Introduction

Samy Tindel

Purdue University

Stochastic calculus and Machine Learning - MA 595

PURDUE UNIVERSITY

Outline

Outline

My Purdue information

History:

- 12th year:
  ↪ as Professor at Purdue
- Before that:
  ↪ in Nancy (France)

Office: 434, Math building

Email: `stindel@purdue.edu`

Webpage: `https://www.math.purdue.edu/ stindel/`

Outline

A motivating example: generative AI

Deep generative models: Often in the headlines for deep fakes
$\hookrightarrow$ Also useful in physics and computational chemistry

Our guiding example: Continuous normalizing flows

Three steps:

1. Sample $\mathbf{x}_0$ according to the distribution μ_0 to be faked
2. Add noise to $\mathbf{x}_0$, thanks to a simple enough equation
   $\hookrightarrow$ This produces a process $\mathbf{X}_t$
3. Reverse time in the equation, starting from $\mathbf{X}_1$
   $\hookrightarrow$ This produces a new sample $\hat{\mathbf{X}}_1 \sim \mu_0$

Picture of the algorithm

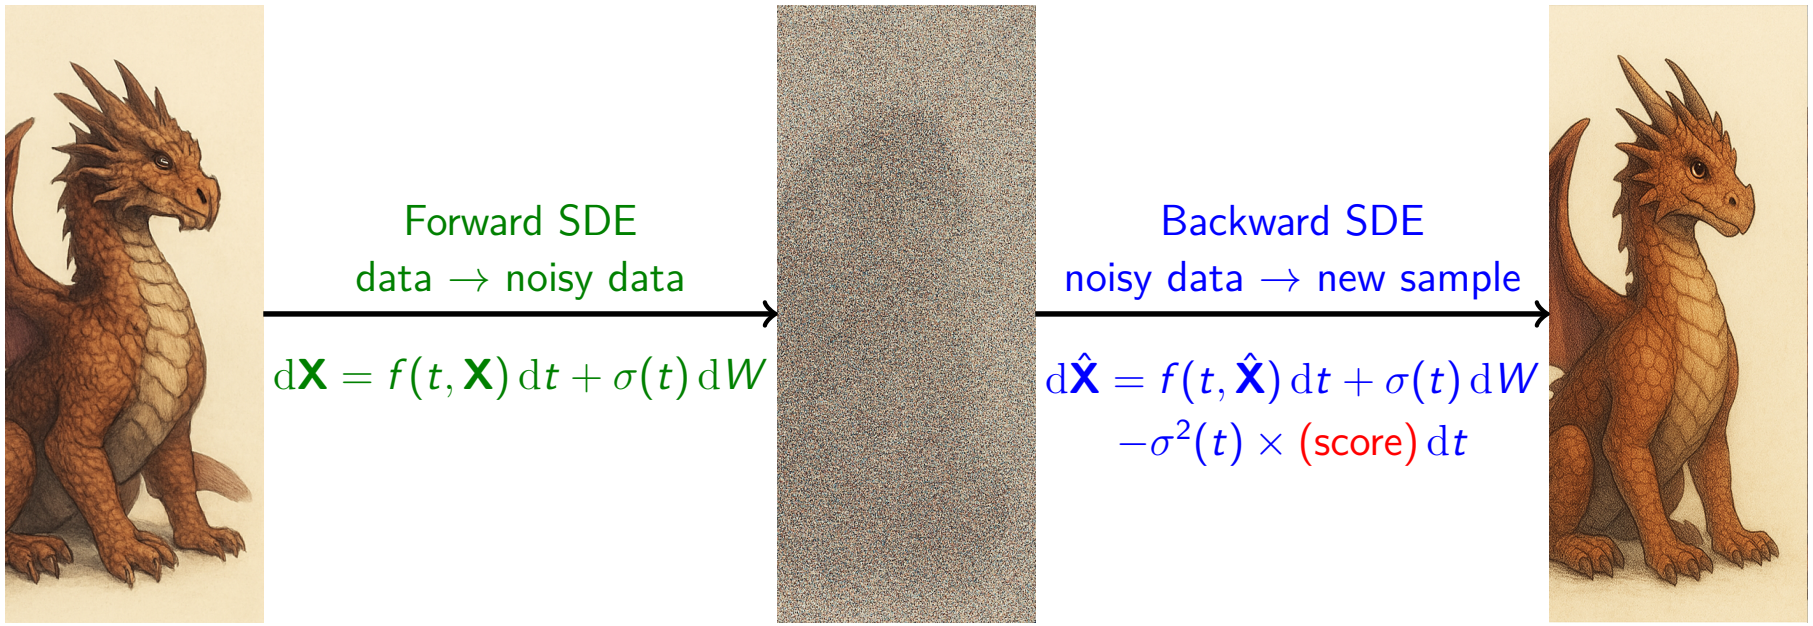

Successfully implemented in numerous domains

Adding noise: the forward equation

Noise addition: For $i = 1, \ldots, d$, one solves

$$\mathbf{X}_t^i = \mathbf{x}_0^i + \int_0^t f\left(s, \mathbf{X}_s^i\right) \mathrm{d}s + \int_0^t \sigma\left(s\right) \mathrm{d}\mathbf{W}_s^i$$

Ingredients:

- $f : [0,1] \times \mathbb{R} \to \mathbb{R}$ and $\sigma : [0,1] \to \mathbb{R}$
  $\hookrightarrow$ parameters of the algorithm.
- $\mathbf{W}$ is a d-dimensional Wiener process.

Remark: The dimension only enters through the initial data $\mathbf{x}_0$
$\hookrightarrow$ We may focus on a generic coordinate $X_t = \mathbf{X}_t^i$

Reversing time: the backward equation

Reversed process: $\hat{X}_t = X_{1-t}$ solves

$$\hat{X}_t = X_1 + \int_0^t \left[f\left(s, \hat{X}_s\right) - \sigma^2(s) \nabla_x \ln\left(p_s(\hat{X}_s)\right) \right] \mathrm{d}s + \int_0^t \sigma(s) \, \mathrm{d}\hat{W}_s$$

Notation:

- p_s is the density of the random variable X_s.
- $\hat{W}$ is another Brownian motion.

Conclusion: One gets $\hat{X}_1 \overset{(d)}{=} x_0$

$\hookrightarrow$ The deep fake is produced

Aim of the course

Observation: The two equations above involve

- A Brownian motion W.
- An integral $\int \sigma(s)\,\mathrm{d}W_s$ with respect to W.
- A stochastic differential equation, forward and backward in time.
- A score function $\nabla_x \ln p_s$, estimated by a neural network.

Main goal of the course:
$\hookrightarrow$ Understand those objects and their properties

Similar objects show up in
$\hookrightarrow$ control, reinforcement learning, signatures

Brief outline of the course

Chapters covered:

1. Basics on Brownian motion
2. Ito's formula
3. Stochastic differential equations
4. Reversed diffusion processes and generative AI algorithms
5. Stochastic control and neural networks
6. Rough paths and the signature method for machine learning
7. Relaxed control and reinforcement learning

Outline

Webpage

Course webpage:

`https://www.math.purdue.edu/stindel/teaching/ma595/ma595.html`

Contents:

- Announcements
- Calendar and schedule
- Slides
- Class notes in book form
- Written notes from class

Grades

Total score calculation:

• Project 1	200 pts.
• Project 2	200 pts.
• Bonus: attendance	10 pts.
• Bonus: answering questions	10 pts.
• Bonus: asking questions during presentations	10 pts.

Projects

Rules for the projects:

- Groups of 4 will be assigned
- Choose a proper topic.
  - Select a reasonable bibliography.
  - The theoretical content has to be substantial.
- Write a 10 pages report on your topic.
  $\hookrightarrow$ Math steps must be rigorous and well explained.
- Coding some of the methods is allowed
  $\hookrightarrow$ Not substitute for math contents.
- There will be a 30mn presentation with slides for each group.
  - No more than 20 slides
  - Slides for the presentations should not be too dense.

Participation bonus (1)

Participation/questions in class:

- Questions will be asked in class
- Volunteers will get some points towards the bonus
- You are expected to participate, not to give an exact answer
- Stupid answers don't exist
- Aim: get to know everyone

Participation bonus (2)

Participation/asking questions:

- Ask questions during your classmate's presentations
- Aim: stay active during presentations

Participation bonus (3)

Participation/attendance:

- Aim: incentive to come to class
- 4 absences allowed throughout the semester
  $\hookrightarrow$ without justification