Deep Learning
Some case studies

L. Blondé, A. Mollaysa, P. Strasser, A. Kalousis

Mar2019

Haute Ecole Spécialisée
de Suisse occidentale
Fachhochschule Westschweiz
University of Applied Sciences and Arts
Western Switzerland

UNIVERSITÉ DE GENÈVE
FACULTÉ DES SCIENCES
Département d'informatique

Outline

Machine Learning
- Buzzwords
- Supervised Learning
- Neural Networks
- Reinforcement Learning
- Imitation Learning

Case Studies
- Modelling Pathological Gaits
- Learning to Walk
- Learning to Generate Molecules

The Holy Trinity of Hype

AI **Artificial Intelligence**: design "intelligent" programs
ML **Machine Learning**: learn from data without being explicitly programmed
DL **Deep Learning**: use (deep) neural networks as function approximators in ML

ML (and especially DL) is responsible for the biggest recent successes in AI.

This talk

(1) Coarse overview of ML concepts
(2) Introduction to 3 case studies of Deep Learning

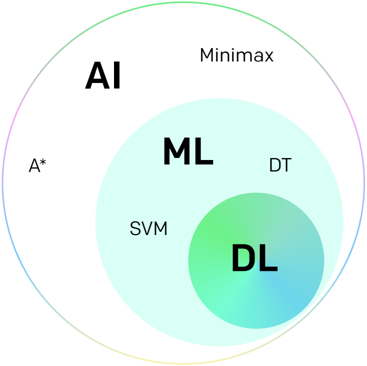

Figure: Russian dolls of AI

Crowded Area

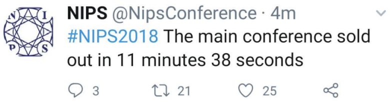

Figure: Tweet from the NIPS (now NeurIPS) conference organisation committee

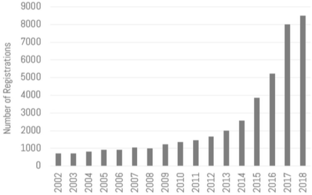

Figure: Number of registrations to the NeurIPS conference over the years

Branches of ML

- Supervised Learning (case **1**)
- Unsupervised Learning (case **3**)
- Reinforcement Learning (case **2**)

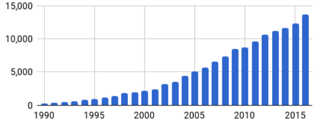

Figure: Number of RL papers published over the last 30 years

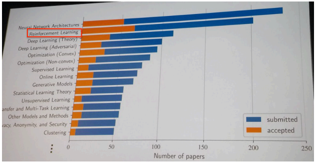

Figure: Distribution of papers by sub-domains published in ICML 2018

Learning (v1.0)

Ability for an **agent** to improve performance after observing **data**.

The **agent** is the learning system ("the" AI)

Agents must...

- Be naturally **adaptive**:
  No "hard-coded" behaviour
- **Encode** domain-specific knowledge underlying the data
- Remain **flexible** enough to enable adaption and improvement

Learning Theory is not Statistics

Statistics is model-centric: model the data

Learning Theory is algorithm-centric: in classification, we model the decision boundary, not the data distribution

Figure: An agent observing data.

Data

Represented by **pairs** (x, y), where

- $x \in \mathcal{X}$ is the **input**
- $y \in \mathcal{Y}$ is the **output**

A **dataset** is defined as a finite set of input-output pairs $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{n}$.

Prediction

For every input x, the agent **predicts** an output $\hat{y}$, according to her internal decision model $f : \mathcal{X} \to \mathcal{Y}$.

Performance Evaluation

The **prediction error** is measured with some point-wise **loss** function ℓ.

ℓ measures the gravity of the mistake.

Agent's Objective

Find the **relationship** $f : x \mapsto \hat{y}$ between elements of $\mathcal{X}$ and $\mathcal{Y}$ that minimises the **prediction error**: $\ell(\hat{y}, y) = \ell(f(x), y)$:

$$\min_{f} \ \ell(f(x), y)$$

Types of SL

The types differ by their **outputs**:

- **Classification**: Binary (**labels**)
  $\mathcal{Y} = \{0, 1\}$ corresponding to **classes**.
  Extension: multi-class (3+ classes)
- **Regression**: Real-valued outputs
  $\mathcal{Y} \subseteq \mathbb{R}$

Examples

- Image classification problem
  $\mathcal{Y} = \{\text{elephant}, \text{dog}, \text{cat}, \text{gorilla}\}$

"Cat"

What is the object in the image?

- Facebook "likes" count prediction
  $\mathcal{Y} = \mathbb{R}_+$

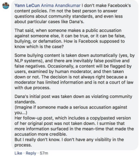

Yann LeCun Anima Anandkumar I don't make Facebook's content policies. I'm not the best person to answer questions about community standards, and even less about particular cases like Dana's.

That said, when someone makes a public accusation against someone else, it can be true, or it can be false, bullying, or defamation. How is Facebook supposed to know which is the case?

Some bullying content is taken down automatically (yes, by NLP systems), and there are inevitably false positive and false negatives. Occasionally, a content will be flagged by users, examined by human moderator, and then taken down or not. The decision is not always right because a moderator has limited information and is not a court of law with due process.

Dana's initial post was taken down as violating community standards.
(Imagine if someone made a serious accusation against you...)
Her follow-up post, which includes a copy/pasted version of her original post was not taken down. I surmise that more information surfaced in the mean-time that made the accusation more credible.
But I really don't know. I don't have any visibility in the process.

Like · Reply · 57m

8.73

How many "likes" will this post get?

"Big" Data

To fit more data, we need **more complex** models that are **easy to scale up**.

Best Candidate

(Deep) Neural Networks

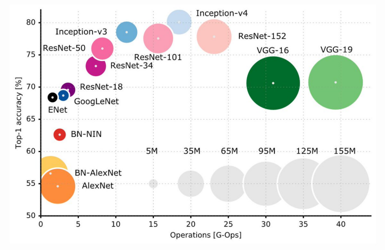

Figure: Various successful neural architectures.

Neural Network

- Hierarchical structure organised in **layers**
- Propagates information from an **input layer** to an **output layer**
- Layers in-between are called **hidden**
- Each layer is composed of **neurons**
- Neurons of a hidden layer are **connected** both to the previous and next layer

Components

- **Key hyperparameter**: depth and width
- **Learned variables**: weights and biases

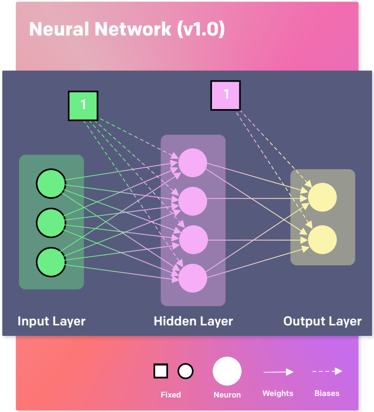

Surgeon's Dream

Anticipate the outcome of a **given** surgical operation on a patient's walking pattern ("gait") **without** carrying it out in the real world.

Problem

Model the **gait** of a given patient given the patient's **clinical data**.

Formulation

x Clinical data

y Walking gait

f Mapping from clinical data to walking gait: $y = f(x)$.

Goal Learn f.

Figure: Surgeon's dream: knowing f.

Measured Data

Measured by hand by a physiotherapist

- **Range of motion:** min/max joint angles (contraction)
- **Spasticity:** "stiffness" of the muscles tied to a joint
- **Selectivity:** controlability over individual muscles
- **Anthropomorphy:** dimensions of the body (weights, sizes, lengths)

Recorded Data[1]

Recorded via a marker+camera system in the form of **time series**

- Joint positions
- Joint angles
- Joint velocities
- Joint angular velocities

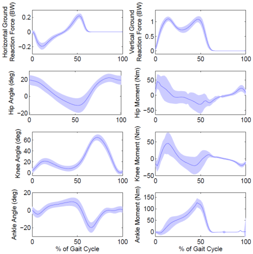

Figure: Joint angles (left), and joint moments (right) of a gait cycle, starting at heel strike.

Figure: Backside camera view of HUG's system. `https://www.unige.ch/medecine/chiru/en/research-groups/943armand/`

[1] Attias et al., "Feasibility and reliability of using an exoskeleton to emulate muscle contractures during walking", 2016.

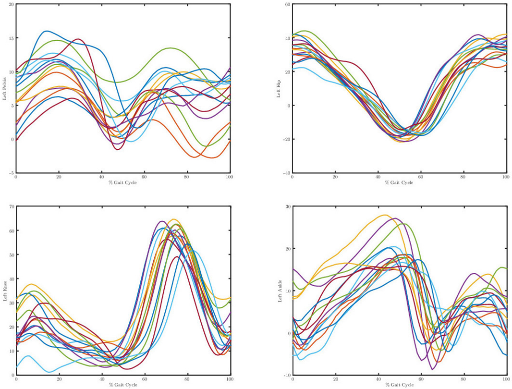

Figure: Evolution of the angle values or different joints, for a **single** patient.

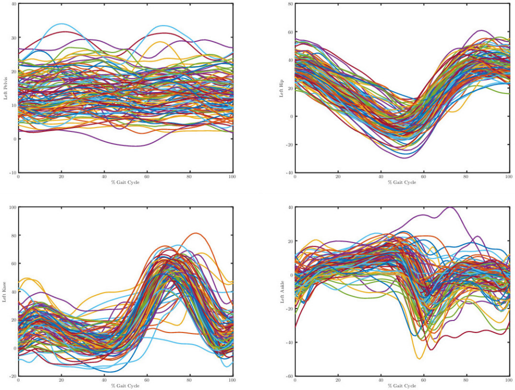

Figure: Evolution of the angle values or different joints, across **several** patient.

Difficulties

Huge variance **across** patients, but also within gaits of a **single** patient.

Model v1.0: MLP

Mulit-Layer Perceptron.

$$y_{0:T} = f(x)$$

Model v2.0: RNN

Recurrent Neural Network. Relies on the sequential nature of the data.

$$\begin{cases} h_0 & = f(x) \\ y_t & = g(h_t) \qquad \forall t \in [0, T-1] \\ h_{t+1} & = \tau(h_t, y_t) \qquad \forall t \in [0, T-1] \end{cases}$$

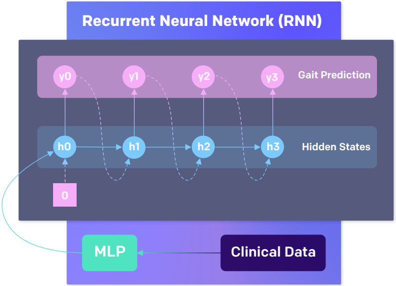

Method	Error (MSE)
Median	100.21
Recurent Neural Network	92.16
MLP	107.4

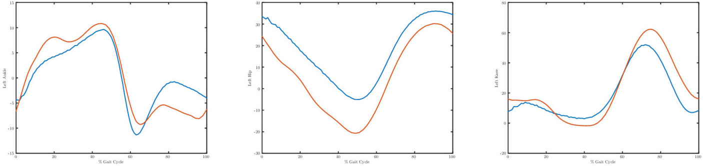

Figure: Gait predictions.

So, is it good?

Not really.

- **Issue**: not enough data
- **Consequence**: generates gaits of low fidelity with the patients' gaits

Attempted Solution

Bring in a **simulator** to generate more data.

- **New issue**: the simulator needs to be **specifically tuned** for the patient/pathology
- **New approach**: use Reinforcement Learning and Imitation Learning.

Data
Clinical data of the same nature as case study 1, but specifically for patients suffering from **Cerebral Palsy**.

Figure: Backside camera view of HUG's system.

New
Use a **simulator**.

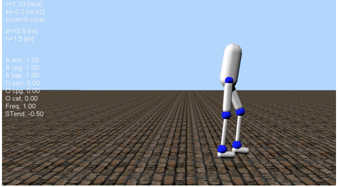

Figure: EPFL BIOROB Team's simulator.
`https://biorob.epfl.ch`

Problem
Find the **parametrisation** of the simulator that corresponds to the pathology displayed in the patient data.

RL is the field of **sequential decision-making under uncertainty**.

An **agent** (decision maker) **interacts** with a previously unknown **environment** and receives **rewards** upon interaction.

Agent's Objective (v1.0)

Maximise the **long-term cumulative reward**.

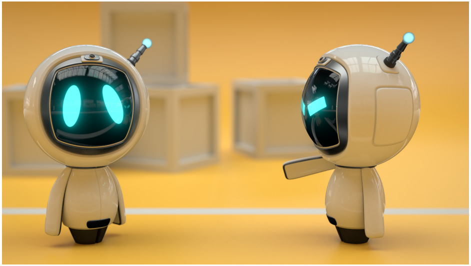

Figure: An agent interacting with its environment.

[2] Sutton and Barto, "Reinforcement Learning: An Introduction", 1998.

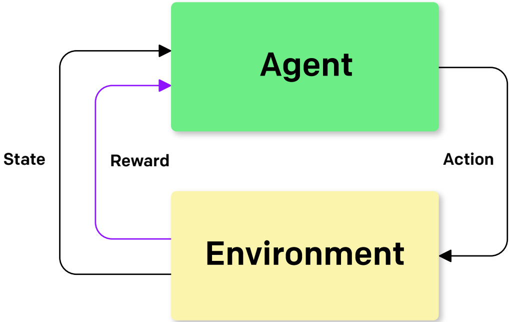

Figure: Interaction diagram (v1.0)

RL Caveats

Interactive Nature

Heavy interaction with the environment while learning: **safety**[3,4] issues in real world scenarios.

Reward Design

Preliminary burden: handcraft reward signal(s) to induce the desired behavior. This is called **reward shaping**[5].

$$r(b_z^{(1)}, s^P, s^{B1}, s^{B2}) = \begin{cases} 1 & \text{if } \mathrm{stack}(b_z^{(1)}, s^P, s^{B1}, s^{B2}) \\ 0.25 & \text{if } \neg\mathrm{stack}(b_z^{(1)}, s^P, s^{B1}, s^{B2}) \wedge \mathrm{grasp}(b_z^{(1)}, s^P, s^{B1}, s^{B2}) \\ 0.125 & \text{if } \neg(\mathrm{stack}(b_z^{(1)}, s^P, s^{B1}, s^{B2}) \vee \mathrm{grasp}(b_z^{(1)}, s^P, s^{B1}, s^{B2})) \wedge \mathrm{reach}(b_z^{(1)}, s^P, s^{B1}, s^{B2}) \\ 0 & \text{otherwise} \end{cases}$$

Figure: Reward from the *block-stacking paper*[6]

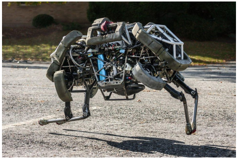

Figure: Boston Dynamics *WildCat* robot

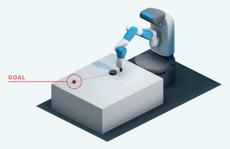

Figure: Hindsight Experience Replay[7]

[3] Amodei et al., "Concrete Problems in AI Safety", 2016.
[4] Held et al., "Probabilistically Safe Policy Transfer", 2017.
[5] Ng, Harada, and Russell, "Policy invariance under reward transformations: Theory and application to reward shaping", 1999.
[6] Popov et al., "Data-efficient Deep Reinforcement Learning for Dexterous Manipulation", 2017.
[7] Andrychowicz et al., "Hindsight Experience Replay", 2017.

Instead of receiving rewards upon interaction, the agent is initially provided with **demonstrations** from an **expert** and does not receive any external feedback while learning.

Objective

Mimic the demonstrated behaviour.

- π_e: expert policy
- Demonstration τ_e: trajectory from π_e
- Trajectory: trace of interaction with the MDP, i.e. state-action pairs (or just states[8][9]) collected during one episode: $\{(s_0, a_0), \ldots, (s_T, a_T)\}$.

Figure: Kinaesthetic Teaching in Virtual Reality[11]

[8] Liu et al., "Imitation from Observation: Learning to Imitate Behaviors from Raw Video via Context Translation", 2017.

[9] Merel et al., "Learning human behaviors from motion capture by adversarial imitation", 2017.

[11] Zhang et al., "Deep Imitation Learning for Complex Manipulation Tasks from Virtual Reality Teleoperation", 2017.

[12] Bagnell, *An invitation to imitation*, 2015.

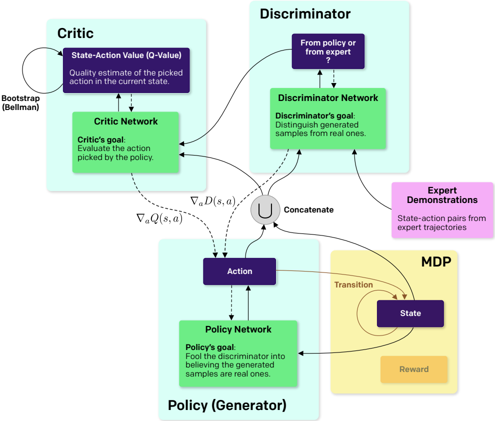

Figure: Sample-efficient Adversarial Mimic.

Sample-Efficient Adversarial Mimic

Video Demo

Figure: https://arxiv.org/abs/1809.02064

Figure: https://youtu.be/-nCsqUJnRKU

Molecule Design[13]

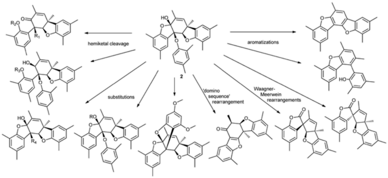

Figure: Diversity-oriented synthesis: producing chemical tools for dissecting biology

Conditional Image Generation[14]

Figure: Class-conditional samples generated with generative Adversarial Nets (GAN)

Style Transfer

credit to @DmitryUlyanovML

[13] O'Connor, Beckmann, and Spring, "Diversity-oriented synthesis: producing chemical tools for dissecting biology", 2012.
[14] Brock, Donahue, and Simonyan, "Large scale gan training for high fidelity natural image synthesis", 2018.

Problem

Can we do the same with molecules?

Goal

- Generate molecules with specific properties
- Enable style (property) transfer from molecule to molecule.

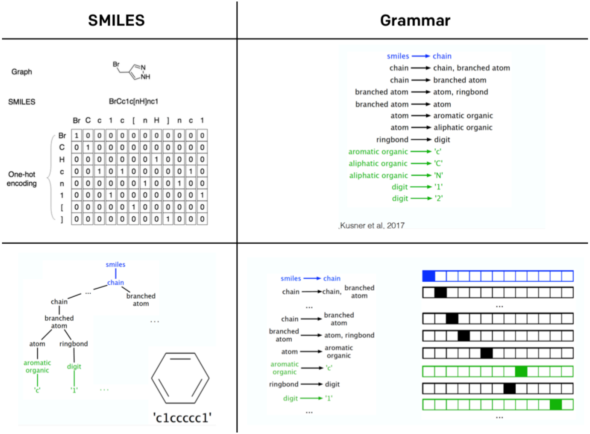

Figure: Introduction of a **grammar** in a SMILES molecule representation.

Data

Pairs (x, y) where

- x is the **molecule**
- y is the **property**

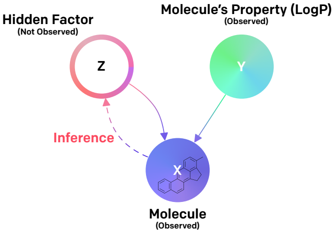

Figure: Graphical model.

Benefits

With the learned $p_\theta(x|z, y)$, we can:

- Generate molecules that have a desired property $y^\star$
- Achieve a **diverse** generation of molecules with property $y^\star$ by sampling from $p_\theta(x|z, y^\star)$ with various z values from $z \sim \mathcal{N}(0, 1)$
- Modify the property of molecule x to have the property $y^\star$ while staying close to the original molecule in terms of structure by sampling from $p_\theta(x|z, y^\star)$ with various z values from $z \sim q_\phi(z|x)$

Conditional Generation

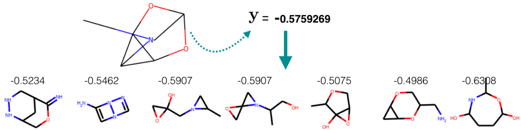

Figure: Generated molecules with property value within a 15% range from the desired value.

Style Transfer

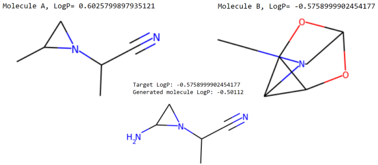

Figure: Property transfer over molecules.

References I

- M. Attias et al. "Feasibility and reliability of using an exoskeleton to emulate muscle contractures during walking". In: *Gait & Posture* 50 (2016), pp. 239–245. ISSN: 0966-6362. DOI: https://doi.org/10.1016/j.gaitpost.2016.09.016. URL: http://www.sciencedirect.com/science/article/pii/S0966636216305781.
- Richard S Sutton and Andrew G Barto. "Reinforcement Learning: An Introduction". 1998.
- Dario Amodei et al. "Concrete Problems in AI Safety". In: (June 2016). arXiv: 1606.06565 [cs.AI]. URL: http://arxiv.org/abs/1606.06565.
- David Held et al. "Probabilistically Safe Policy Transfer". In: (May 2017). arXiv: 1705.05394 [cs.RO]. URL: http://arxiv.org/abs/1705.05394.
- Andrew Y Ng, Daishi Harada, and Stuart Russell. "Policy invariance under reward transformations: Theory and application to reward shaping". In: *International Conference on Machine Learning (ICML)*. 1999, pp. 278–287. URL: http://www.robotics.stanford.edu/~ang/papers/shaping-icml99.pdf.
- Ivaylo Popov et al. "Data-efficient Deep Reinforcement Learning for Dexterous Manipulation". In: (Apr. 2017). arXiv: 1704.03073 [cs.LG]. URL: http://arxiv.org/abs/1704.03073.
- Marcin Andrychowicz et al. "Hindsight Experience Replay". In: (July 2017). arXiv: 1707.01495 [cs.LG]. URL: http://arxiv.org/abs/1707.01495.
- Yuxuan Liu et al. "Imitation from Observation: Learning to Imitate Behaviors from Raw Video via Context Translation". In: (July 2017). arXiv: 1707.03374 [cs.LG]. URL: http://arxiv.org/abs/1707.03374.
- Josh Merel et al. "Learning human behaviors from motion capture by adversarial imitation". In: (July 2017). arXiv: 1707.02201 [cs.RO]. URL: http://arxiv.org/abs/1707.02201.
- Tianhao Zhang et al. "Deep Imitation Learning for Complex Manipulation Tasks from Virtual Reality Teleoperation". In: (Oct. 2017). arXiv: 1710.04615 [cs.LG]. URL: http://arxiv.org/abs/1710.04615.
- J Andrew Bagnell. *An invitation to imitation*. Tech. rep. Carnegie Mellon, Robotics Institute, Pittsburgh, 2015. URL: https://www.ri.cmu.edu/pub_files/2015/3/InvitationToImitation_3_1415.pdf.
- Cornelius J O'Connor, Henning SG Beckmann, and David R Spring. "Diversity-oriented synthesis: producing chemical tools for dissecting biology". In: *Chemical Society Reviews* 41.12 (2012), pp. 4444–4456.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. "Large scale gan training for high fidelity natural image synthesis". In: *arXiv preprint arXiv:1809.11096* (2018).

Propagation Rule

The **activation** of a neuron is a direct expression of the activations of the neurons from previous layers.

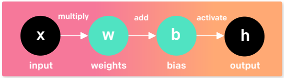

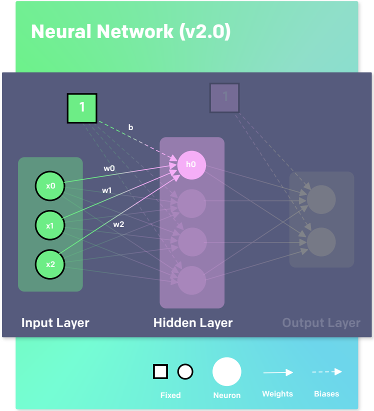

$$
\begin{aligned}
h_0 &= f\left(w^T x + b\right) \\
&= f\left(w_0 x_0 + w_1 x_1 + w_2 x_2 + b\right)
\end{aligned}
$$

f is an **activation function** ("non-linearity")

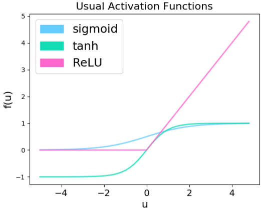

Usually no non-linearity at the **output layer**

Network Types

- **Feed-Forward** NN (Deep NN or DNN)
- **Recurrent** NN (RNN)

Layer Types

- Locally-connected
    - Convolutional
    - Sub-sampling (Pooling)
    - Upsampling (Deconvolutional)
- Fully-connected (FC or Dense)

From Task to Architecture

Design your NN for the task at hand

- Need to reason only at **local** scale: **locally-connected**
- Need to reason only at **global** scale: **fully-connected**
- Need to reason at **both** scales: **locally-connected + fully-connected**

Common Architectures

"Architecture" is most popular sub-area of ML research (#1 in published papers)

- **Fully Convolutional Network** (FCN)
- **Multi-Layer Perceptron** (MLP)
- **Convolutional Neural Network** (CNN or ConvNet)

An FCN is composed **exclusively** of locally-connected layers.

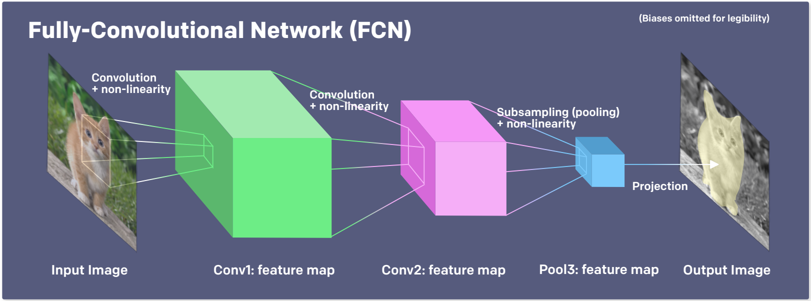

FCNs were first introduced for **Semantic Image Segmentation**.

An MLP is composed **exclusively** of fully-connected layers.

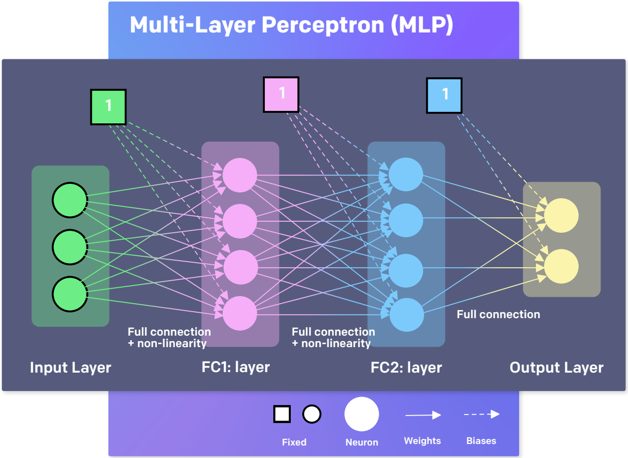

MLPs are the most common architectures as a whole but are also extremely commonly used as **building blocks** for more complex architectures.

A ConvNet has a **hybrid** layer composition: locally-connected layers followed by fully-connected layers.

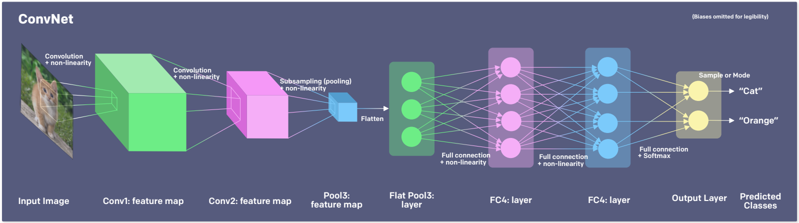

ConvNets are the **gold standard** architecture for **vision tasks** (images as inputs).

Supervised VAE

Encode a **property** as part of the latent representation via a Variational Auto-Encoder (VAE).

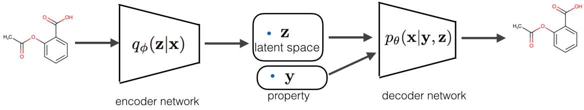

Property-enforcing Regularizer

Force the model to take into account the property information.

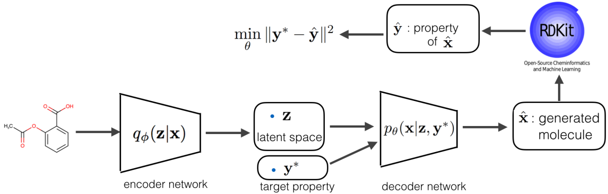

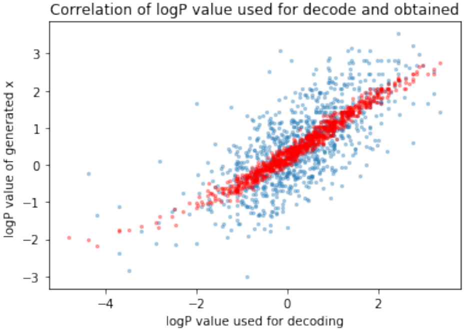

Figure: Correlation plot of the property of the original molecule and the property of the generated molecule.