A new perspective and contextualisation of the group equivariant CNN

Expanding the framework of G-CNNs

Oscar Carlsson*, Elias Nyholm, Daniel Persson

Mackey functions and $\mathcal{I}_k$ (the induced representation)

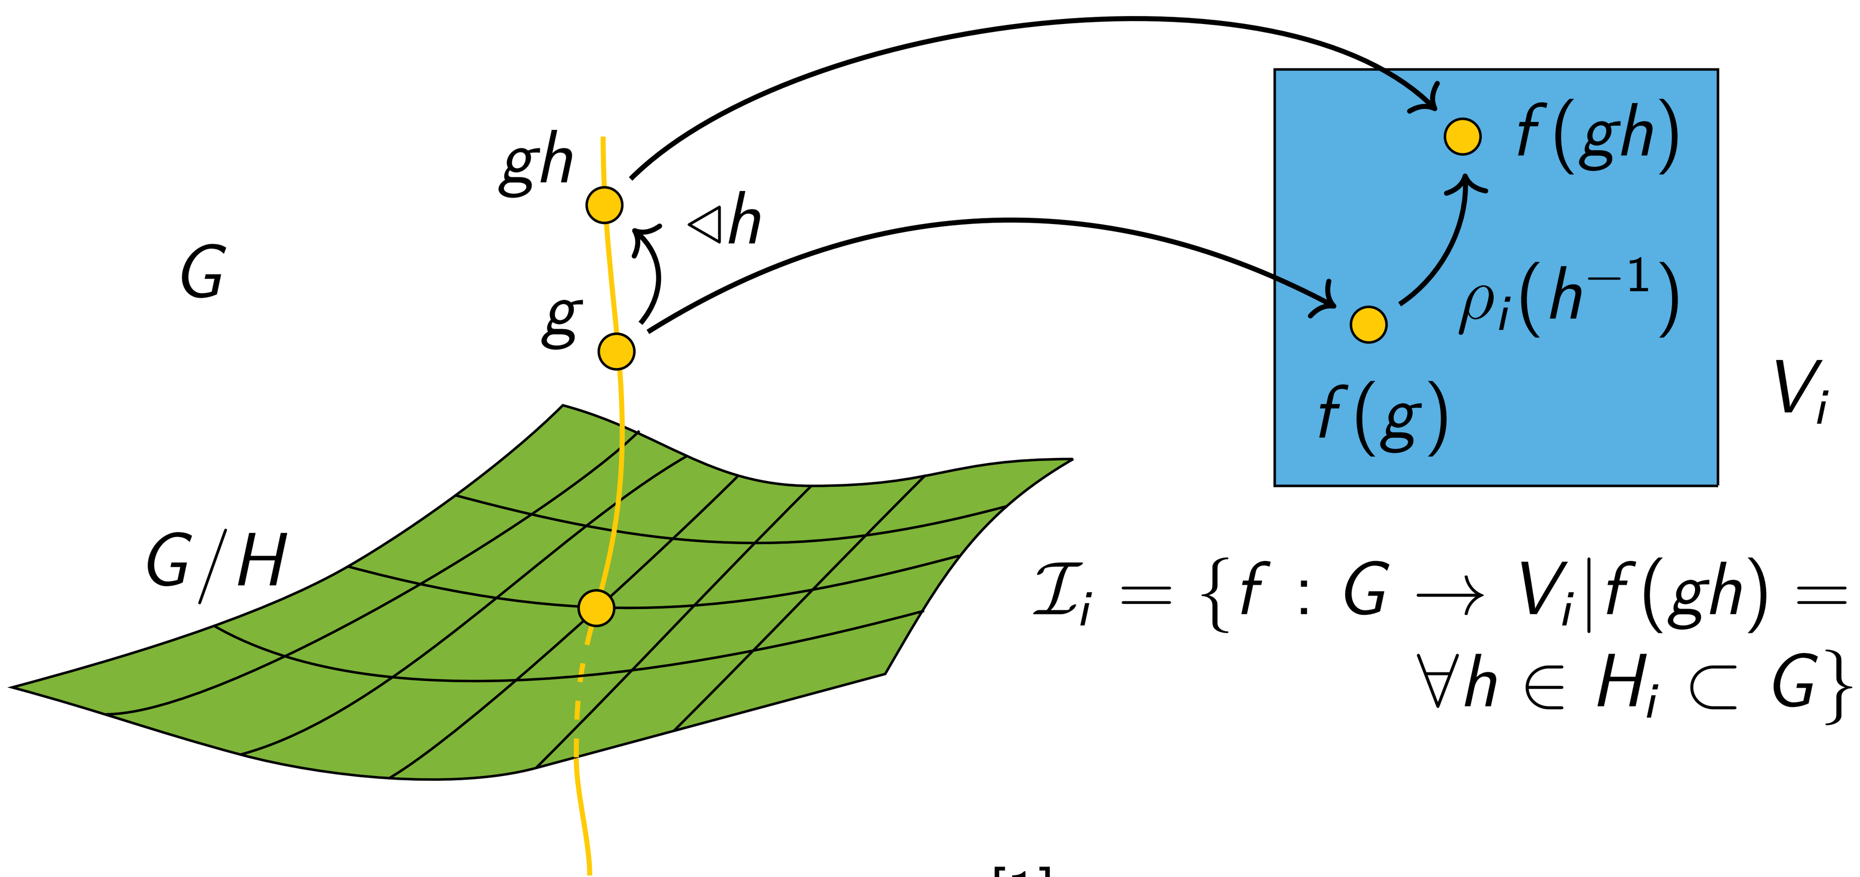

$$\mathcal{I}_i = \{f : G \to V_i | f(gh) = \rho_i(h^{-1})f(g) \quad \forall h \in H_i \subset G\}$$

Group equivariant CNN[1]

The group equivariant CNN

$$\phi : \mathcal{I}_1 \to \mathcal{I}_2, \quad [\phi f](g) = \int_G \widehat{\kappa}(g^{-1}g')f(g')dg', \quad \underbrace{k[\phi f](g) = [\phi(kf)](g)}_{G-\text{equivariance of } \phi}$$

is a special case of the, not necessarily equivariant, general linear map

$$\psi : L^2(G; V_1) \to L^2(G; V_2) \quad \text{by} \quad [\psi f](g) = \int_G \underbrace{\kappa(g, g')}_{\text{matrix}} \underbrace{f(g')}_{\text{vector}} dg'.$$

General case

The general linear map is a special case of the (non-linear) map

$$\psi : L^2(G; V_1) \to L^2(G; V_2) \quad \text{by} \quad [\psi f](g) = \int_G \omega(f, g, g')dg',$$

where $\omega(f, g, g') = \kappa(g, g')f(g')$. See Elias Nyholm's poster for more details.

Equivariance vs the Mackey constraint

Mathematically, the equivariance of ϕ and the Mackey properties of the input and output can be imposed independently in the general linear map.

Kernel domain

The domain of the integration kernel is inherently connected to the properties of the integral transform. Certain properties on the map can be enforced by choosing an appropriate projection and defining the kernel on the new space. The same properties can be obtained by restricting the space of kernels defined on the larger spaces.

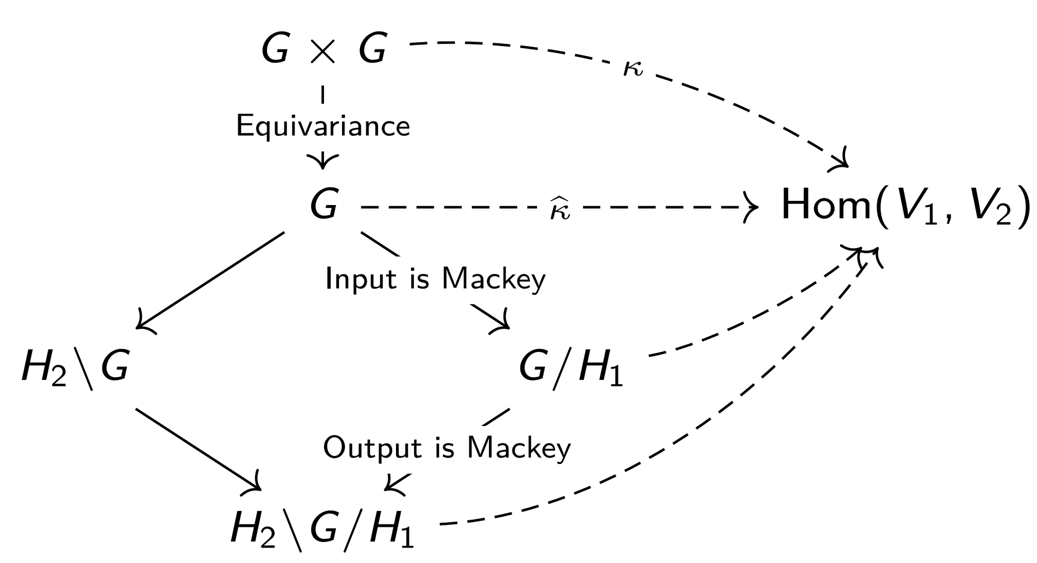

Functions and sections

There is a natural connection between the Mackey functions and sections of the corresponding associated vector bundle.

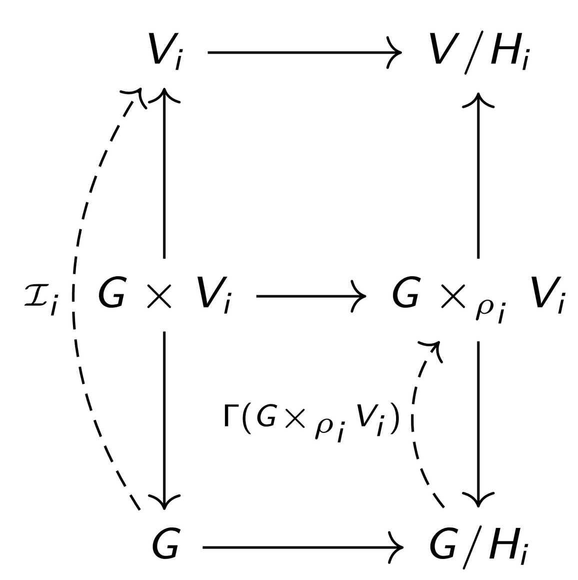

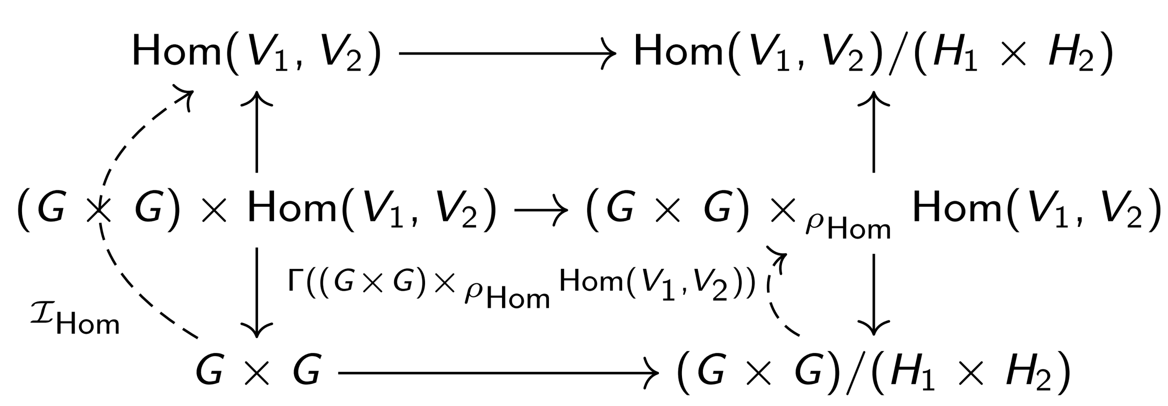

Email: osccarls@chalmers.se

Work in progress.

References

[1] Cohen, Taco et al. *A General Theory of Equivariant CNNs on Homogeneous Spaces*, NeurIPS 2019