

## Peace Engineering: Demystifying Machine Learning

Manel Martínez-Ramón; Ramiro Jordán; Donna Koechner; Xochiquetzal García-Necochea;  
Jeremiah Dacpano and Ryan García.

*Department of Electrical and Computer Engineering, The University of New Mexico,  
Albuquerque, New Mexico 87131-0001, USA  
Corresponding Author Email: [manel@unm.edu](mailto:manel@unm.edu)*

---

### ABSTRACT

#### CONTEXT

Artificial intelligence (AI) is a discipline that is developed and applied by scientists and engineers of multiple areas, so it should not be circumscribed only to the domain of Computer Science (CS). Nevertheless, the ins and outs of AI are generally unknown by the average higher education student, and the pedagogical practices make AI learning mostly out of reach for undergraduate students. Most AI courses are restricted to their theoretical and technological components, but the sociological part of AI is usually omitted while, as in many other applied science areas, deontology is an unavoidable component, and it would not be acceptable to ignore this aspect.

#### PURPOSE

The proposed approach is intended to change how AI is introduced. We constructed a course to teach AI to undergraduate students. We hypothesize that it is possible to teach AI in a comprehensive and useful way to students of engineering in various areas with limited mathematical backgrounds. The methodology and contents are designed to teach AI as a socio-technical system, where ethical and sociological components are introduced.

#### APPROACH

The class is organized into blocks containing the concepts of structure, criterion, and algorithm to provide an AI landscape to the students. The AI constructs are contextualized over the base of socio-technical systems rather than explaining them as pure artifacts or pieces of code. The practice contains experiments oriented to reproduce AI models and to study the bias, disinformation, trustability, manipulation, introduction of ethics, and regulation of AI.

#### ANTICIPATED OUTCOMES

The student's skills will be evaluated in terms of their ability to identify and describe the main structures and algorithms used in AI from a theoretical point of view, analyze the results of several experiments with already constructed learning machines, and design machines in practice sessions. The students will be assessed on the societal aspects of AI, where they will be asked to discuss the different topics and plan ways to insert values into AI.

#### SUMMARY

We present a new approach to the study of AI at undergraduate levels. We justify that the junior and senior student backgrounds are sufficient to successfully follow the class, and we introduce AI as a socio-technical system to introduce the necessary ethical and societal elements to construct the needed deontological practices in the development and use of AI.

#### KEYWORDS

socio-technical system, disinformation, artificial intelligence

## Introduction

We propose the design of a single introductory course in AI suitable for undergraduate students across campus to provide the student with basic AI skills and competencies (with emphasis in DL), promote awareness of the social implications of AI and the need for viewing it as a socio-technical construct, and to foster a culture of fairness in designing, conducting and reporting AI experiments and applications. We propose an approach consistent with the development of AI through history, which allows us to see it not as a computer artifact, but as a socio-technical system. Indeed, the concept of Artificial intelligence (AI) was introduced in 1955 by John McCarthy, during a conference organized at Dartmouth College “*to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it*” (McCarthy et al., 2018).”

In spite of its previous approach as expert systems (Jackson, 1990), the AI is not based on machine learning, initially inspired by the structure of the nervous tissue as described by S. Ramon y Cajal and C. Golgi (Ramon y Cajal & Azoulay, 1955). Warren McCulloch and Walter Pitts (McCulloch, & Pitts, 1943) neural model caught the attention of scientists such as John Von Neumann or Norbert Wiener, among others. In 1958, Frank Rosenblatt constructed the first learning machine, the perceptron (Rosenblatt, 1957). Paul Werbos (Werbos, 1974, 1988) and David Rumelhart et al. (Rumelhart et al., 1986) developed the backpropagation (BP) algorithm that made it possible to train multilayer perceptrons (MLP) constructed with artificial neurons.

The next step was the image recognition by the convolutional neural networks (CNN), were derived from the findings about the visual cortex by David Hubel and Torsten Wiesel (Hubel & Wiesel, 1958). Yan LeCun and coworkers (LeCun et al., 1989) simulated the function of the vision. Indeed, the study of the persistence of the memory was followed by recurrent structures as Hopfield networks or Boltzmann networks (Barra et al., 2012) and the Restricted Boltzmann Machine of Geoffrey Hinton (Hinton et al., 1984; Ackley et al., 1985). The first recurrent neural network (RNN) able to process and retain information on past events was published by Elman (1990) and Jordan (1986). Later, the long-short-term memory (LSTM) network was published by Hochreiter et al. (1997), and used by Sutskever et al. (2011) in text generation.

The first successful and widely used form of generative artificial intelligence for images was the Generative Adversarial Network (GAN) (Goodfellow et al., 2014), but of the main next quantum leaps in this evolution of artificial intelligence came with the inception of the transformers by Vaswani et al. (2017). In this approach, the machine synthesizes the so-called attention mechanisms, which focus on these elements of the input pattern that are relevant to the machine learning task at hand. These devices are at the core of the present large language models (LLM) as the extremely popular ChatGPT. The LLMs have produced an overwhelming quantity of social impact, partly given by the fact that these machines are indistinguishable from human behavior, i.e. they can pass the Turing test, yet they are not trained to provide knowledge or *tell the truth to their best knowledge*, but to produce responses that *look convincing, or good*. This behavior has been scholarly described as bullshit [*sic*] (Rudolph, 2023, Hicks, 2024). Blake Lemoine, a former Google AI engineer, raised ethical concerns about the use of these models (Zhang, 2018; Lemoine, 2022). Geoff Hinton, the most prominent among the *fathers of AI*, left Google and spoke about the dangers of AI (Sloane, 2024). Others, such as Yoshua Bengio, who shares the 2018 Turing Award with Hinton and LeCun, or Max Tegmark, warn about the downsides of AI.

The social controversy comes with confusion and disinformation about what AI really is –see, e. g., Dubosson (2020). It is common to find in debates in the media and forums about any aspects of AI, a common sense that AI is represented by LLMs such as ChatGPT, or image generators such as Microsoft DALL-E, where actual AI, the fundamentals, basic usage, applications, potential and limitations are mostly ignored. Given the importance, potential, and dangers of AI, the introduction of AI literacy in the public in general is needed, as well as the evolution in which AI and ML is studied in higher education to introduce the societal, ethical, and deontological aspects of this discipline at undergraduate levels, so students can acquire AI competencies with the use of their pre-existing skills. This is what we call *Demystifying Machine Learning*.

## Background

In November 2018, WEEF-GEDC ([www.ifees.net/weef-gedc-2018](http://www.ifees.net/weef-gedc-2018)) held the First Global Peace Engineering conference to challenge the mindset of all concerned citizens, academia, governments, industry, NGOs, etc. to incorporate consequences into the design of new solutions (Jordan et al., 2018), also to break silos across existing disciplines and investigate creating new ones so we can address the global challenges by 2030. An outcome of the global conference was the creation of the Peace Engineering Consortium (PEC) in 2018. Transparency and verifiable trusted data are essential.

Another outcome of the Global Peace Engineering Conference is that the University of New Mexico, School of Engineering created the Peace Engineering Minor (PEM), open to all students from all disciplines. The goal of the PEM is to educate *street smart* professionals of all disciplines, not only engineers, who can adapt and pivot quickly based on the context in which they are placed. They must have the necessary tools to be agile, resilient, resourceful and ethical, along with the leadership and communications skills to work across disciplines and cultures. As a required part of this minor, we present a new approach to the study of AI at undergraduate levels containing the concepts of structure, criterion, and algorithm to provide an AI landscape to the students, contextualized over the base of socio-technical systems, where ethical and societal elements are introduced to construct a deontological practices in the development and use of AI.

There is significant effort to introduce AI from different points of view and diversity of students. Lu et al. (2023) study the challenge of establishing AI majors across universities, with emphasis on the problems of the fast evolution of AI. Chan et al. (2023) cite the need of establishing AI education in elementary through high school. Lee et al. (2021) study the development of AI literacy in middle schools as a socio-technical system. Korte et al. (2024) tackle the problem of introducing AI literacy to prepare the population for AI to contribute to the future society. In medicine, Perchik (2023) proposes short course structures to introduce AI education in radiology. The potential use and necessity of professionals skilled in AI are treated in Soai (2024), while Ali (2024) studies the introduction of AI education in pharmacy. Pillay (2018), tackles the need for introducing AI in engineering as a preparation for the so-called 4th industrial revolution.

Flores-Alonso (2023) establishes undergrad AI education which recognizes the existence of different backgrounds and diversity and promotes inclusion. King (2019) studies the inclusion of AI in undergraduate electrical and computer engineering. Southworth (2023) proposes the introduction of AI across the curricula. Stolpe (2024) proposes a framework for AI literacy and socio-ethical understanding of AI, and Zhang (2024) discerns the possible contributions of AI in the development of ideological and political perspectives in higher education.

## Description of the course

The proposed approach is intended to change how AI is introduced. We construct a course to teach AI that is suitable for junior or senior students, taking into account their background and skills. It is possible to teach AI to higher ed students in various areas with limited mathematical backgrounds. We designed the methodology and contents to teach AI as a socio-technical system, so the fundamental ethical and sociological components are naturally introduced.

Demystifying AI is a challenging task, but our experience suggests that it is possible. The approach that we propose is based on the idea that AI is a of the course quest to mimic the human brain, something stated in the press release written by the Nobel Academy related to the Physics Nobel Prize Awarded to AI pioneers Hinton and Hopfield (see [www.nobelprize.org/prizes/physics/2024](http://www.nobelprize.org/prizes/physics/2024)). This approach allows us to teach the very fundamentals of AI with the use of a few standard algebraic concepts that are actually High School level: the dot product and the matrix multiplication. The derivative is needed to develop the backpropagation algorithm, but we start by exemplifying the optimization process by explaining and demonstrating the LMS algorithm. On it, the optimization rule is “update the

weights with the data times the error". This is something that can be generalized to any AI machine.

We can therefore explain the fundamentals of AI learning with basic algebra. The generalization of a training criterion, that holds the previous principle, is understood by the Bayes Rule, and we need to introduce or refresh the concept of probability density, prior, posterior and likelihood. While these concepts are not trivial, students are able to reproduce and interpret the Bayes rule after two weeks of theory and practice. We do not need any other mathematical concepts.

## Prerequisites

The course assumes that the student is in Junior or Senior year. Engineering students need to pass the courses Introduction to Probability, Signals and Systems, and Linear Algebra. The required knowledge for the probability course is restricted to the definition of the probability distribution, conditional probability, and Bayes Theorem. Similarly, the student needs to have been exposed to the definition and practice of the concept of convolution, which is part of a Signals and Systems course, and they need to be able to describe and reproduce the concepts of inner product and matrix manipulations as matrix addition, product, and transpose operator. No other mathematical concepts are required. These concepts are reviewed for all students in class.

## Technological component

The student is exposed to basic Python programming and Matlab. The class sessions contain the necessary demonstrations in Python, also provided in Google® Colab through Canvas, for them to learn through practice the essential mechanisms of programming. In particular, the materials provided in the course include a module demonstrating the basics of Python and Matlab. In that module, students are trained to start a simple Google® Colab notebook or a Jupyter Notebook and perform basic operations, construct a script, designing a function, creating a class, and declaring an object. The construction of libraries or other higher-level abstraction structures is not necessary for this course, but the demonstrations have examples of how to call and use the NumPy, Scikit Learn, and PyTorch libraries. Students can access this module as an additional reference, while the main activities include practice on these concepts.

Each module is organized into a maximum of four sections, which are the structure, criterion, and algorithm. The structures refer to the neuron, with sigmoid and ReLU activation, the MLP, the CNN, the RNN, and the basic attention mechanism. Therefore, only four main deep learning (DL) structures are presented. To describe the operation implemented with these structures, the instructor needs to use matrix addition, product, transpose operator, and convolution. The introduction of these operations is aided by simple demonstrations in Google Colab.

The only criterion to train a learning machine that will be explained is the Maximum A Posteriori. The student will be introduced to the concept of maximum likelihood, prior and posterior through the simple review of the Bayes Theorem, this is the most advanced theoretical concept used in class. Therefore, the instructor will dedicate a sufficient amount of time for the student to be able to repeat, reproduce, and explain these three concepts.

The only algorithm to be introduced and utilized in the class is the BP. This concept is the most advanced practical one, and it will be introduced after the explanation of the criterion. The instructor will take a significant amount of time to explain and demonstrate it after the introduction of the MLP. Practice about the specific high-level abstraction software that facilitates the automatic application of the BP is provided. During the introduction of the rest of the structures, the instructor proves that the BP can be applied to them without modification and, therefore, an analysis and development of this algorithm with these structures is not necessary.

## Socio-Technical component

The course contains three specific modules that explain AI from a societal point of view. The first module, called *AI and Society*, analyzes, from the eyes of experts in these areas, the present and

anticipated social impact of the penetration of AI in society. The necessity of regulation (Amariles & Baquero, 2023), the concepts of disinformation, bias, privacy violation, and other downsides of AI (Watch & Doung, 2023) are summarized. The second module introduces the idea of *AI as a socio-technical system* versus a computer artifact, and the implications of both models in the design and treatment of malfunctions that may have a societal impact (Baxter & Sommerville, 2011; Johnson & Verdicchio, 2017). The third module uses the previous concepts to introduce the idea of intended, embedded, and perceived values (Johnson & Verdicchio, 2024) and strategies of design to introduce *values in AI* (Friedman et al., 2019) as defined by the IEEE (Shahriari & Shahriari, 2017) and the European Union (Smuha, 2019).

## Ethical experimental practices

All the experimental assignments are introduced in the class. Besides the demonstration of the necessary experimental tools, the students are reminded of the mandatory elements of a report, including introduction, theory, experiment description, presentation of the results, and discussion. This practice intends to introduce a culture of fairness in experimental reporting. The necessity of these components of the report is justified in terms of the originality and reproducibility of the work. This is assessed as one of the expected student learning outcomes (SLOs).

## Web tools and bibliography

The materials of the course, including slides and class summaries, as well as the assignments, class rubrics, syllabus, and other resources are provided through the Canvas<sup>®</sup> Learning Management System. The contents of the syllabus are approved by the Senate of the University. The data needed for the experiments is also included in Canvas. The students can currently run all the experiments in Python through Google<sup>®</sup> Colab or in Matlab<sup>®</sup> Online through a University License. Therefore, students do not need to install any software on their computers. The instructor strictly needs a computer connected to the internet with a portable file document (pdf) file viewer, a web browser, and a connection to a projector. We recommend the student make use of online tools to edit in LaTeX, such as the free versions of Overleaf<sup>®</sup> or Papeeria<sup>®</sup>. Additionally, all materials will be published under a public license in a GitHub repository. Mandatory references consisting of a corpus of articles covering the socio-technical and ethical aspects of the class are provided in Canvas. A supplementary book in deep learning is recommended but not mandatory. The recommended books are Zhang et al. (2023) (online at <https://d2l.ai/>), Bishop & Bishop (2023), or Martínez-Ramón et al. (2024).

## Organization of the class and activities

The class has a classic structure of short lectures, demonstrations, practice groups, and group and individual assignments, which are the base of the assessment. Each module has a single homework assignment. Modules 1, 3, 5, and 7 contain individual practical experiments with the explained structures, while modules 2, 4, and 6 contain the elaboration of essays about the explained topics, supported by the provided bibliography. The possibility of conducting student presentations is under study but not currently included in the corresponding activities due to time constraints. Activities are graded through specific rubrics available to the students. All the group activities consist of reading, summarizing and discussing the key aspects of the articles provided as reference for the different topics of the modules as described below. Usually, students have to form groups themselves and choose the topic and associated paper to discuss in the days prior to the days of the corresponding activity. Prior to the start of the activity, they are assigned a second topic, developed by another group, for them to assess.

The class does not have midterm or final exams. The theory is organized in the following rather usual blocks:

1) The MLP and training practicalities.

1.1 The perceptron

Activity: The perceptron rule

## 1.2 Structure and optimization

Activity: A simple multilayer perceptron

## 1.3 The Backpropagation

Activity: Low-level observation of the BP behaviour

## 1.4 Training practicalities

Activity: Adding momentum to the BP

The first and second sections correspond to the structure and optimization criterion, and the rest are dedicated to the details of the algorithm. The activities of these sections are coded at low level, so not libraries are needed and their purpose is to show at the same time the details of the behaviour of the algorithms and a basic way to program a simple neural network.

## 2) AI and society.

2.1 Group Activity: Downsides of AI: bias, discrimination, job losses, inequity, impact over underrepresented groups and others.

2.2 Group activity: Case studies in regulation. Debate on the necessity of regulation.

Here students take a break on the technical components and start overviewing the downsides of AI in society. This first group activity serves as a model of the other two activities. During the class section, the students need to provide a summary on their choice among the provided papers, which they have previously studied, upload it to Canvas together with a critical paragraph (analysis). They then will assess the summary and analysis of another group, corresponding to the second topic assigned to them, according to the rubric. The assessment is double blind, conducted with Canvas functions designed for this activity. This module has group activities only.

## 3) The CNN.

3.1 Elements of a convolutional neural network. Activity: Convolution of an image.

3.2 Operational elements of a CNN. Activity: Convolution using Pytorch

3.3 Training a CNN. Activity: The LeNet handwritten digit classifier.

As the rest of the modules, the CNN is presented in three parts, corresponding to structure (3.1 and 3.2), criterion and algorithm (3.3). Here the presented formulation is different from the one in the main textbooks, and the design of the formulation and the presentation is such that the students can identify every element of the CNN backpropagation to the backpropagation of the MLP, in order to induce in the students the skills necessary to explain and reproduce in theory the process. The practice does not contain the details of the BP, but instead we introduce a high level of abstraction libraries to reproduce an arbitrary CNN.

## 4) AI as a socio-technical system.

4.1 Group activity: Theory of socio-technical systems

4.2 Group activity: AI as a socio-technical system

Here we introduce, in the first activity, the socio-technical systems as understood prior to the advent of AI, and once its concepts are worked out by the student, we particularize the provided scholarship to the area of AI.

## 5) RNN.

4.1 The Elman RNN. Activity: visualizing the time recursion of an RNN

4.2 Training an RNN. Activity: visualizing the backpropagation through time.

4.3 Long-short memory networks. Activity: LSTM for forecasting a real world signal

In this module the criterion is essentially kept, as the student already interiorized that the criterion is identical to the one used in the two previous structures, and we treat the new backpropagation algorithm in greater detail. We provide real time signals for the students to experiment. For example, students can use the power load records of the ISO New England or the California ISO databases available for experimentation in forecasting the electric load needed for the next day in a given location of the USA. An exciting work could be to perform an actual forecast of the load and assess the accuracy of the prediction during the next day.

6) Values in AI.

6.1 Group activity: Values as defined by the IEEE and the EU

6.2 Group Activity: Intended, embedded and perceived values in AI.

This module is based on the current definitions of values in AI and about mechanisms to design, introduce and assess these values in AI algorithms.

7) Transformers.

7.1 Attention Mechanisms. Activity: Construction and visualization of an attention layer

7.2 Transformers for language. Activity: Describing the structure of a standard transformer

7.3 Transformers for vision. Activity: experimenting with a transformer for vision

A transformer cannot be programmed and trained in short time and, therefore this module is intended to introduce the theoretical concepts and structure of such systems, with the purpose of understanding its behaviour, this is, explaining and justifying the use of the attention mechanism, explaining and interpreting the structure of a large language model at a block level, and explaining how a transformer is modified to deal with images.

## Anticipated Student Learning Outcomes

The evaluation is based on individual and group assignments. The student is provided with a rubric for each assignment, that can be used to self-evaluate their performance. The assignments are currently submitted through the Canvas Learning Management System, provided with a similarity (plagiarism) feature, and a placeholder for detailed feedback, provided together with the rubric. The specific measurable SLOs and the ways to measure them are described below.

**1. Analyze the basics of machine learning from a theoretical point of view.** In each assignment, the students will analyze the results of the experiments with respect to the theoretical developments given in each module. For example, students need to compare an MLP and the CNN algorithms in the assignment *Image Recognition with CNN* and explain the different performances from the point of view of the performance of the two used models. Students will reproduce and summarize the basic theory of both algorithms. Optionally, students can present a paper that identifies the main criterion to construct both machines (similarity between both models), describe both algorithms (differences), and criticize their drawbacks.

**2. Reproduce and implement the algorithms presented in class.** In various assignments, there is a mandatory practical part, where students develop the algorithms introduced in class and conduct experiments that demonstrate their characteristics. They need to analyze the results.

**3. Apply machine learning to real-life problems.** Some of the assignments are related to real data, used to solve common machine learning problems. The student is required to design and develop the right experiment, defend their choices, and evaluate the results.

**4. Describe and criticize the ethical and sociological elements of Artificial Intelligence.** The assignments oriented to evaluate this SLO contain description, interpretation, and criticism of the socio-technical system concepts, regulations, bias, and other risks, and the introduction of values. Students need to properly describe one or more of the following: AI regulation, concepts of discrimination, bias, privacy; interpret one or more of the following: case studies of AI

regulation., AI values; criticize one or more of the following: case studies of discrimination, bias, privacy violation, gender and race issues, impact of AI n minorities and protected communities.

**5. Create original, fair, and reproducible work.** The work has to be original. All materials that are not original must be cited. Both the theory and the results of the experiment must be reproducible, that is, the information necessary for the theoretical developments, production of the data, and reproduction of the results must be present in the assignments.

## Conclusion

The evolution of AI, particularly DL, is partly a quest to imitate the function of the brain, although other approaches, such as Kernel Learning, do not follow this path. The advances in AI technologies have made them usable in practice, and their penetration in the different segments of society is rapidly increasing, which carries a high social and ethical impact. We presented a proposal for an undergraduate course that introduces AI to higher education students with a structured design that allows the learning of the discipline over basic analytical concepts. It balances the learning of the technical and theoretical aspects of AI and the awareness of the social and ethical implications of the implementation and use of AI and fosters ethically acceptable practices in the design and experimentation of AI-related technologies.

## References

- Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive science*, 9(1), 147-169.
- Ali, S. (2024). Transformative education: Nurturing AI proficiency in pharmacy. *Journal of the American Pharmacists Association*, 64(1), 5-6.
- Amariles, D. R., & Baquero, P. M. (2023). Promises and limits of law for a human-centric artificial intelligence. *Computer Law & Security Review*, 48, 105795.
- Barra, A., Bernacchia, A., Santucci, E., & Contucci, P. (2012). On the equivalence of Hopfield networks and Boltzmann machines. *Neural Networks*, 34, 1-9.
- Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interacting with computers*, 23(1), 4-17.
- Bishop, C. M., & Bishop, H. (2023). *Deep learning: Foundations and concepts*. Springer Nature.
- Chan, H. A., Chan, L., Hodgson, P., Tam, H., Cheung, W., & Tsang, K. (2023, November). A Design of Artificial Intelligence Education. In *2023 IEEE International Conference on Teaching, Assessment and Learning for Engineering (TALE)* (pp. 1-8). IEEE.
- Dubosson, M., Fragnière, E., Rochat, D., Sitten, M., & Berdeaux, E. (2020). Confusion between Artificial Intelligence and digitization at work—ignorance or blind trust? *Revue Européenne d'Économie et Management des Services*, 2020(9), 97-118.
- Elman, J. L. (1990). Finding structure in time. *Cognitive science*, 14(2), 179-211.
- Flores-Alonso, S. I., Diaz, N. V. M., Kapphahn, J., Mott, O., Dworaczyk, D., Luna-Gracia, R., & Paredes, A. S. R. (2023, October). Introduction to AI in Undergraduate Engineering Education. In *2023 IEEE Frontiers in Education Conference (FIE)* (pp. 1-4). IEEE.
- Friedman, B., & Hendry, D. G. (2019). *Value sensitive design: Shaping technology with moral imagination*. MIT Press.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26(2), 38.



- Hinton, G. E., Sejnowski, T. J., & Ackley, D. H. (1984). *Boltzmann machines: Constraint satisfaction networks that learn*. Pittsburgh, PA: Carnegie-Mellon University, Department of Computer Science.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1), 106.
- Jackson, P. (1990). *Introduction to expert systems*. Addison-Wesley Longman Publishing Co., Inc.
- Johnson, D. G., & Verdicchio, M. (2017). Reframing AI discourse. *Minds and Machines*, 27, 575-590.
- Johnson, D. G., & Verdicchio, M. (2024). The sociotechnical entanglement of AI and values. *AI & SOCIETY*, 1-10.
- Jordan, M. I. (1986). Serial order: a parallel distributed processing approach. technical report, June 1985-march 1986 (No. AD-A-173989/5/XAB; ICS-8604). California Univ., San Diego, La Jolla (USA). Inst. for Cognitive Science.
- Jordan, R., Agi, K., Maio, E., Nair, I., Koechner, D., & Ballard, D. (2018). Invitation to Shape Peace Engineering from November 12-16, 2018. *WEEF-GEDC*.
- King, S. O. (2019, April). How electrical engineering and computer engineering departments are preparing undergraduate students for the new big data, machine learning, and AI paradigm: A three-model overview. In *2019 IEEE Global Engineering Education Conference (EDUCON)* (pp. 352-356). IEEE.
- Korte, S. M., Cheung, W. M. Y., Maasilta, M., Kong, S. C., Keskitalo, P., Wang, L., ... & Gu, M. M. (2024). Enhancing artificial intelligence literacy through cross-cultural online workshops. *Computers and Education Open*, 6, 100164.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., & Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4), 541-551.
- Lee, I., Ali, S., Zhang, H., DiPaola, D., & Breazeal, C. (2021, March). Developing middle school students' AI literacy. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education* (pp. 191-197).
- Lemoine, B. (2022). Is LaMDA sentient?—an interview. *Medium*.
- Lu, C., & Lu, A. J. (2023, October). Attempts and Challenges in Establishing an Undergraduate AI Major. In *2023 IEEE Frontiers in Education Conference (FIE)* (pp. 1-8). IEEE.
- Martinez-Ramon, M., Ajith, M., & Kurup, A. R. (2024). *Deep Learning: A Practical Introduction*. John Wiley & Sons.
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2018). A proposal for the Dartmouth summer research project on artificial intelligence (1955). *Reprinted online at <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>*.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5, 115-133.
- Perchik, J. D., Smith, A. D., Elkasseem, A. A., Park, J. M., Rothenberg, S. A., Tanwar, M., ... & Sotoudeh, H. (2023). Artificial intelligence literacy: developing a multi-institutional infrastructure for AI education. *Academic radiology*, 30(7), 1472-1480.
- Pillay, N., Maharaj, B. T., & van Eeden, G. (2018, November). AI in engineering and computer science education in preparation for the 4th industrial revolution: A South African perspective.

In *2018 World Engineering Education Forum-Global Engineering Deans Council (WEEF-GEDC)* (pp. 1-5). IEEE.

- Ramon y Cajal, S., & Azoulay, L. (1955). Histologie du système nerveux de l'homme et des vertébrés. In *Histologie du système nerveux de l'homme et des vertébrés* (pp. 2v-2v).
- Rosenblatt, F. (1957). The perceptron, a perceiving and recognizing automaton Project Para. Cornell Aeronautical Laboratory.
- Rudolph, J., Tan, S., & Tan, S. (2023). ChatGPT: Bullshit spewer or the end of traditional assessments in higher education?. *Journal of applied learning and teaching*, 6(1), 342-363.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *nature*, 323(6088), 533-536.
- Shahriari, K., & Shahriari, M. (2017, July). IEEE Standard Review—Ethically aligned design: A vision for prioritizing human wellbeing with artificial intelligence and autonomous systems. In *2017 IEEE Canada International Humanitarian Technology Conference (IHTC)* (pp. 197-201). IEEE.
- Sloane, M. (2024). Controversies, contradiction, and “participation” in AI. *Big Data & Society*, 11(1), 20539517241235862.
- Smuha, N. A. (2019). The EU approach to ethics guidelines for trustworthy artificial intelligence. *Computer Law Review International*, 20(4), 97-106.
- Sohail, N., Puyana, C., Zimmerman, L., & Tsoukas, M. M. (2024). AI in dermatology: Bridging the gap in patient care and education. *Clinics in Dermatology*.
- Southworth, J., Migliaccio, K., Glover, J., Reed, D., McCarty, C., Brendemuhl, J., & Thomas, A. (2023). Developing a model for AI Across the curriculum: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence*, 4, 100127.
- Stolpe, K., & Hallström, J. (2024). Artificial intelligence literacy for technology education. *Computers and Education Open*, 6, 100159.
- Sutskever, I., Martens, J., & Hinton, G. E. (2011). Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)* (pp. 1017-1024).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Wach, K., Duong, C. D., Ejdys, J., Kazlauskaitė, R., Korzynski, P., Mazurek, G., ... & Ziemba, E. (2023). The dark side of generative artificial intelligence: A critical analysis of controversies and risks of ChatGPT. *Entrepreneurial Business and Economics Review*, 11(2), 7-30.
- Werbos, P. J. (1974). Beyond regression: New tools for prediction and analysis in the behavioral sciences. *PhD thesis, Committee on Applied Mathematics, Harvard University, Cambridge, MA*.
- Werbos, P. J. (1988). Generalization of backpropagation with application to a recurrent gas market model. *Neural networks*, 1(4), 339-356.
- Zhang, B. H., Lemoine, B., & Mitchell, M. (2018, December). Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 335-340).
- Zhang, T., Lu, X., Zhu, X., & Zhang, J. (2023). The contributions of AI in the development of ideological and political perspectives in education, *HELIYON*, 9 (3), E13403.

## **Acknowledgments**

We thank the International Federation of Engineering Education Societies, the Global Engineering Deans Council, the Ibero-American Science and Technology Education Consortium (ISTEC, Inc.), ISTEC-ECE, SensorComm Technologies Inc. and the University of New Mexico for their support and trust in us. We also wish to thank Mathworks®.

## **Copyright Statement**

Copyright © 2024 Martínez-Ramón; Jordán; Koechner; García-Necochea; Dacpano and García. The authors assign to the Australasian Association for Engineering Education (AAEE) and educational non-profit institutions a non-exclusive license to use this document for personal use and in courses of instruction provided that the article is used in full and this copyright statement is reproduced. The authors also grant a non-exclusive license to AAEE to publish this document in full on the World Wide Web (prime sites and mirrors), on Memory Sticks, and in printed form within the AAEE 2024 proceedings. Any other usage is prohibited without the express permission of the authors.