{E}[g_1 g_1^\top] & \mathbb{E}[g_1 g_2^\top] \\ \mathbb{E}[g_2 g_1^\top] & \mathbb{E}[g_2 g_2^\top] \end{bmatrix} =: \begin{bmatrix} B_{uu} & B_{uv} \\ B_{vu} & B_{vv} \end{bmatrix}. \end{aligned}$$

Their inverse matrices are given by the Schur component.

$$\begin{aligned} \begin{bmatrix} A_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix}^\dagger &= \begin{bmatrix} \tilde{A}_{uu}^\dagger & -\tilde{A}_{uu}^\dagger A_{uv} A_{vv}^\dagger \\ -\tilde{A}_{vv}^\dagger A_{vu} A_{uu}^\dagger & \tilde{A}_{vv}^\dagger \end{bmatrix}, \quad \text{where} \begin{cases} \tilde{A}_{uu} := A_{uu} - A_{uv} A_{vv}^\dagger A_{uv}^\top \\ \tilde{A}_{vv} := A_{vv} - A_{vu} A_{uu}^\dagger A_{vu}^\top \end{cases} \\ \begin{bmatrix} B_{uu} & B_{uv} \\ B_{vu} & B_{vv} \end{bmatrix}^\dagger &= \begin{bmatrix} \tilde{B}_{uu}^\dagger & -\tilde{B}_{uu}^\dagger B_{uv} B_{vv}^\dagger \\ -\tilde{B}_{vv}^\dagger B_{vu} B_{uu}^\dagger & \tilde{B}_{vv}^\dagger \end{bmatrix}, \quad \text{where} \begin{cases} \tilde{B}_{uu} := B_{uu} - B_{uv} B_{vv}^\dagger B_{uv}^\top \\ \tilde{B}_{vv} := B_{vv} - B_{vu} B_{uu}^\dagger B_{vu}^\top \end{cases} \end{aligned} \quad (29)$$

With all these, the cooperative open gain can be computed with the formula (27),

$$\left(\mathbb{E}[x_t x_t^\top] \otimes \mathbb{E}[W_x^{t+1} W_x^{t+1\top}] \right)^\dagger \text{vec} \left(\begin{bmatrix} P_u^t & \mathbf{0} \\ \mathbf{0} & P_v^t \end{bmatrix} \right) = \text{vec} \left(\begin{bmatrix} B_{uu} & B_{uv} \\ B_{vu} & B_{vv} \end{bmatrix}^\dagger \begin{bmatrix} P_u^t & \mathbf{0} \\ \mathbf{0} & P_v^t \end{bmatrix} \begin{bmatrix} A_{uu} & A_{uv} \\ A_{vu} & A_{vv} \end{bmatrix}^{-\top} \right). \quad (30)$$

Substituting (29) into (30), after some algebra we will arrive at the KFAC of the cooperative open gain suggested in (15).

$$\begin{aligned} \tilde{\mathbf{k}}_t &\approx \text{vec}(\tilde{B}_{uu}^\dagger P_u^t \tilde{A}_{uu}^{-\top} + \tilde{B}_{uu}^\dagger B_{uv} B_{vv}^\dagger P_v^t (\tilde{A}_{uu}^\dagger A_{uv} A_{vv}^\dagger)^\top) \\ &= \text{vec}(\tilde{B}_{uu}^\dagger (P_u^t + B_{uv} B_{vv}^\dagger P_v^t A_{vv}^{-\top} A_{uv}^\top) \tilde{A}_{uu}^{-\top}) \\ &= \text{vec}(\tilde{B}_{uu}^\dagger (P_u^t + B_{uv} \text{vec}^\dagger(\mathbf{I}_t) A_{uv}^\top) \tilde{A}_{uu}^{-\top}), \end{aligned} \quad (31)$$

where the last equality follows by another KFAC approximation $\mathbf{I}_t \approx (A_{vv} \otimes B_{vv})^\dagger \text{vec}(P_v^t) = \text{vec}(B_{vv}^\dagger P_v^t A_{vv}^{-\top})$. Finally, recalling the expression, $\tilde{\mathbf{k}}_t := \tilde{P}_{uu}^\dagger (P_u^t - P_{uv}^t \mathbf{I}_t)$, from (11) and rewriting (31) by

$$\begin{aligned}\tilde{\mathbf{k}}_t &\approx \text{vec}(\tilde{B}_{uu}^\dagger (P_u^t + B_{uv} \text{vec}^\dagger(\mathbf{I}_t) A_{uv}^\top) \tilde{A}_{uu}^{-\top}) \\ &= (\tilde{A}_{uu} \otimes \tilde{B}_{uu})^\dagger \text{vec}(P_u^t + B_{uv} \text{vec}^\dagger(\mathbf{I}_t) A_{uv}^\top)\end{aligned}$$

imply the KFAC representation $\tilde{P}_{uu}^t \approx \tilde{A}_{uu} \otimes \tilde{B}_{uu}$ in (14). Hence we conclude the proof. \square

D.3. Proof of Theorem 6

We first show that setting $P_{uv}^t := \mathbf{0}$ in the update (11) yields (9). To begin, observe that when P_{uv}^t vanishes, the cooperative gains $(\tilde{\mathbf{k}}_t, \tilde{\mathbf{K}}_t)$ appearing in (11) degenerate to $\tilde{\mathbf{k}}_t = P_{uu}^\dagger P_u^t$ and $\tilde{\mathbf{K}}_t = P_{uu}^\dagger P_{ux}^t$. Therefore, it is sufficient to prove the following result.⁴

Lemma 9. *Suppose $Q_{t,n}$ in (6) and P_t in (7) are expanded up to second-order along the same local trajectory $\{(\mathbf{x}_t, \theta_{t,n}, \dots, \theta_{t,N}) : \forall t \in [T]\}$, then we will have the following relations when $P_{uv}^t := \mathbf{0}$ at all stages.*

$$\forall t, \quad Q_{\theta^1}^t = P_u^t, \quad Q_{\theta\theta}^t = P_{uu}^t, \quad Q_{\theta x}^t = P_{ux}^t, \quad Q_{\theta^2}^t = P_v^t, \quad Q_{\theta\theta}^{t,2} = P_{vv}^t, \quad Q_{\theta x}^{t,2} = P_{vx}^t, \quad (32)$$

where $(\mathbf{u}_t, \mathbf{v}_t) \equiv (\theta_{t,1}, \theta_{t,2})$ denotes the actions for Player 1 and 2. Furthermore, we have

$$\forall t, \quad W_t = \sum_{n=1}^N V_{t,n}. \quad (33)$$

Proof. We will proceed the proof by induction. At the terminal stage $T-1$, we have

$$P_{T-1} = \sum_{n=1}^2 \ell_{T-1,n} + W_T \circ F_{T-1} = \sum_{n=1}^2 (\ell_{T-1,n} + \phi_n \circ F_{T-1}) = \sum_{n=1}^2 Q_{T-1,n},$$

since $\phi_n = V_{T,n}$. This implies that when solving the second-order expansion for $\pi_{T-1,1}$ and $\pi_{T-1,2}$, we will have

$$\min_{\pi_{T-1,1}, \pi_{T-1,2}} P_{T-1} = \min_{\pi_{T-1,1}} Q_{T-1,1} + \min_{\pi_{T-1,2}} Q_{T-1,2}$$

since the cross-correlation matrix P_{uv}^{T-1} is discarded. Therefore, all equalities in (32) hold at this stage. Furthermore, substituting $P_{uv}^{T-1} := \mathbf{0}$ into (24) yields the following GR value function

$$\begin{aligned}W_{T-1} &= P_{T-1} - \frac{1}{2} ((P_u^{T-1})^\top (P_{uu}^{T-1})^\dagger P_u^{T-1} + (P_v^{T-1})^\top (P_{vv}^{T-1})^\dagger P_v^{T-1}) \\ &= \sum_{n=1}^2 \left(Q_{T-1,n} - \frac{1}{2} (Q_{\theta}^{T-1,n})^\top (Q_{\theta\theta}^{T-1,n})^\dagger Q_{\theta}^{T-1,n} \right) = \sum_{n=1}^2 V_{T-1,n}.\end{aligned}$$

So (33) also holds. Now, suppose (32, 33) hold at $t+1$, then

$$P_t = \sum_{n=1}^2 \ell_{t,n} + W_{t+1} \circ F_t = \sum_{n=1}^2 (\ell_{t,n} + V_{t+1,n} \circ F_t) = \sum_{n=1}^2 Q_{t,n}.$$

Together with $P_{uv}^t := \mathbf{0}$, we can see that all equalities in (32) hold. Furthermore, it implies that

$$W_t = P_t - \frac{1}{2} (P_u^\top P_{uu}^\dagger P_u + P_v^\top P_{vv}^\dagger P_v) = \sum_{n=1}^2 \left(Q_{t,n} - \frac{1}{2} (Q_{\theta}^{t,n})^\top (Q_{\theta\theta}^{t,n})^\dagger Q_{\theta}^{t,n} \right) = \sum_{n=1}^2 V_{t,n}.$$

Hence we conclude the proof. \square

Next, we proceed to the second case, which suggests that running (9) with $(Q_{\theta x}^t, Q_{\theta\theta}^t) := (\mathbf{0}, \mathbf{I})$ yields SGD. Since the FNE update in this case degenerates to $\delta\pi_{t,n}^* = Q_{\theta}^{t,n}$, it is sufficient to prove the following lemma.

⁴ We consider the two-player setup for simplicity, yet the methodology applies to the multi-player setup.

Lemma 10. Suppose $H_{t,n}$ in (5) and $Q_{t,n}$ in (6) are expanded up to second-order along the same local trajectory $\{(x_t, \theta_{t,n}, \dots, \theta_{t,N}) : \forall t \in [T]\}$, then we will have the following relations when $(Q_{\theta x}^{t,n}, Q_{\theta\theta}^{t,n}) := (\mathbf{0}, \mathbf{I})$ for all stages.

$$\forall t, \quad Q_{\theta}^{t,n} = H_{\theta}^{t,n}, \quad V_{\mathbf{x}}^{t,n} = \mathbf{p}_{t,n}. \quad (34)$$

Proof. Again, we will proceed the proof by induction. First, notice that $V_{\mathbf{x}}^{T,n} = \phi_{\mathbf{x}}^n = \mathbf{p}_{T,n}$ no matter whether or not $Q_{\theta x}^{t,n}$ and $Q_{\theta\theta}^{t,n}$ degenerate. At the terminal stage $T-1$, we have

$$Q_{\theta}^{T-1,n} = \ell_{\theta}^{T-1,n} + (F_{\theta}^{T-1})^{\top} V_{\mathbf{x}}^{T,n} = \ell_{\theta}^{T-1,n} + (F_{\theta}^{T-1})^{\top} \mathbf{p}_{T,n} = H_{\theta}^{T-1,n}.$$

Also, when $Q_{\theta x}^{T-1,n} := \mathbf{0}$, (10) becomes

$$V_{\mathbf{x}}^{T-1,n} = Q_{\mathbf{x}}^{T-1,n} = (F_{\mathbf{x}}^{T-1})^{\top} V_{\mathbf{x}}^{T,n} = (F_{\mathbf{x}}^{T-1})^{\top} \mathbf{p}_{T,n} = H_{\mathbf{x}}^{T-1,n} = \mathbf{p}_{T-1,n}.$$

Hence, (34) holds at $T-1$. Now, suppose these relations hold at $t+1$, then

$$Q_{\theta}^{t,n} = \ell_{\theta}^{t,n} + (F_{\theta}^t)^{\top} V_{\mathbf{x}}^{t+1,n} = \ell_{\theta}^{t,n} + (F_{\theta}^t)^{\top} \mathbf{p}_{t+1,n} = H_{\theta}^{t,n}$$

and similarly

$$V_{\mathbf{x}}^{t,n} = Q_{\mathbf{x}}^{t,n} = (F_{\mathbf{x}}^t)^{\top} V_{\mathbf{x}}^{t+1,n} = (F_{\mathbf{x}}^t)^{\top} \mathbf{p}_{t+1,n} = H_{\mathbf{x}}^{t,n} = \mathbf{p}_{t,n}.$$

Hence, we conclude the proof. \square

Finally, the last case follows readily by combining Lemma 9 and 10, so we conclude all proofs. \square

E. More on the Experiments

All experiments are run with Pytorch on the GPU machines, including GTX 1080 TI, GTX 2070, and TITAN RTX. We preprocessed all datasets with standardization. We also perform data augmentation when training CIFAR100. Below we detail the setup for each experiment.

Classification (Table 2 and 3). For CIFAR10 and CIFAR100, we use standard implementation of ResNet18 from https://pytorch.org/hub/pytorch_vision_resnet/. As for SVHN and MNIST, the residual network consists of 3 residual blocks. The residual block shares a similar architecture in Fig. 2 except with the identity shortcut mapping and without BN. We use 3×3 kernels for all convolution filters. The number of feature maps in the convolution filters is set to 12 and 16 respectively for MNIST and SVHN. Meanwhile, the inception network consists of a convolution layer followed by an inception block (see Fig. 6), another convolution layer, and two fully-connected layers. Regarding the hyper-parameters used in baselines, we select them from an appropriate search space detailed in Table 9. We use the implementation in <https://github.com/Thrandis/EKFAC-pytorch> for EKFAC and implement our own EMSA in PyTorch since the official code released from Li et al. (2017a) does not support GPU parallelization.

Table 9. Hyper-parameter search in Table 2

Standard Baselines	Learning Rate (LR)
SGD	(7e-2, 5e-1)
Adam & RMSprop	(7e-4, 1e-2)
EKFAC	(1e-2, 3e-1)

Ablation study (Fig. 5) Each grid in Fig. 5 corresponds to a distinct combination of baseline and dataset. Its numerical value reports the performance difference between the following two training processes.

- Accuracy of the baseline run with the best-tuned configuration which we report in Table 2 and 3.
- Accuracy of DGNOpt with its parameter curvature set to the precondition matrix implied by the above best-tuned setup.

For instance, suppose the learning rate of EKFAC on MNIST is best-tuned to 0.01, then we simply set $Q_{\theta\theta}^{t,n} \approx 0.01 \times Q_{\theta}^{t,n} Q_{\theta}^{t,n\top}$ for all t . From Theorem 6, these two training procedures only differ in the presence of $Q_{\theta x}^{t,n}$, which allows EKFAC to adjust its update based on the change of $\mathbf{x}_t \in \eta_{t,n}^C$.

Runtime and memory complexity (Fig. 7). The numerical values are measured on the GTX 2070.

Feedback analysis (Fig. 8). We use the same inception-based network in Table 3.

Remark for EMSA (Footnote 3). Extended Method of Successive Approximations (EMSA) was originally proposed by Li et al. (2017a) as an OCP-inspired method for training *feedforward* networks. It considers the following minimization,

$$\theta_t^* = \arg \min H_t^p(z_t, z_{t+1}, p_t, p_{t+1}, \theta_t),$$

$$\text{where } H_t^p(z_t, z_{t+1}, p_t, p_{t+1}, \theta_t) := H_t(z_t, p_{t+1}, \theta_t) + \frac{1}{2}\rho \|z_{t+1} - f_t(z_t, \theta_t)\|_2 + \frac{1}{2}\rho \|p_t - \nabla_{z_t} H_t\|_2 \quad (35)$$

essentially augments the original Hamiltonian H_t with the feasibility constraints on both forward states and backward co-states. EMSA solves the minimization (35) with L-BFGS per layer at each training iteration. In Table 2 and 3, we extend their formula to accept $H_{t,n}$. Due to the feasibility constraints, the resulting modified Hamiltonian $H_{t,n}^p$ depends additionally on x_{t+1} and $p_{t,n}$; hence being different from the original Hamiltonian $H_{t,n}$. As a result, the ablation analysis using Theorem 6 is not applicable for EMSA.

Cooperative training (Fig. 9, Fig. 11, and Table 4). The network consists of 4 convolutions followed by 2 fully-connected layers, and is activated by ReLU. We use 3×3 kernels with 32 feature maps for all convolutions and set the batch size to 128.

Adaptive alignment with bandit (Fig. 10 and Table 5). We use the same ResNet18 as in classification for CIFAR10, and a smaller residual network with 1 residual block for SVHN. The residual block shares the same architecture as in Fig. 2 except without BN. All convolution layers use 3×3 kernels with 12 feature maps. Again, the batch size is set to 128. Note that in this experiment we use a slightly larger learning rate compared with the one used in Table 2. While DGNOpt achieves better final accuracies for both setups, in practice, the former tends to amplify the stabilization when we enlarge the information structure during training. Hence, it differentiates DGNOpt from other baselines.

Alg. 2 presents the pseudo-code of how DGNOpt can be integrated with any generic bandit-based algorithm (marked as blue). For completeness, we also provide the pseudo-code of EXP3++ in Alg. 3. We refer readers to Seldin & Slivkins (2014) for the definition of $\xi_k(m)$ and η_k (do not confuse with $\eta_{t,n}$ in the main context).

Algorithm 2 DGNOpt with Multi-Armed Bandit (MAB)

Input: dataset \mathcal{D} , network $\{f_i(\cdot, \theta_i)\}$, number of alignments M
 Initialize the multi-armed bandit `MAB.init(M)`
repeat
 Draw an alignment $m \leftarrow \text{MAB.sample}()$.
 Construct $F \equiv \{F_t : t \in [T]\}$ according to m .
 Compute x_t by propagating $x_0 \sim \mathcal{D}$ through F
 for $t = T-1$ **to** 0 **do** ▷ Solve FNE or GR
 Solve the update $\delta\pi_{t,n}^*$ with (9) or (11)
 Solve $(V_{\mathbf{x}}^{t,n}, V_{\mathbf{x}\mathbf{x}}^{t,n})$ or $(W_{\mathbf{x}}^t, W_{\mathbf{x}\mathbf{x}}^t)$ with (10) or (25)
 end for
 Set $x'_0 = x_0$
 for $t = 0$ **to** $T-1$ **do** ▷ Update parameter
 Apply $\theta_{t,n} \leftarrow \theta_{t,n} - \delta\pi_{t,n}^*(\delta x_t)$ with $\delta x_t = x'_t - x_t$
 Compute $x'_{t+1} = F_t(x'_t, \theta_{t,1}, \dots, \theta_{t,N})$
 end for
 Compute the accuracy r on validation set.
 Run `MAB.update(r)`.
until converges

Algorithm 3 EXP3++ (Seldin & Slivkins, 2014)

function `init(M)`
 $(k, M) \leftarrow (1, M)$
 $\forall m, L_k(m) = 0$
end function

function `sample()`
 $\forall m, \epsilon_k(m) = \min\{\frac{1}{2M}, \frac{1}{2}\sqrt{\frac{\ln M}{kM}}, \xi_k(m)\}$
 $\forall m, \rho_k(m) = e^{-\eta_k L_k(m)} / \sum_{m'} e^{-\eta_k L_k(m')}$
 $\forall m, \tilde{\rho}_k(m) = (1 - \sum_{m'} \epsilon_k(m'))\rho_k(m) + \epsilon_k(m)$
 Sample action according to $\tilde{\rho}_k(m)$
end function

function `update(r_k^m)`
 $\ell_k^m = (1 - r_k^m) / \tilde{\rho}_k(m)$
 $L_{k+1}(m) = L_k(m) + \ell_k^m$
 $k \leftarrow k + 1$
end function


Additional Experiments.


Figure 13. (a) Training curve and (b) final accuracy as we vary the number of player (N) as a hyper-parameter of game-extended EKFAC. Similar to Fig. 9, we also observe that $N=2$ gives the best final accuracy on both datasets.