

Lens Semantics of Machine Learning


Fabio Zanasi
University College London

Based on joint work with
Geoff Cruttwell, Bruno Gavranovic, Neil Ghani, Paul Wilson


IFIP WG1.3 Meeting
January 2022

String Diagrams in Computer Science


Neural networks




Bayesian networks



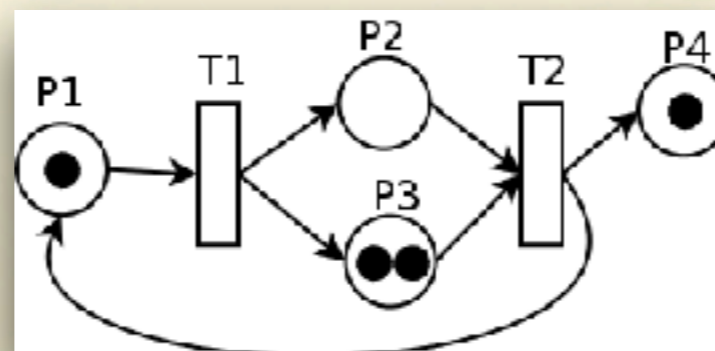
Quantum circuits



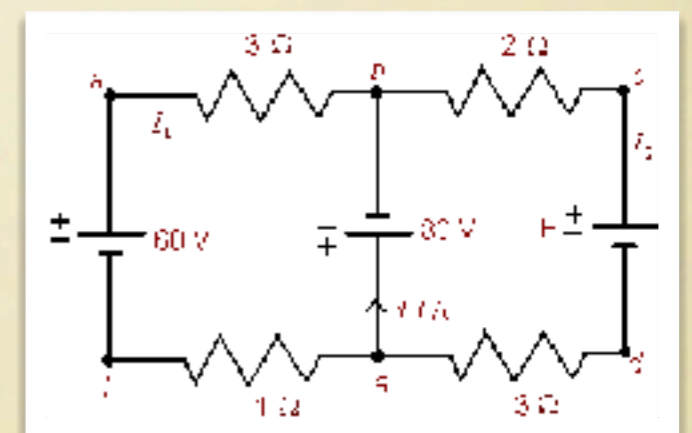
Signal flow graphs



Petri nets






Electrical Circuits



Key Ideas

Graphical notation becomes **formal syntax**
and it is given **compositional** semantics

Example:
Signal Flow Graphs



↓
?

↓
?

↓
 R_1

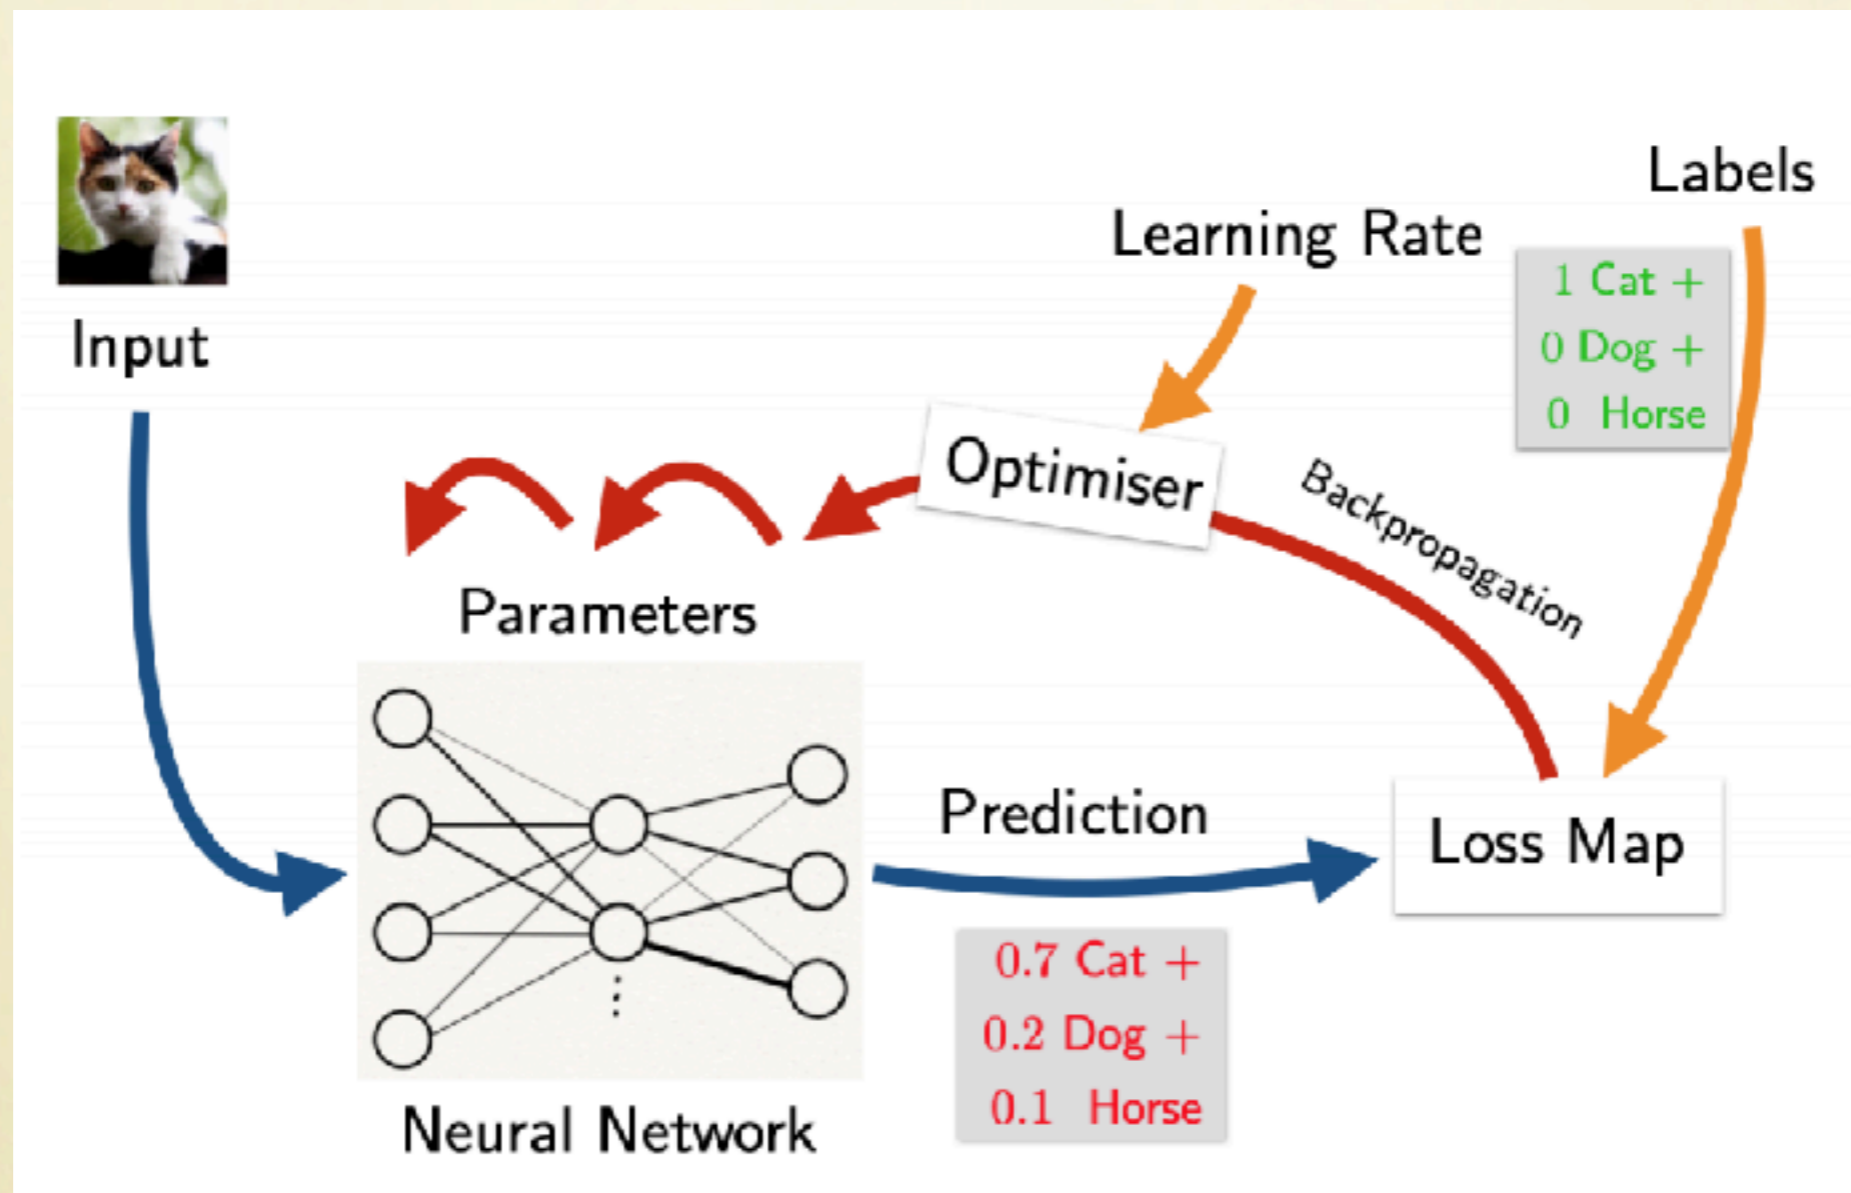
;

↓
 R_2

=

R

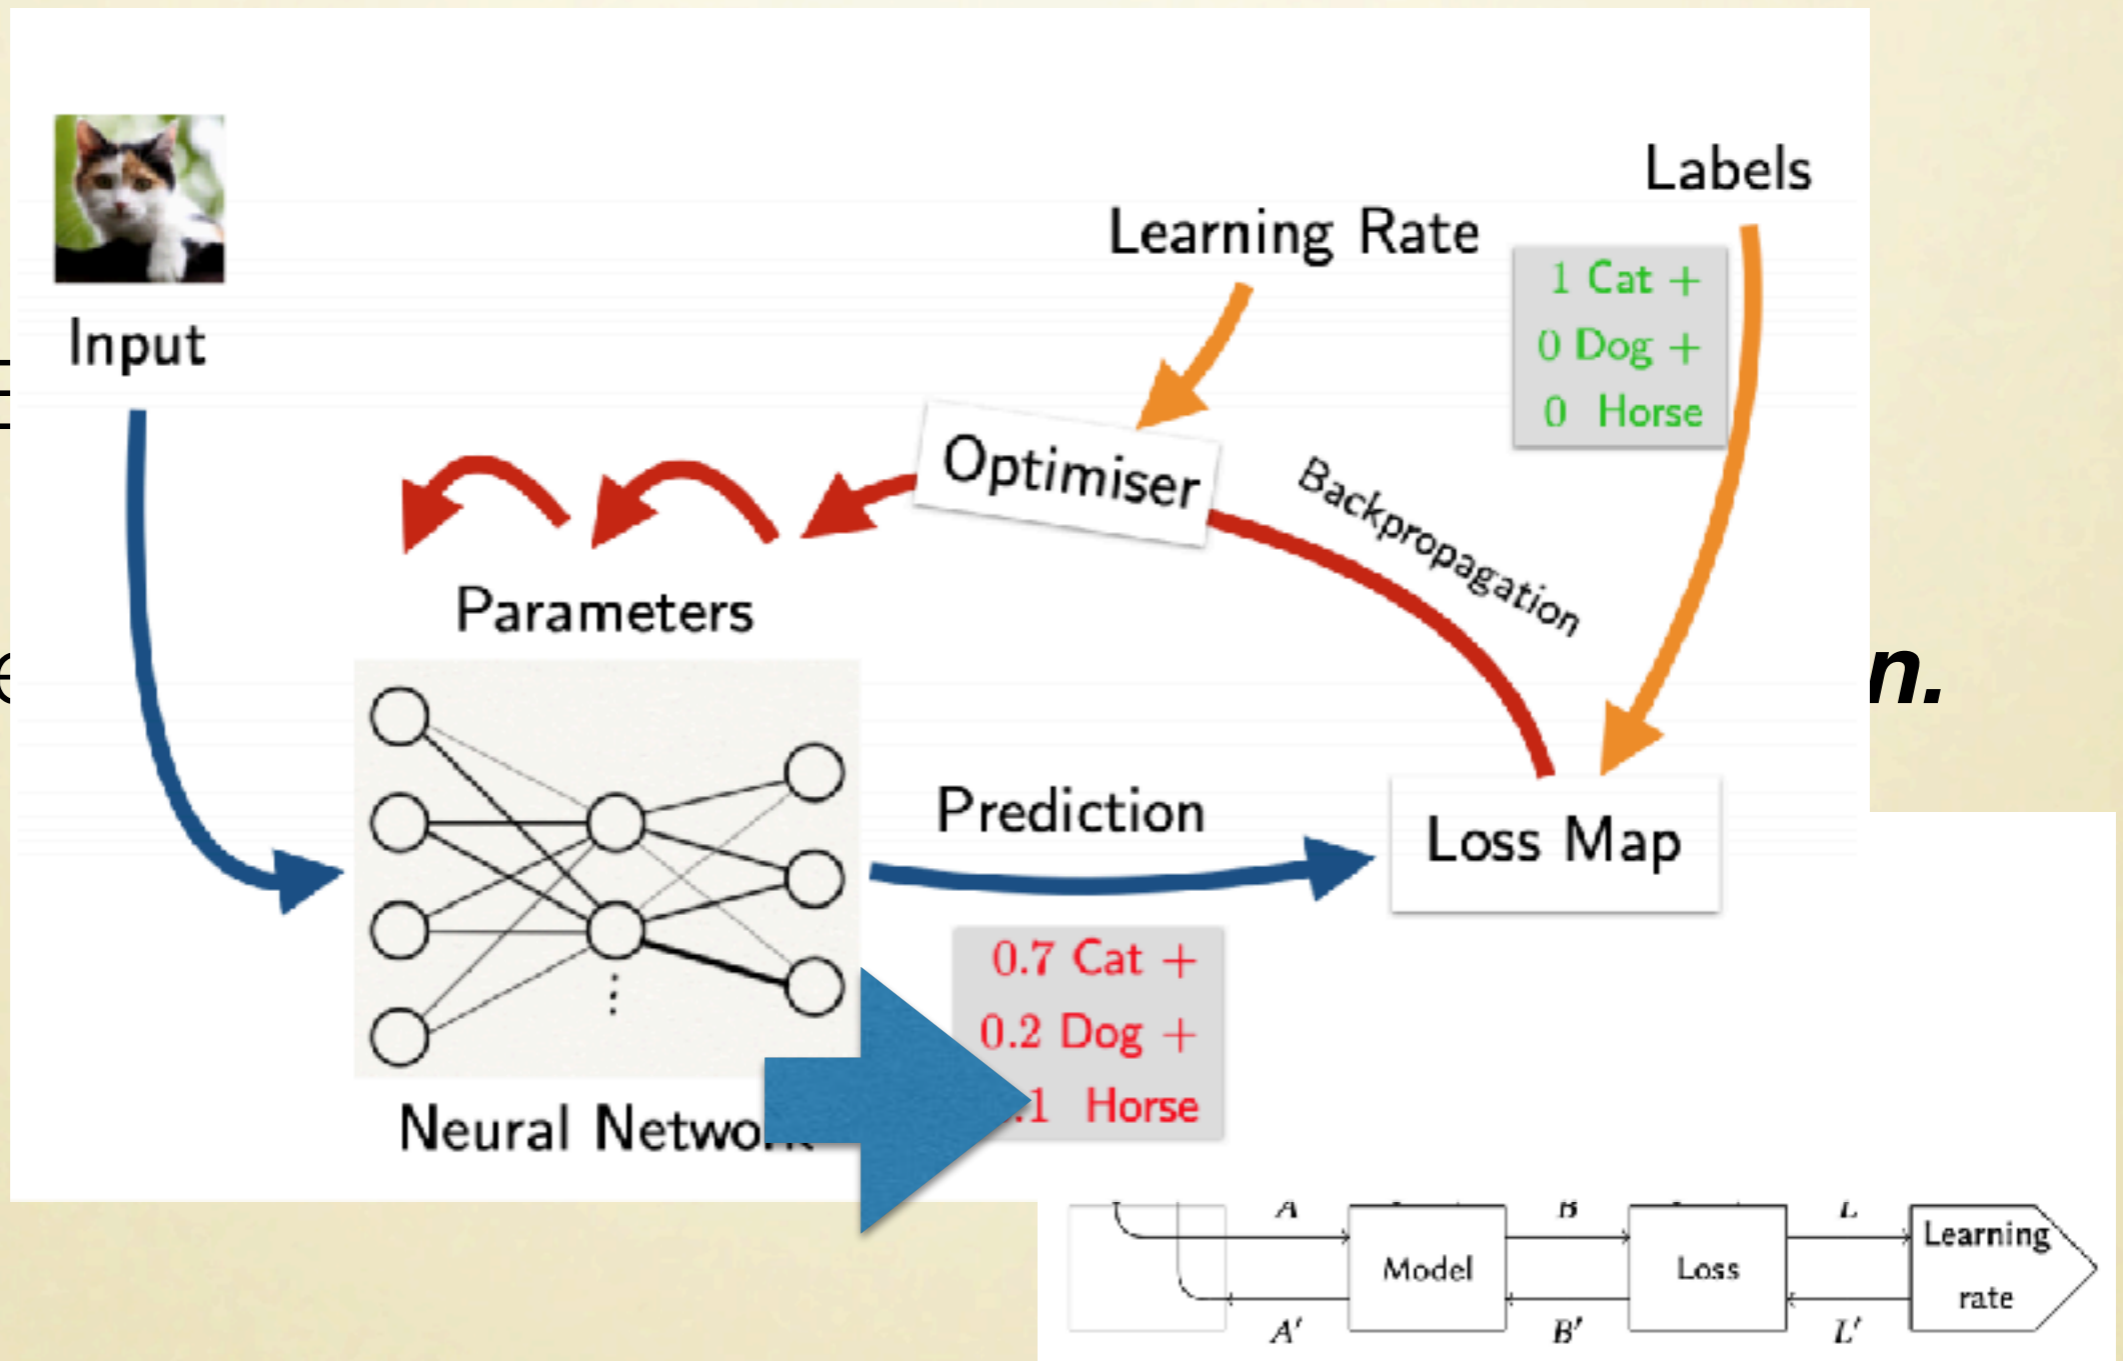
Case study: Gradient-Based Learning





What are the fundamental semantic structures underpinning (gradient-based) learning?

Roadmap

Plan:



Parametric Maps





Para(C)

C

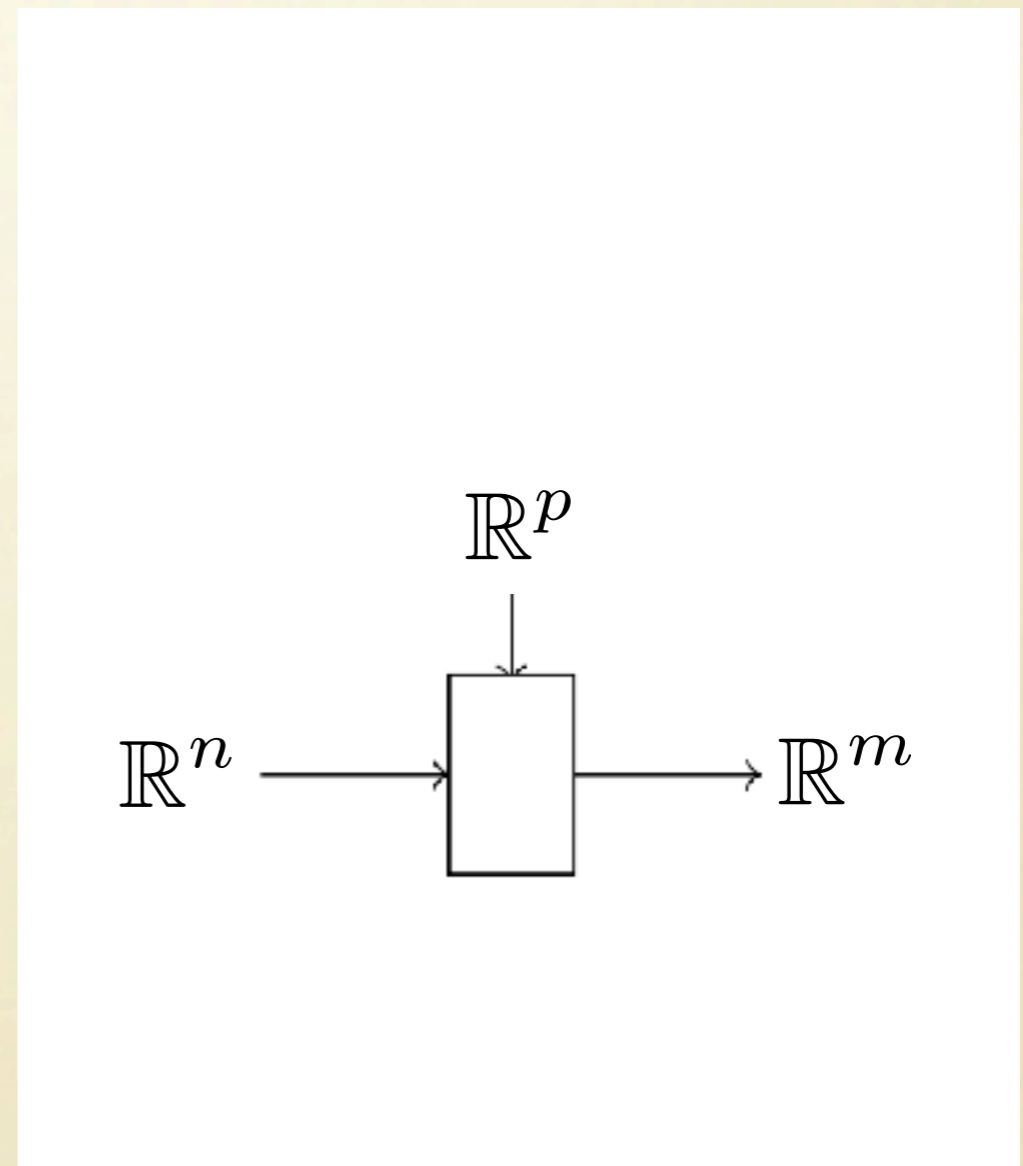
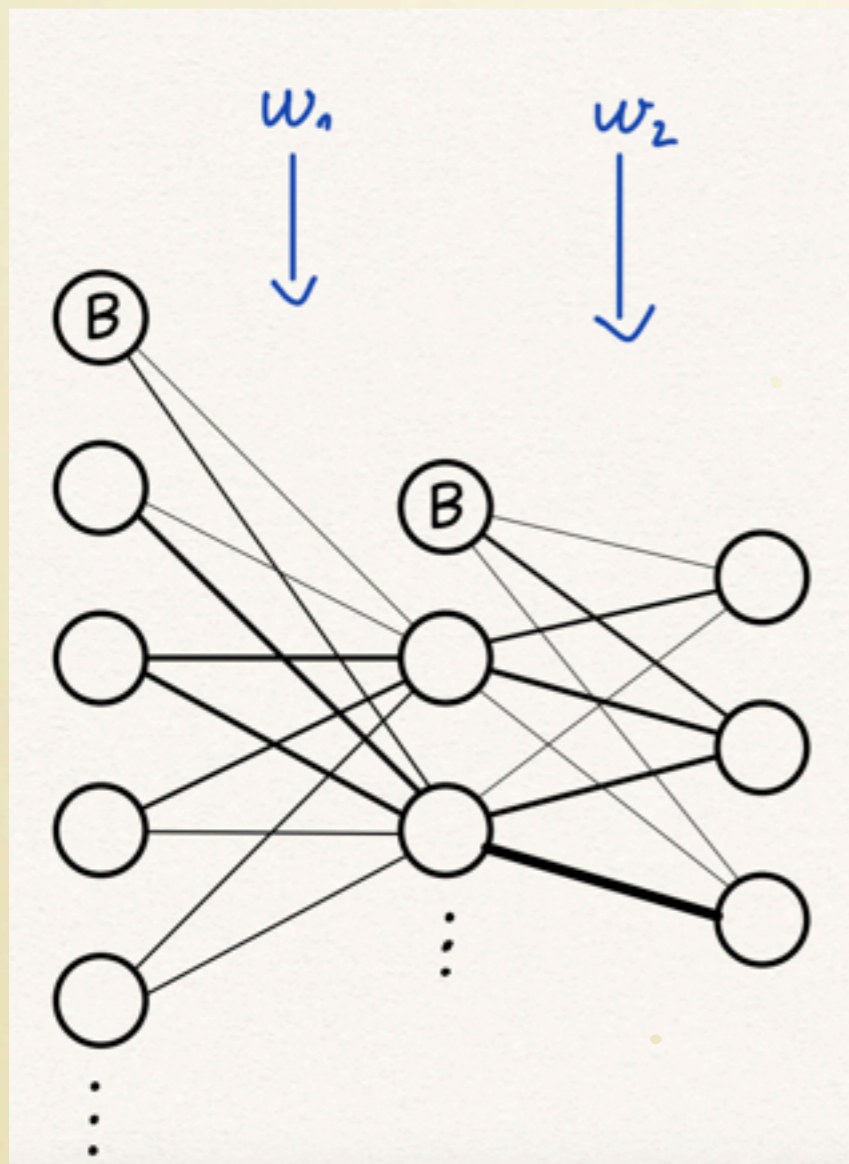
$$\frac{f : A \rightarrow B}{f : P \times A \rightarrow B}$$

Composing Parametric Maps





Parametric Maps

Example: a neural network is a parametric map in **Smooth**



Parametric Maps

Example: a neural network is a parametric map in **Smooth**

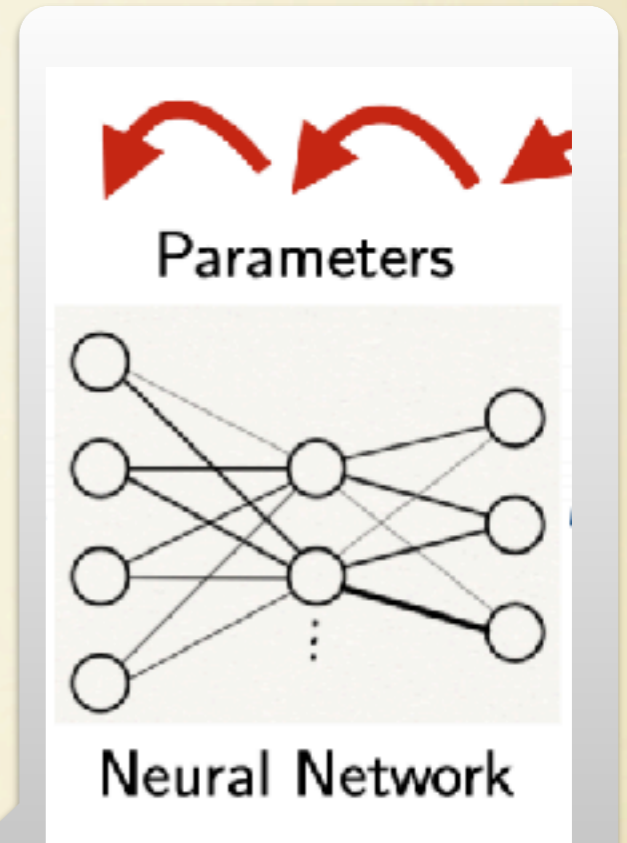


Lenses

$$\text{Lens}(\mathcal{C}) \quad l: (A, A') \rightarrow (B, B')$$

$$\text{get}_l: A \rightarrow B$$

$$\text{put}_l: A \times B' \rightarrow A'$$



a neural network is a lens in **Smooth**

$$f: A \rightarrow B$$

$$R[f]: A \times B' \rightarrow A'$$

$$(\vec{x}, \vec{y}) \mapsto J_f(\vec{x})^T \cdot \vec{y}$$

Reverse derivative categories *

Robin Cockett

University of Calgary, Department of Computer Science, Canada
robin@umcgary.ca

Geoffrey Cruttwell

Norumbia University, Department of Mathematics and Computer Science, Canada
geoffcruttwell@norumbia.ca

Jonathan Gallagher

Dalhousie University, Department of Mathematics and Statistics, Canada
jonathan.gallagher@dal.ca

Jean-Simon Pacaud Lemay

University of Oxford, Department of Computer Science, UK
jean-simon.lemay@ke.ox.ac.uk

Benjamin MacAdam

University of Calgary, Department of Computer Science, Canada
benjamin.macadam@ucalgary.ca

Gordon Plotkin

Google Research
gtp@inf.ed.ac.uk



Dorette Pronk

Dalhousie University, Department of Mathematics and Statistics, Canada
dorette.pronk@dal.ca

String Diagrams for Lenses


We use string diagrams in $\text{Tamb}(\mathcal{C})$

Faithful embedding of $\text{Lens}(\mathcal{C})$ in $[\text{Lens}(\mathcal{C})^{op}, \text{Set}] \cong \text{Tamb}(\mathcal{C})$





Parametric Lenses

Parametric lenses = parametric maps in ***Lens(C)***
= maps in ***Para(Lens(C))***




Composing Parametric Lenses

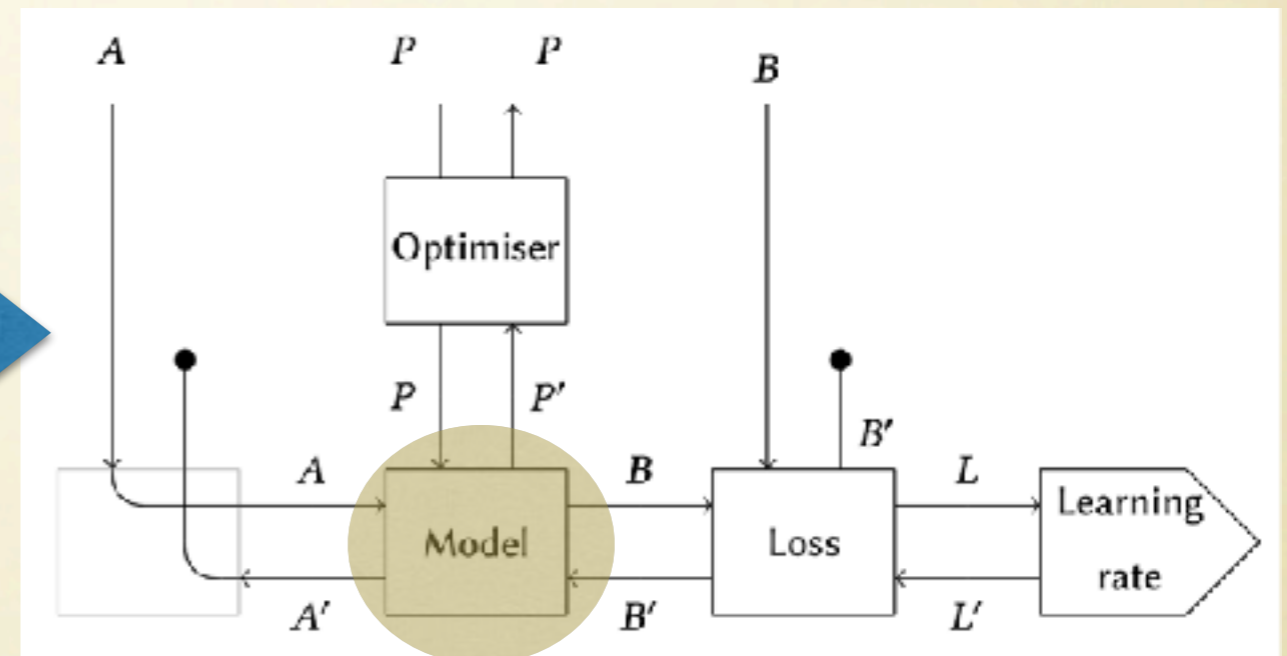



Parametric Lenses


Example: a neural network is a parametric lens in **Smooth**



Where we are, so far



Optimiser as reparametrisation



Lens(Smooth)

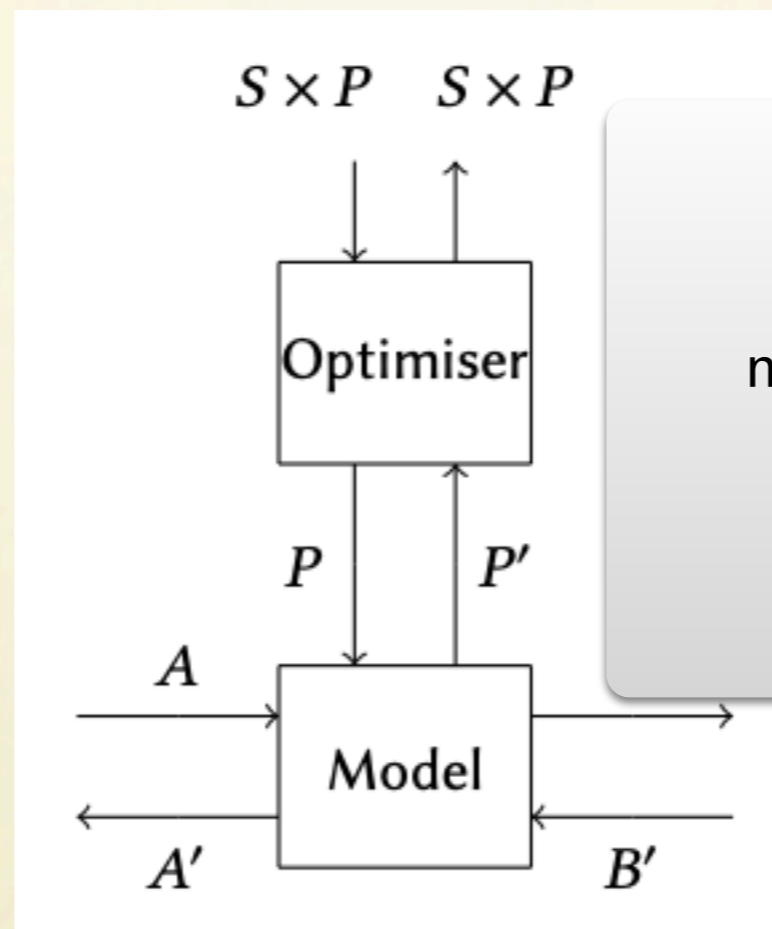
$gd: (P, P) \rightarrow (P, P')$

$get_{gd}: p \mapsto p$

$put_{gd}: (p, p') \mapsto p + p'$

Basic gradient descent

Optimiser as reparametrisation



Lens(Smooth)


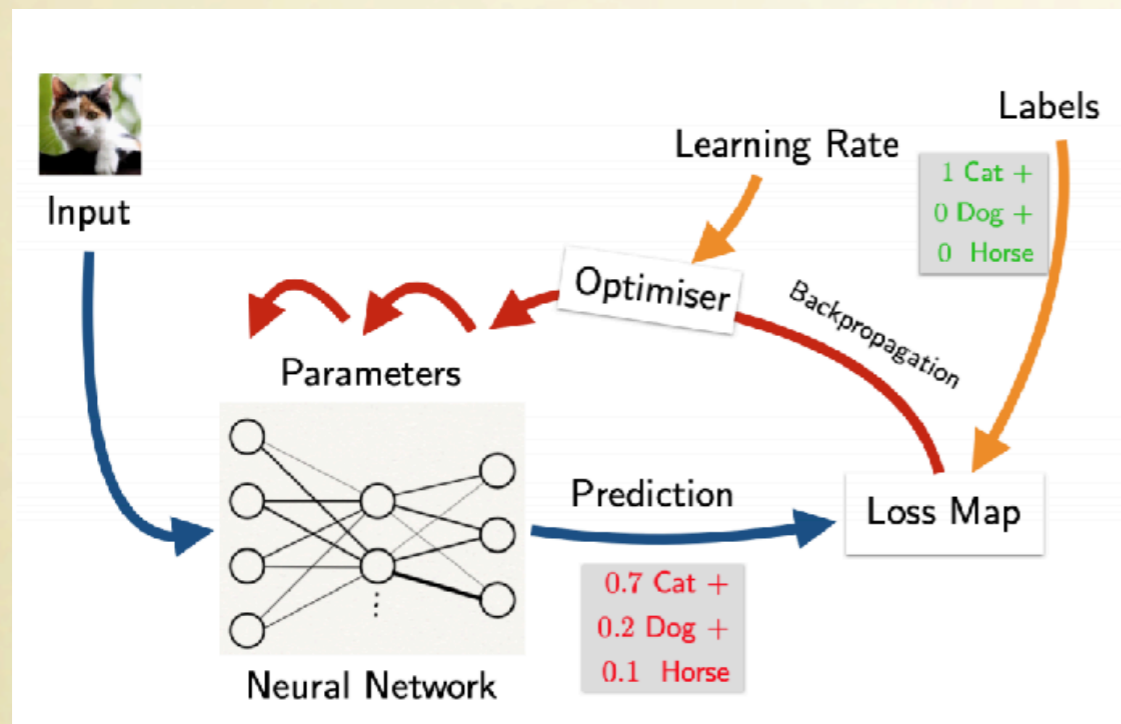
$$\text{nm} : (S \times P, S \times P) \rightarrow (P, P')$$

$$\text{get}_{\text{nm}} : (s, p) \rightarrow p - \gamma s$$

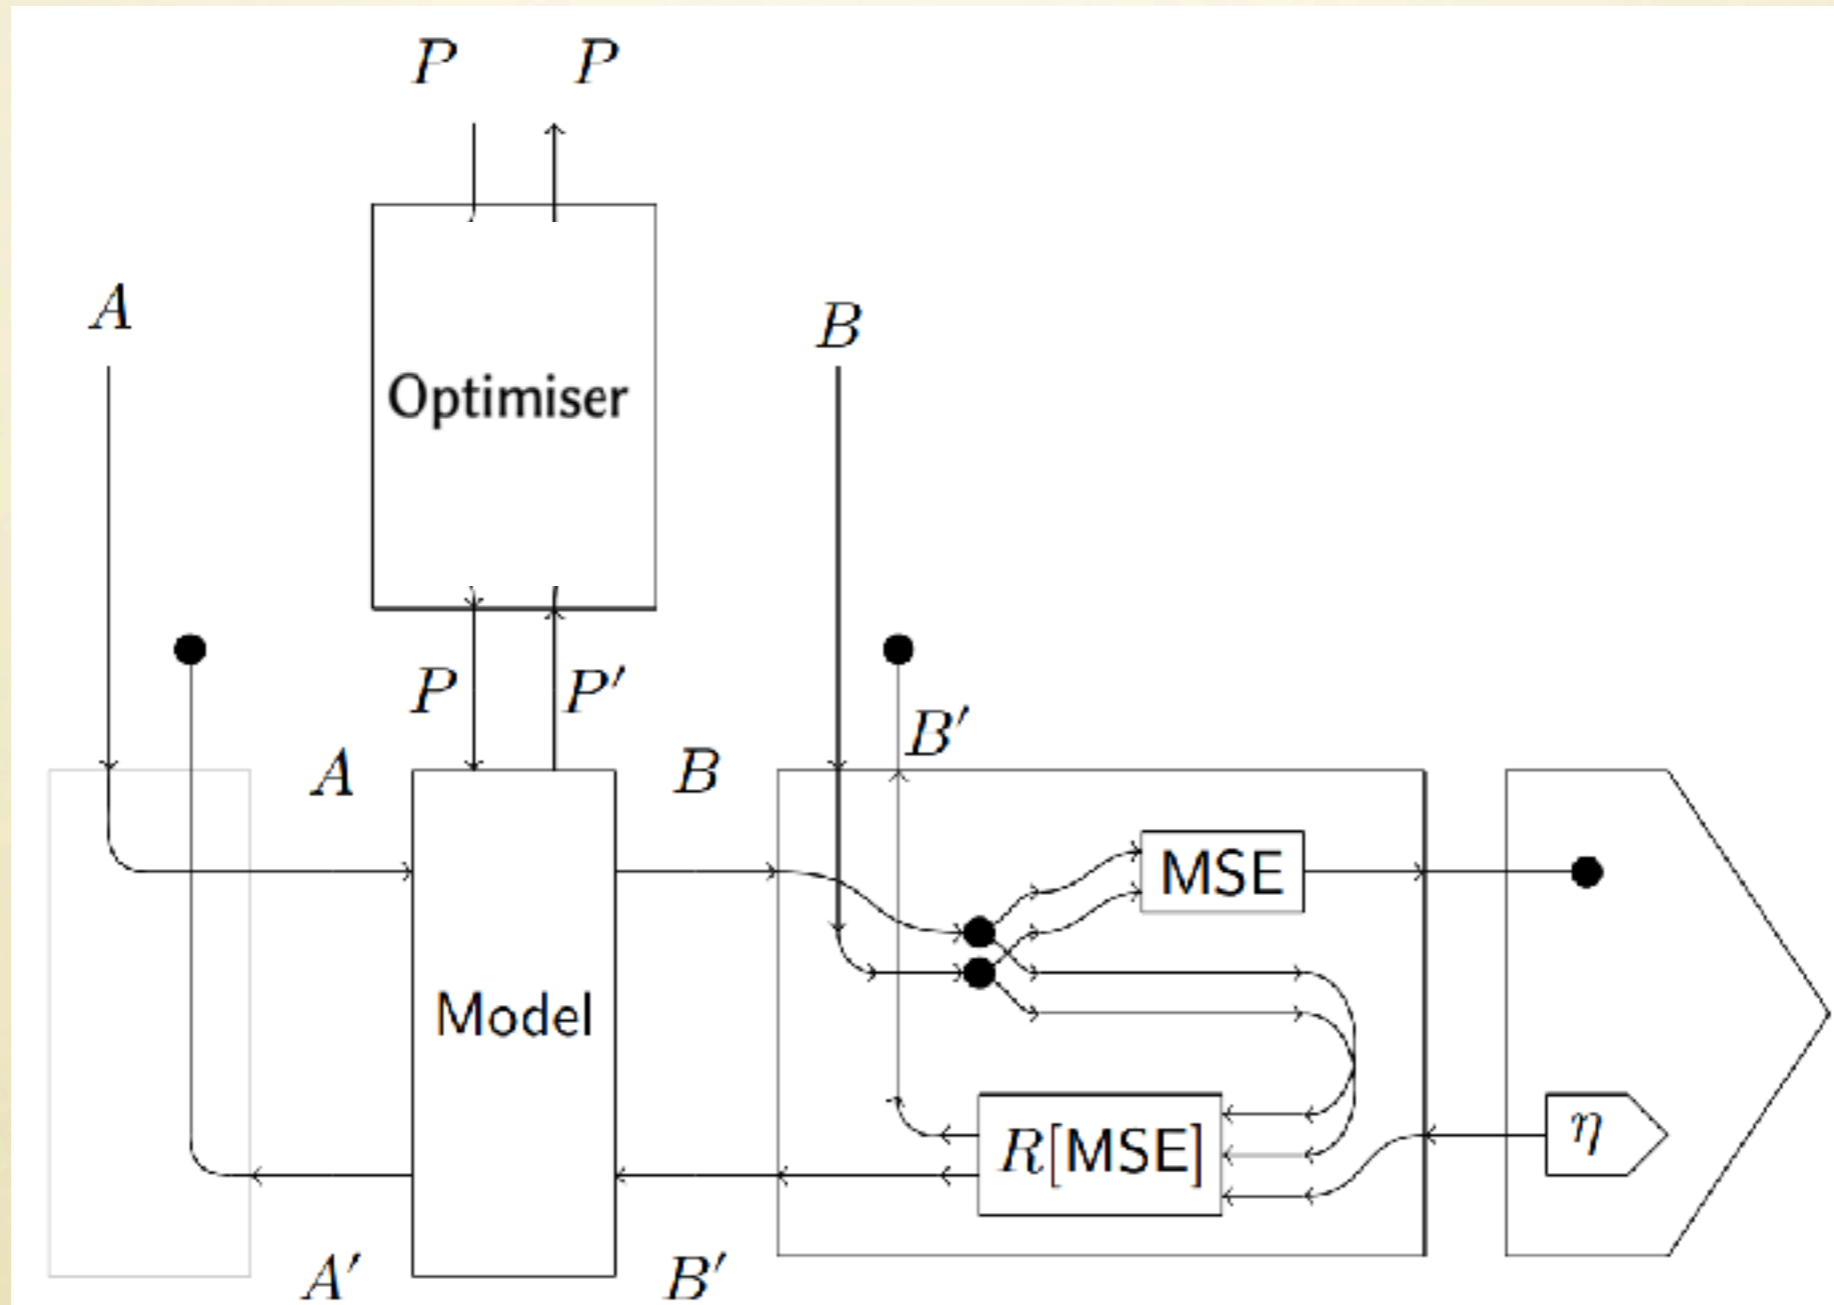
$$\text{put}_{\text{nm}} : (s, p, p') \mapsto (\gamma s + p', p - s)$$

‘Stateful’ optimiser such as **Adagrad**, **Adam**, **Nesterov**, **Momentum**, etc. can be all formalised uniformly as reparametrisations.

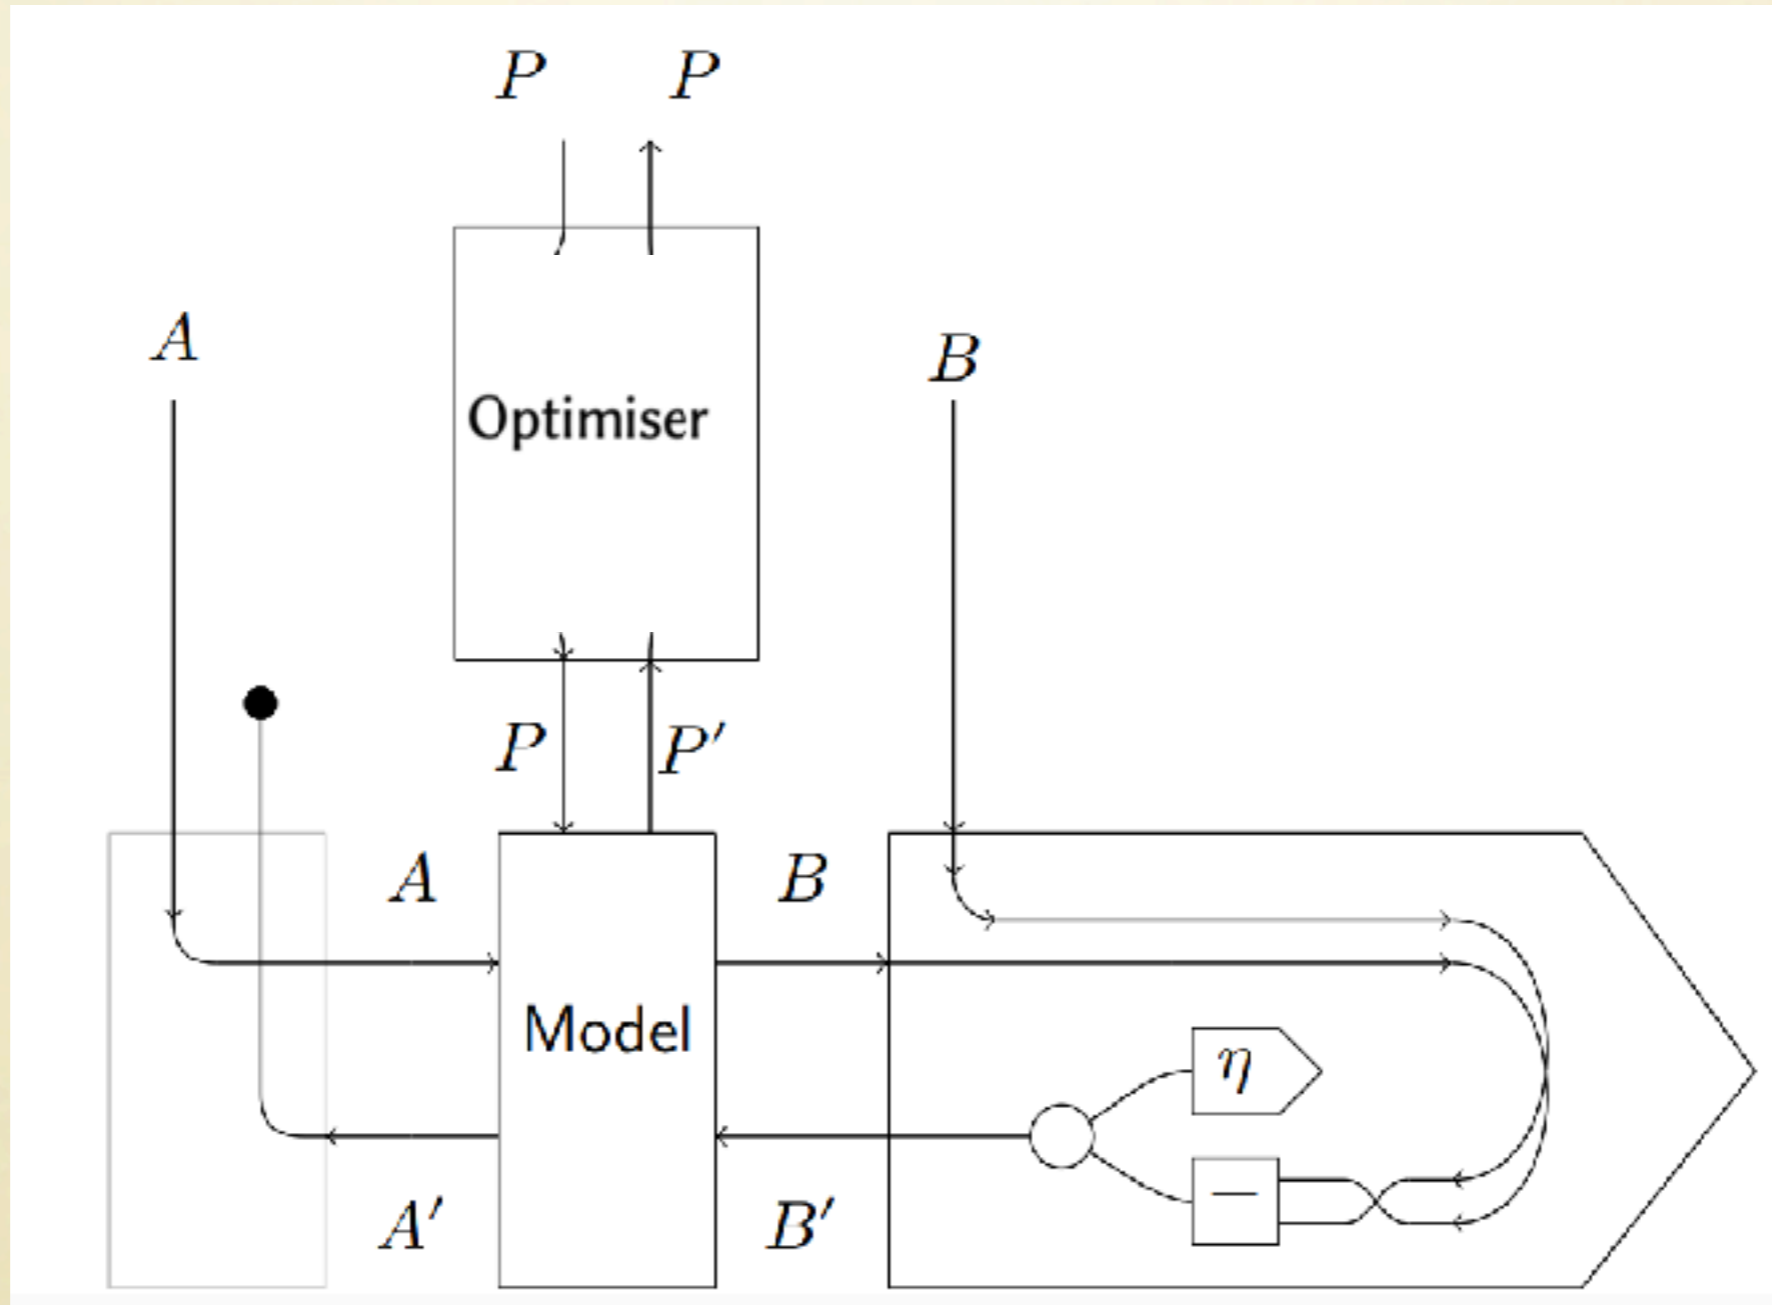
Where we are, so far



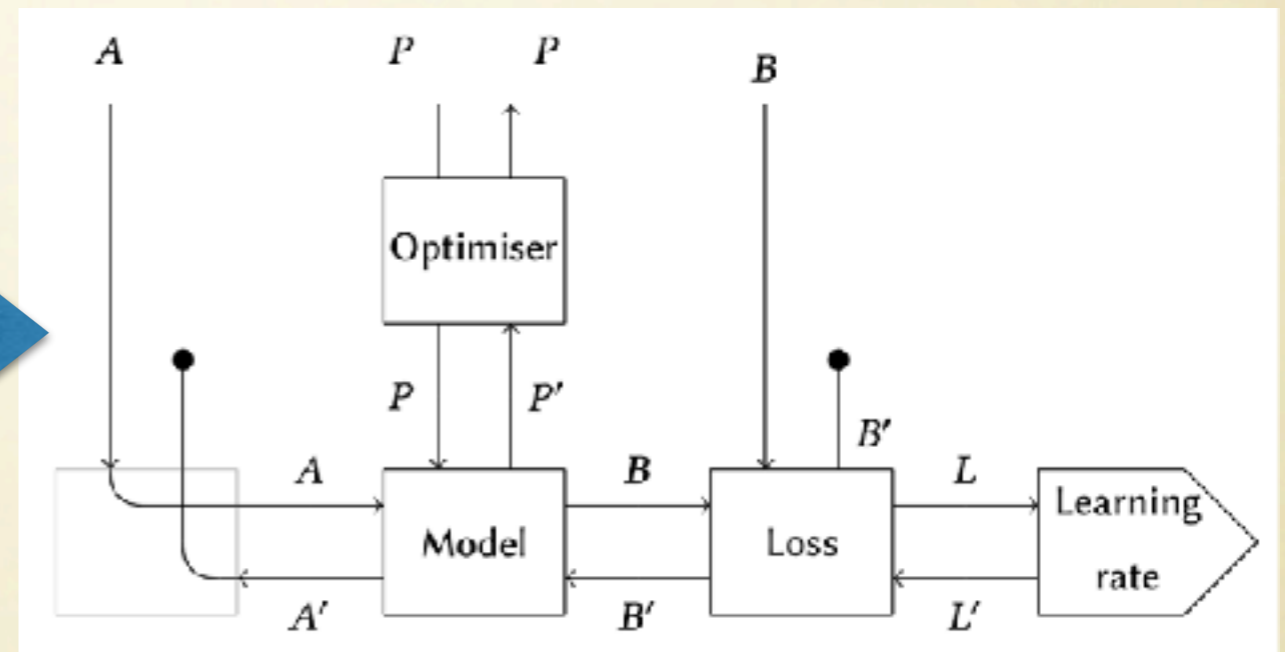
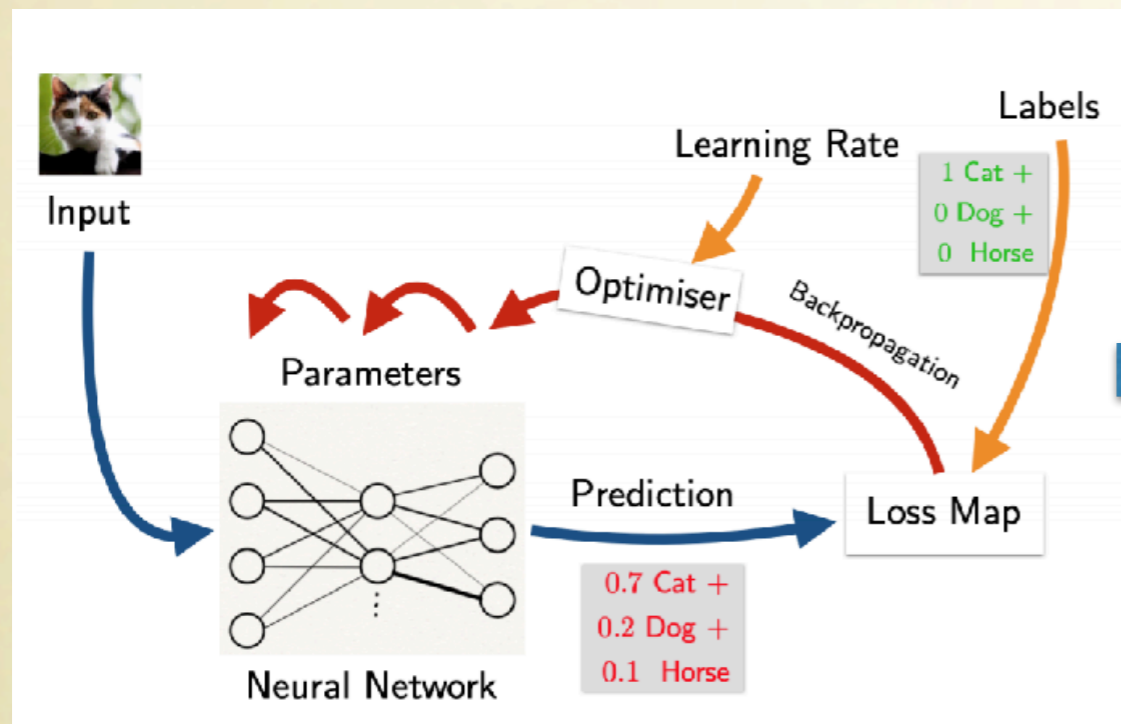
Mean Squared Error



Mean Squared Error




Where we are, so far



Variations I: Deep Dreaming




Input
Image




``Over-
processed''
Image



Variations I: Deep Dreaming




Parameter Learning




Deep dreaming

Variations II: Learning Boolean circuits



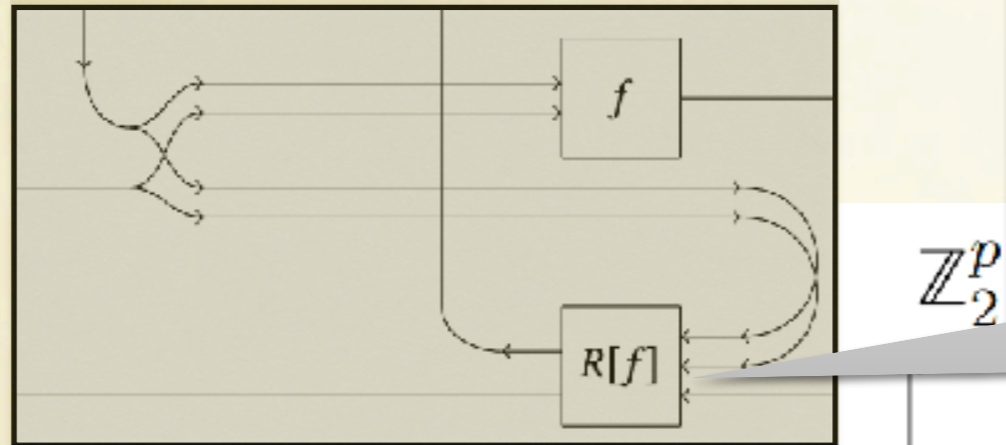
Neural networks



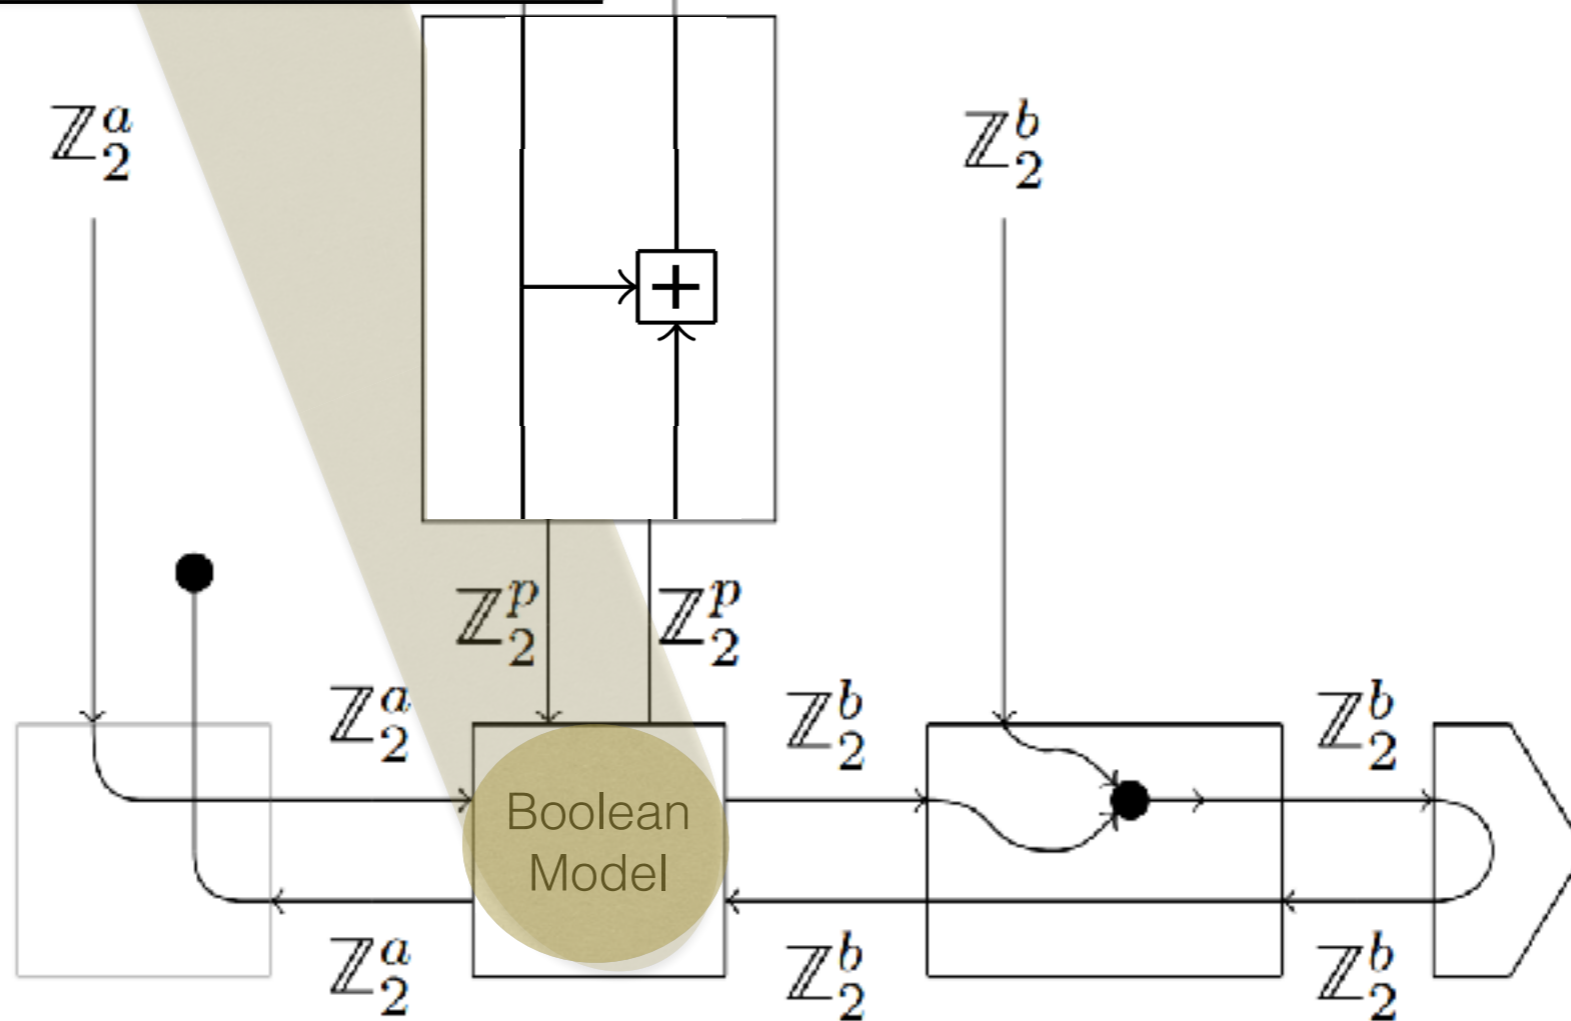

Boolean circuits

Instead of ***Smooth*** we work in ***Bool***


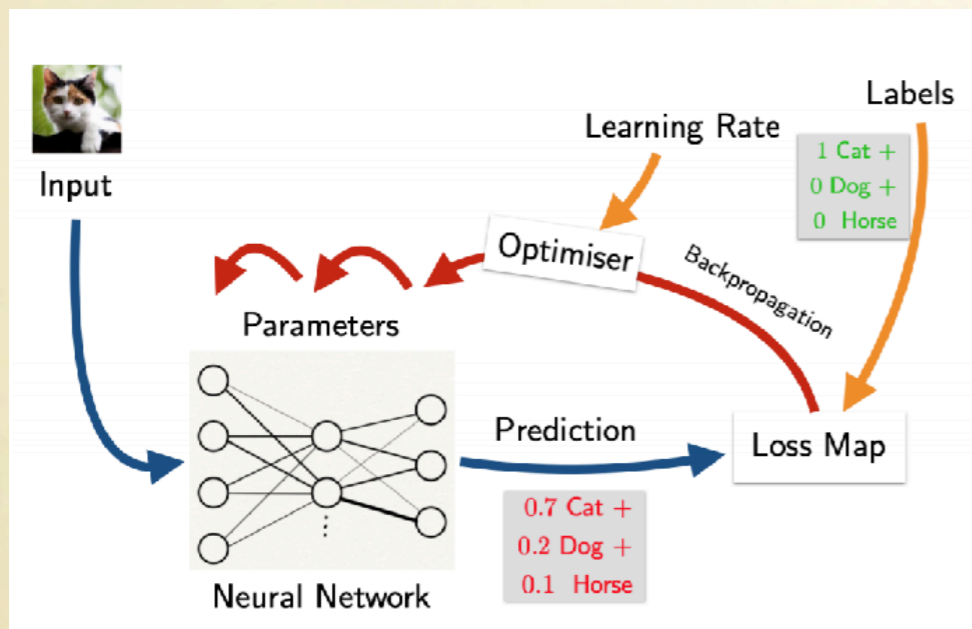
Variations II: Learning Boolean circuits



$R[f]$ is defined **inductively** on circuit syntax:



Discussion



An algebraic framework for gradient-based learning.

Discussion

Uniformity

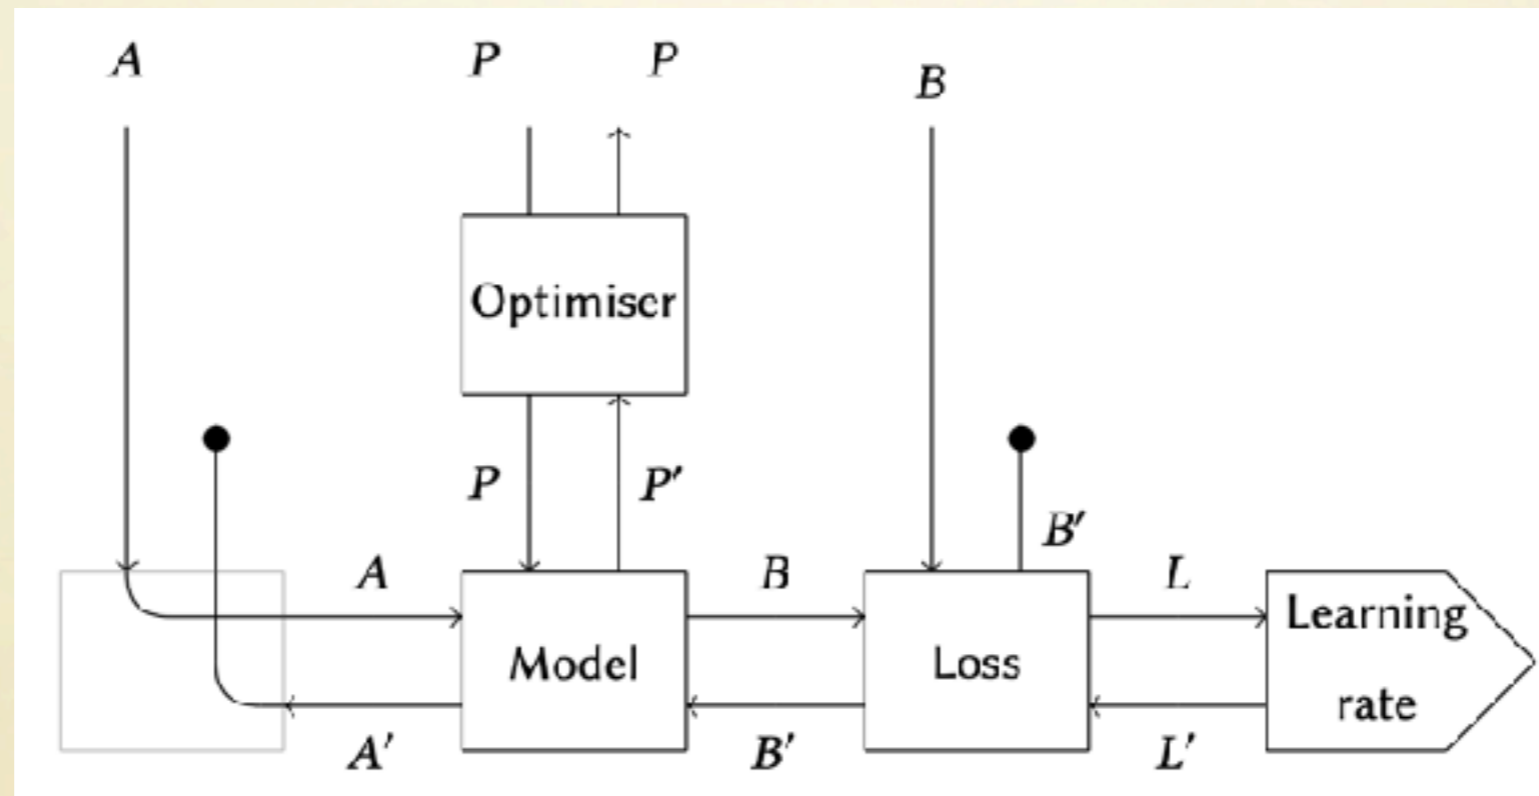
All components at work in the learning process are reduced to a single concept: parametric lenses.

Abstraction

A wide range of optimisers, models, loss maps, etc. are instances of the framework.


Discussion

Compositional reasoning



Discussion

Resource-sensitive Formalism



Discussion

Going forward

Learning-Theoretic Learning beyond gradient-descent

Learning
Neural
Networks

Learning
Boolean
Circuits

Bayesian
Update

Automata
Learning

Learning
Dynamical
Systems

Discussion

Implementation

<https://github.com/statusfailed/numeric-optics-python/>

<http://catgrad.com/p/reverse-derivative-ascent>

Bibliography

G. Cruttwell, B. Gavranovic, N. Ghani, P. Wilson, F. Zanasi - *Categorical Foundations of Gradient-Based Learning*, ESOP 2022

P. Wilson, F. Zanasi - *Reverse Derivative Ascent: a Categorical Approach to Learning Boolean Circuits*, ACT 2020