White Paper

BUILDING TOMORROW'S AI WORKFORCE: PATHWAYS FROM COMMUNITY COLLEGE TO CAREER

Institute for California AI Policy (I-CAP),
Silicon Valley Leadership Group Foundation

Authors:
Akshay Mendon, Hamed Honari

Table of Contents

Strategies for Building Tomorrow's AI Workforce: Community College to Career

Authors: Akshay Mendon and Hamed Honari
I-CAP Fellows for the Silicon Valley Leadership Group Foundation

Executive Summary

This white paper documents the development of curriculum in Artificial Intelligence (AI) specifically tailored for community college students. Its purpose is to serve as a resource for other educational institutions across California looking to implement similar AI and machine learning curricula.

An initiative of the Silicon Valley Leadership Group Institute for California AI Policy (I-CAP) and the SVLG Foundation, in partnership with the Foothill-DeAnza Community College District, the curriculum described below aims to empower a diverse student body with the skills and knowledge necessary to succeed in the rapidly evolving AI landscape. Equipping students with in-demand AI skills supports workforce development, fosters innovation, and helps meet the growing industry demand for AI talent across Silicon Valley.

Why Community Colleges?
Community colleges play a critical role in addressing workforce shortages and improving educational equity. They serve diverse learners—including first-generation students, working adults, and underrepresented populations—who may not otherwise have access to high-growth fields like AI and Machine Learning (ML). By offering hands-on, industry-aligned training, this curriculum ensures students can develop practical AI skills, preparing them for both immediate employment and further academic opportunities.

In close collaboration with the Foothill College Computer Science and Business departments, I-CAP Fellows Hamed Honari and Akshay Mendon developed AI and ML curriculum in direct response to the Bay Area's critical shortage of AI professionals—an "undersupply" estimated at 13,575 in the Bay Region and 4,511 in the Silicon Valley Sub-Region [1]. To pinpoint the exact nature of the skill gaps stalling workforce development in Artificial Intelligence (AI) and Machine Learning (ML), the fellows began this effort with detailed surveys and interviews with both students and employers. From these efforts, one finding stood out above all others: education must be made more inclusive and practical if a program is to meet the growing AI demands of regional industries.

Program Overview

To provide a comprehensive community college educational pathway, the fellows developed:

1. Two “Principles of AI and ML” Certificates – Offering foundational theory and immediate practical skills in short, focused programs.
2. Associate of Science (AS) Degree in Artificial Intelligence – A robust, transfer-oriented credential designed to equip community college students for both direct workforce entry and seamless transition to four-year institutions.
3. Bachelor’s Degree in AI for Business – A forthcoming program that merges advanced technical coursework with essential business acumen, preparing graduates for leadership roles in AI-driven enterprises.

Course examples highlighted in this whitepaper include *CS F011A: Introduction to Artificial Intelligence and CS F012A: Introduction to Machine Learning*. These courses distinguish themselves in three principal ways:

- Hands-on Learning: Courses integrate real-world projects, such as analyzing San Jose traffic data, to ensure students gain industry-relevant experience.
- Workforce Readiness: Students develop in-demand technical skills—including data preprocessing, model development, and ethical AI implementation—ensuring they are competitive for AI jobs.
- Ethical AI Integration: Responsible AI principles, including bias mitigation and transparency, are embedded throughout the curriculum to promote socially responsible AI development.

Introduction

Key Considerations

The below considerations animate this project:

Figure 1: Key observation and motivations for developing a comprehensive syllabi and AI curriculum

- *Growing demand for AI talent:* According to Gartner report [5], it is predicted that 80% of the workforce will require significant upskilling in AI by 2027, underscoring a critical skill gap. LinkedIn's 2023 Emerging Job Report [6], identified AI roles as one of the fast-growing careers, further emphasizing the high demand for AI-skilled professionals. Further, according to the 2024 Work Trend Index Annual Report from Microsoft and LinkedIn, three out of four people use AI at work [11]. AI-related careers postings have seen an annual growth rate of 37% according to the reports from LinkedIn and World Economic Forum [9,10]. Thus, with companies ranging from startups to tech giants searching for talents with AI skills, community colleges are in a unique position to fill this gap not only regionally but also nationwide [14].

- *Empowering diverse populations:* A diverse AI workforce is essential for ethical and responsible AI development [7,8]. Providing accessible AI education at the community college level empowers students from diverse backgrounds with the skills, knowledge, and hands-on experience necessary to succeed in an AI-driven economy.

- *Gap in AI curricula across institutions:* The field of AI is advancing rapidly, and many AI curricula are outdated or fragmented. Industry reports show that the rapid pace of change in AI means that curricula often lag behind the latest developments. A community college-specific AI curriculum could be designed to be more agile and responsive to industry needs [12-13].

- *Regional workforce needs:* Community colleges play a vital role in training the local workforce. With the rapid rise in demand for AI and Machine Learning professionals across various industries in the Bay Area, developing an AI curriculum to equip and upskill the local workforce is essential.

- *AI across disciplines:* AI is inherently interdisciplinary, with practical applications in economics, healthcare, art, and beyond. Existing programs often lack an interdisciplinary perspective that transcends solely technology. Syllabi samples below will demonstrate integrating multiple focus areas into AI curriculum.

Demand-driven Training

Rather than a single occupation, AI and ML jobs represent a constellation of in-demand roles across various industries.

Recent job posting analytics from Lightcast 2024.4 (Nov. 2023 – Oct. 2024) reveal the wide-ranging demand for AI/ML proficiency in the Bay Area and Silicon Valley. As shown in the tables below, roles such as *Machine Learning Engineer and Data Scientists* are consistently among the top job titles. Meanwhile, positions like *Platform Software Engineer, Site Reliability Engineer, and Firmware Engineer* further illustrate just how broad and multifaceted AI-centric careers have become.

Job Title	Bay Area	Job Title	Bay Area
Machine Learning Engineers	1,978	Marketing Managers	591
Product Managers	1,906	Product Marketing Managers	546
Data Scientists	1,741	Site Reliability Engineers	516
Data Analysts	868	Directors of Product Management	486
Platform Software Engineers	853	Firmware Engineers	485
Sales Managers	780	Systems Engineers	474
Business Development Managers	765	Marketing Product Managers	410
Machine Learning Software Engineers	751	Directors of Business Development	402
Principal Product Managers	695	Territory Sales Managers	317

Table 1: Top Job Titles in Job Postings for Artificial Intelligence for Business Occupations in the Bay Region

Title	Silicon Valley	Title	Silicon Valley
Machine Learning Engineers	1,193	Google Cloud Architects	255
Product Managers	846	Business Development Managers	244
Data Scientists	754	Directors of Product Management	234
Machine Learning Software Engineers	524	Systems Engineers	218
Platform Software Engineers	398	Sales Managers	216
Firmware Engineers	394	Product Marketing Managers	184
Data Analysts	343	Marketing Managers	169
Site Reliability Engineers	330	Technical Product Managers	165

Table 2: Top Job Titles in Job Posting for Artificial Intelligence for Business Occupations in Silicon Valley.

These high-volume postings reinforce the idea that AI expertise is no longer confined to a single role or industry; it cuts across technical, managerial, and innovative functions. For students and jobseekers, this diversity means there is ample opportunity to align AI and ML skills with multiple career pathways—so long as they can access the right training and educational credentials.

Diversity in AI

In addition to addressing the regional AI workforce gap, the team set out to design community college AI curriculum to address the lack of diversity in the AI sector.

AI systems are only as good as the data they are trained on and the perspectives of those who build them. Without diversity in the workforce, AI models risk being biased, leading to costly exposure for businesses and real-world consequences such as algorithmic discrimination in hiring, lending, healthcare, and law enforcement. A 2018 study by MIT Media Lab researcher Joy Buolamwini found that facial recognition systems from major technology companies had error rates of up to 34.7% for dark-skinned women, compared to error rates of 0.8% for light-skinned men, highlighting significant biases due to lack of diversity in training data and development teams. A 2019 study by the National Institute of Standards and Technology (NIST) found that facial recognition systems had significantly higher error rates for people of color, largely due to a lack of diversity among the engineers and datasets used in development. Finally, more globally, McKinsey & Company's 2020 report found that companies with higher gender and ethnic diversity in

leadership outperform their less diverse counterparts by 25% and 36%, respectively, in profitability.

Despite this, the AI workforce remains predominantly male and white—women make up only 26% of AI-related roles, and Black and Hispanic professionals hold less than 10% of technical AI positions, according to the World Economic Forum's 2020 Global Gender Gap Report.

To address these disparities, it is essential to look to community colleges as a key part of the solution. Community colleges serve as accessible entry points for a diverse student population, including women, first-generation students, and underrepresented minorities, and they play a critical role in expanding the AI talent pipeline. Investing in AI education at the community college level can help cultivate the next generation of diverse AI professionals, ensuring that the field is driven by a broad range of perspectives and experiences that lead to more equitable and effective technological solutions.

Why Community Colleges?: A Deep Dive

The National Perspective: Community Colleges as Engines of Opportunity

Community colleges are a cornerstone of higher education in the United States, offering accessible, affordable, and high-quality education to a diverse population. With 1,462 community colleges nationwide, these institutions serve as the primary gateway to postsecondary education for millions of students [3].

Female students make up 57% of the community college population, and 45% are first-generation college students, highlighting the critical role these institutions play in expanding access to education for historically underrepresented groups [3]. By offering a structured pathway into AI, community colleges can ensure that more women, first-generation students, and underrepresented learners gain direct access to one of the fastest growing and most impactful fields in the world.

Beyond access, community colleges also inspire greater public confidence than traditional four-year universities. A recent survey found that 48% of Americans trust community colleges, compared to only 33% for four-year institutions [3]. This trust stems from their student-centered approach, which includes smaller class sizes, more personalized instruction, and strong workforce alignment.

While introductory courses at universities often exceed 300 students, community colleges maintain class sizes of 25 to 35 students, allowing for more engagement, direct faculty interaction, and individualized learning experiences—all of which are essential for complex subjects like AI and Machine Learning [3].

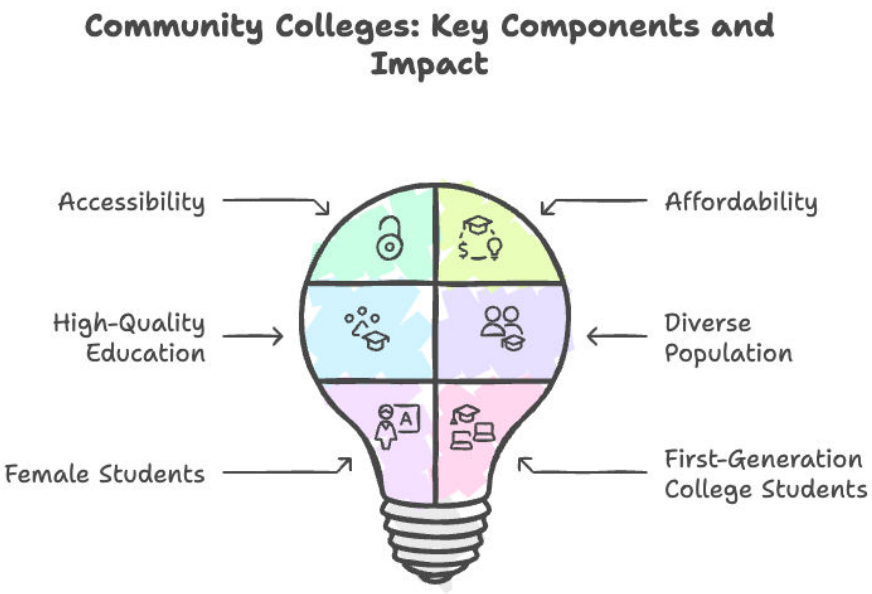

Fig. 2: Community Colleges – Key Components and Impact

California: The Largest and Most Diverse Community College System in the U.S.

California boasts the nation's largest community college system, with 116 institutions serving over 1.9 million students—far surpassing any other state. This system is a model of educational diversity and accessibility, reflecting broader national trends while playing a unique role in preparing the workforce of tomorrow.

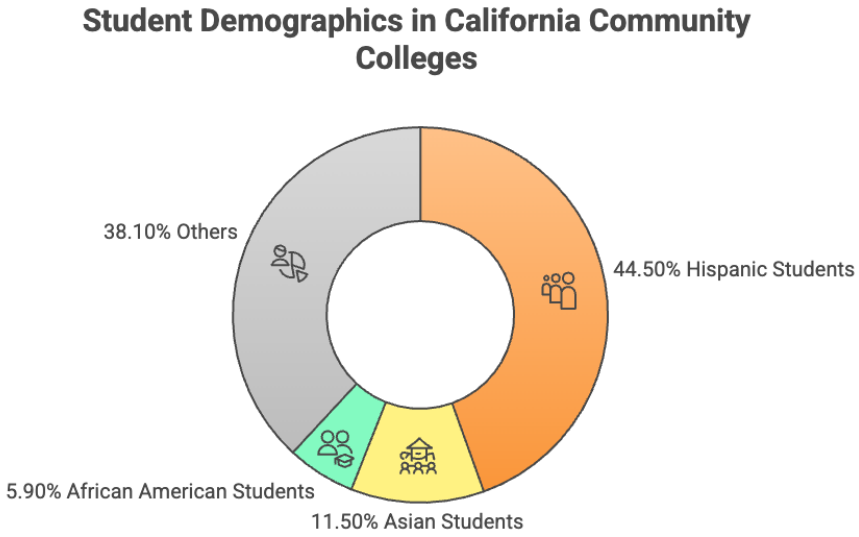

Fig. 3: Student Demographics in California Community Colleges

California community colleges are among the most diverse in the nation, with nearly half (44.5%) of students identifying as Hispanic, alongside large Asian (11.5%) and African American (5.9%) populations [3]. This demographic richness makes them the ideal setting for fostering inclusive AI education, ensuring that underrepresented groups have equal access to emerging technologies.

The diversity of Silicon Valley community colleges is even more pronounced, with 78% of SV community college students identifying as students of color, and 55% women and nonbinary.

The state's close ties to Silicon Valley make it an ideal location for cutting-edge AI education, allowing students to connect with industry leaders, secure internships, and enter high-demand careers faster than if they were required to transfer to a traditional four-year university.

Pathway to Transfer and Career Advancement

While two stand-alone AI and ML courses initially appeared sufficient to boost student engagement, deeper analysis indicated a far greater opportunity: the development of a dedicated Associate of Science (AS) in AI/ML. 49% of students at community colleges plan to transfer to four-year universities [3]. This more extensive approach not only responds to local industry demands but also addresses systemic challenges within the community college system.

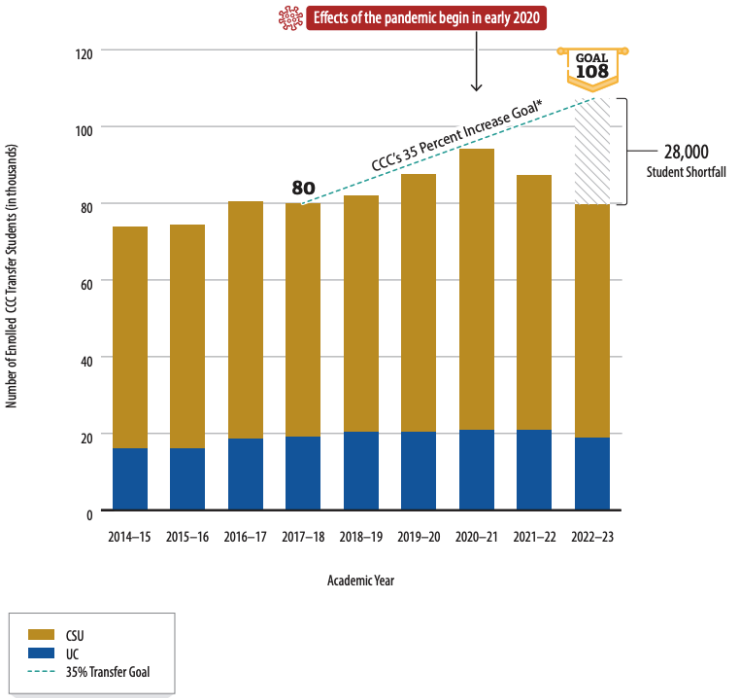

Figure 4: Number of community colleges students in California transferring to CSU and UC [Source: California State Auditor. (2023). California community colleges: A review of enrollment declines and efforts to improve student outcomes.]

As illustrated in Figure 4, the California Community Colleges (CCC) were on track to meet a 35% transfer increase goal, aiming for 108,000 transfers. However, the pandemic abruptly disrupted this progress, resulting in a 28,000-student shortfall [4]. An AS degree in AI/ML can help mitigate that shortfall by:

1. Providing a Clear Credential: Students earn a recognized degree that formally validates their AI/ML competencies, making them more competitive for transfer to four-year universities or direct entry into the workforce.

2. Deepening Technical Mastery: A dedicated degree structure, rather than a few disconnected courses, allows for a more comprehensive curriculum covering data preprocessing, algorithm design, and hands-on projects with real-world datasets.

3. Strengthening Transfer Pathways: By aligning core AI/ML coursework with the admission requirements of UC and CSU systems, students gain a smoother route to bachelor's programs—filling some of the 28,000 transfer seats that have gone unfilled.

4. Reinvigorating Enrollment: Showcasing a distinct, in-demand AS program can attract new learners (including those who paused their education during the pandemic) and retain current students seeking cutting-edge skills.

Designing the Curriculum

The I-CAP Fellow's framework for designing and defining the course objectives prioritized a student-centered approach, emphasizing hands-on experience and practical application. It recognized that some prerequisite concepts might be unfamiliar to a diverse background of students, and therefore integrated "learning-on-the-go" strategies. This involved incorporating practical examples and demonstrations within the lectures to seamlessly introduce and explain these foundational concepts in tandem with demonstrating the applied use cases.

All courses focused on the applicability of AI concepts by showcasing real-world examples and case studies from diverse domains. For instance, in the introductory course on AI, students get to learn about generic types of problems solved by AI such as Natural Language Processing, Computer Vision, Machine Learning, Generative AI, followed by learning about its application and case studies in Business such as marketing and business analytics, healthcare, and art just to name a few example. This approach aimed to make the learning process more engaging and relevant for students, while also preparing them for the evolving demands of the AI workforce.

Given the rapidly evolving nature of Artificial Intelligence, the team also considered it crucial to integrate discussions on ethical considerations and responsible AI practices. By showcasing use cases from various domains, the curricula aimed to familiarize students with the ethical implications of AI and encourage critical thinking and responsible AI development.

Below are key design elements that are woven into the curriculum, further recognizing a diversity of learners and real-world applications that are essential for industry preparation:

- Real-world Relevance: AI concepts are connected to real-world applications that bolster engagement.
- Hands-on Learning: Modules integrate practical skills development through coding exercises, projects, and opportunities for students to work with real-world data.
- Ethical Considerations: Students will discuss the ethical implications of the broader societal impact of AI technology, encouraging critical thinking and bias mitigation.
- Interdisciplinary: Modules reference AI problems spanning across industries to cater to diverse student interests and to provide a comprehensive look into the wide-ranging use-cases of AI.
- Diverse Student Backgrounds: Curriculum recognizes and engages varying levels of prior knowledge and experience.

Curriculum Framework:

The curriculum development team adopted a hybrid approach to the curriculum, combining elements of application-driven and top-down approaches. While the overall structure followed a top-down progression, the fellows incorporated bottom-up approaches for certain foundational concepts to ensure a solid understanding of underlying principles.

The progression the team followed is shown in the following diagram. It begins with a wider lens, a foundation on which to build a broad understanding, replete with an introduction to key concepts and applications of AI, and then a narrowing toward deeper concepts such as Machine Learning.

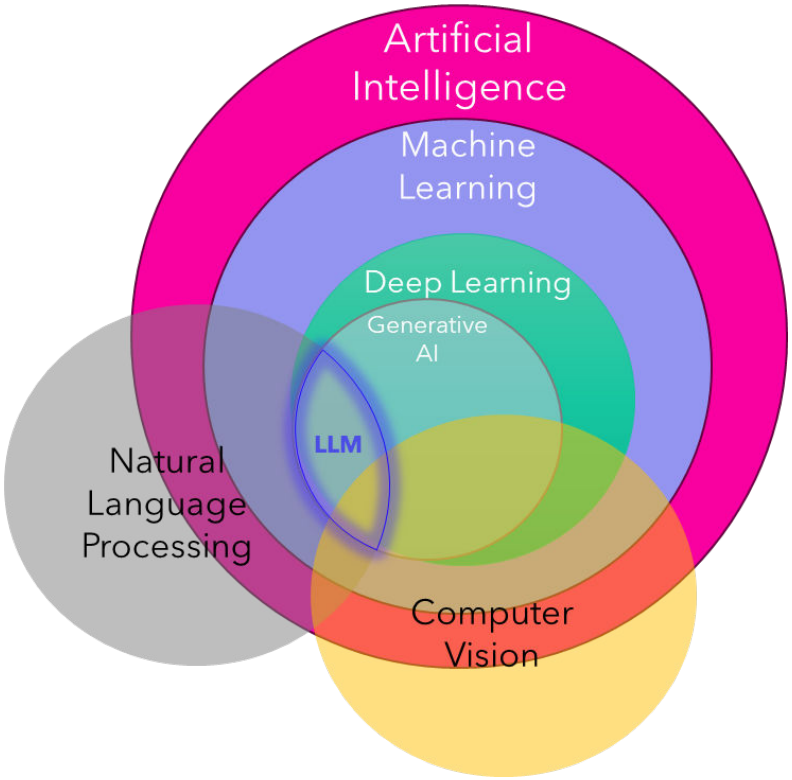

Fig. 5: Schematic of progression and building knowledge in the field of AI

Finally, the team also emphasized a “learn-by-doing” approach in all courses, as hands-on experience is crucial for bridging the gap between academic learning and the demands of the AI workforce. By providing ample opportunities for practical application through coding exercises and projects, the team aimed to equip students with the skills and confidence to excel.

Building CS F011A: Introduction to Artificial Intelligence

Course Design

The I-CAP Fellows and Foothill College CS instructor initiated the syllabus design process with a course on the Introduction to Artificial Intelligence as a foundation and introductory course in the AI curriculum. The goal of the course: to provide students a strong understanding of the wide scope of the field, a conceptual foundation, motivation, and a sense of potential applications. This course will provide students with a broad lens before diving deeper into the technical details, which are discussed next in CS F012A: Introduction to Machine Learning.

This course is designed for undergraduate college students with STEM and non-STEM backgrounds, as well as professionals seeking to dive into the ever-growing and practical field of Artificial Intelligence. It provides learners with an overall view and understanding of the core concepts, terminologies, and tools to further navigate into the practical aspects of AI and be prepared for the workforce. Throughout this course, students are exposed to types of problems solved by AI, and diverse applications of AI in various domains ranging from Business, Healthcare, Art, while simultaneously considering their ethical implications.

Upon finishing this course, students will possess foundational AI knowledge, exposure to various AI problem types, and insight into practical AI applications across diverse fields and real-world scenarios. The course also aims to provide comprehensive content for the students to engage their curiosity, hands-on experiences with lab and quizzes, and tools to prepare them for higher-level coursework.

Course Outline

Weekly Breakdown

Week 1: Introduction to AI

- Topics: Overview of AI and its history, Distinction between AI, machine learning, and deep learning
- Ethical Focus: Awareness of potential biases in AI development from historical perspectives.

- Lab: Introduction to AI tools and environments (basic Python setup, simple AI demonstration).
- Quiz: Key concepts of AI and its historical evolution.

Week 2: Types of AI

- Topics: Narrow AI vs. General AI, Supervised vs. Unsupervised Learning
- Ethical Focus: Implications of general AI development and control.
- Lab: Implementing basic classification models (supervised learning).
- Quiz: Distinction between types of AI and learning paradigms.

Week 3: Generic Types of Problems Solved by AI

- Topics: Natural Language Processing (NLP), Computer Vision, Machine Learning and Deep Learning, Reinforcement Learning, Generative AI
- Ethical Focus: Privacy concerns in NLP and facial recognition technologies.
- Lab: Sentiment analysis with a simple NLP model.
- Quiz: Applications and problem types addressed by AI.

Week 4: AI Applications in Business

- Topics: Case studies in marketing, finance, and operations, AI tools for business analytics
- Ethical Focus: Data privacy and ethical marketing practices.
- Lab: Using AI-based business analytics tools (introduction to a data visualization tool).
- Quiz: Case study analysis on AI in business.

Week 5: AI in Healthcare

- Topics: Applications in diagnostics and patient care, Ethical considerations in healthcare AI
- Ethical Focus: Patient data privacy and biases in diagnostic models.
- Lab: Exploring AI-based healthcare apps (hands-on demonstration).
- Quiz: AI use cases and ethical dilemmas in healthcare.

Week 6: AI in Daily Life

- Topics: Smart assistants, recommendation systems, and personal finance, Interactive demo of AI applications
- Ethical Focus: Surveillance and user data in recommendation systems.
- Lab: Building a basic recommendation system.
- Quiz: Applications of AI in everyday life.

Week 7: AI in Textiles and Art

- Topics: Use of AI in textile design and manufacturing, Generative art: creating artworks using AI algorithms
- Ethical Focus: Copyright and ownership of AI-generated art.
- Lab: Generating art using a simple generative AI tool.
- Quiz: AI in creative industries.

Week 8: Data and AI

- Topics: Importance of data in AI development, Data collection and management practices
- Ethical Focus: Ethical data collection and storage.
- Lab: Data preparation and cleaning for machine learning.
- Quiz: The role of data in AI and its management.

Week 9: Challenges and Limitations of AI

- Topics: Understanding biases in AI systems, Discussion on job displacement and the future of work
- Ethical Focus: Bias mitigation strategies and societal impacts.
- Lab: Detecting and visualizing bias in datasets.
- Quiz: Challenges in AI and bias-related concepts.

Week 10: Future Trends in AI

- Topics: Exploration of emerging technologies, AI's potential impact on society
- Ethical Focus: Long-term societal impacts and ethical foresight.
- Lab: Exploring cutting-edge AI technology.
- Quiz: Future trends and technologies.

Week 11: Ethical Considerations

- Topics: Discussion on AI ethics and accountability, Frameworks for responsible AI usage
- Ethical Focus: Comprehensive review of AI ethics frameworks.
- Lab: Analyzing ethical scenarios with AI decision-making.
- Quiz: Ethical considerations and frameworks.

Week 12: Course Review and Project Presentations

Topics: Review of key concepts, Student presentations on AI projects
Ethical Focus: Incorporation of ethical practices in student projects.
Lab: Presentation preparation and peer feedback.
Quiz: Cumulative review of course content.

Building CS F012A: Introduction to Machine Learning

Course Design

After the first course, the team set out to design the second course in the AI curriculum: CS F012A: Introduction to Machine Learning, a course that stands apart from traditional machine learning offerings by its unique emphasis on real-world applicability, ethical AI integration, and inclusivity. This course is a cornerstone of the program, reflecting a commitment to accessibility, diversity, and workforce readiness.

Unlike many conventional ML courses that cater to students with advanced mathematical or programming backgrounds, CS F012A was crafted to meet the needs of community college students who often juggle multiple responsibilities and come from varied educational and professional experiences. Keeping this in mind, the team designed CS F012A to be highly practical, minimizing abstract mathematics and emphasizing real-world applications. Students work with readily available datasets to learn algorithms through practical application, focusing on solving tangible problems rather than abstract theoretical concepts.

Employers value skills that translate directly to real-world scenarios. This course equips students with the ability to:

- Preprocess and analyze data using industry-standard tools like Python, Pandas, and NumPy.
- Implement core machine learning algorithms such as linear regression, decision trees, and clustering techniques.
- Develop an understanding of algorithmic trade-offs, such as overfitting vs. underfitting and model transparency.

One of the most striking differences between CS F012A and traditional ML courses is its comprehensive integration of Responsible AI principles, a direct response to feedback from employers and students. To honor this, the team ensured that Responsible AI is not just a topic covered in one week but an integral part of every lesson. Each week, students explore themes such as:

- Ethical AI Design
- Transparency and fairness in algorithms

- Addressing bias and racism in AI
- Human-in-the-Loop (HITL) systems for ethical decision-making
- Sustainability in AI and reducing environmental impacts

By weaving these principles into the fabric of the course, students develop a holistic understanding of AI that goes beyond technical skills. Together, the team worked to ensure that students learn to create solutions that are not only innovative but also equitable and sustainable.

Course Outline

Weekly Breakdown

Week 1: Mathematical Foundations for Machine Learning

- Topics: Vectors, Matrices, Probability, Central Tendency, Variance, Eigenvalues, Linear Algebra, and Calculus.
- Ethical Focus: The importance of data integrity and avoiding statistical manipulation in AI models.
- Lab: Introduction to NumPy and Pandas for data manipulation.
- Quiz: Basic statistics and probability concepts.

Week 2: Data Preprocessing and Introduction to Machine Learning

- Topics: Types of Data, Data Quality, Missing Values, Feature Engineering, and Supervised vs. Unsupervised Learning.
- Ethical Focus: Data bias, handling missing values fairly, and responsible feature engineering.
- Lab: Pandas and Data Visualization techniques.
- Quiz: Data attributes, data preprocessing, and ML principles.

Week 3: Regression Analysis & Gradient Descent

- Topics: Linear Regression, Polynomial Regression, Cost Function, Overfitting, Regularization, and Gradient Descent.
- Ethical Focus: Bias in regression models and mitigating algorithmic discrimination.
- Lab: Implementing Linear Regression and Gradient Descent with real-world datasets.
- Quiz: Regression techniques and gradient descent.

Week 4: Logistic Regression, Classification Metrics, and Fair AI

- Topics: Logistic Regression, Decision Boundaries, Precision/Recall, ROC Curves, and Confusion Matrices.

- Ethical Focus: Evaluating AI bias in classification models and fairness in decision-making.
- Lab: Implementing Logistic Regression on real-world data (e.g., loan approvals).
- Quiz: Logistic Regression, classification metrics, and bias evaluation.

Week 5: Support Vector Machines (SVM) & Midterm Review

- Topics: SVM for classification, Linear vs. Non-Linear SVM, Kernel Trick, and Soft Margin SVM.
- Ethical Focus: Using SVM responsibly to avoid reinforcing biases in automated decision systems.
- Lab: SVM classification on an industry dataset (e.g., sentiment analysis).
- Quiz: SVM theory and classification.

Week 6: Decision Trees, Ensemble Learning, and Random Forests

- Topics: Decision Trees, Pruning, Gini Impurity, Bagging, Boosting (AdaBoost), and Random Forests.
- Ethical Focus: Interpretability in AI decision-making and risks of overfitting real-world data.
- Lab: Building Decision Trees and Random Forests using real business datasets.
- Quiz: Decision tree mechanics and ensemble learning.

Week 7: K-Nearest Neighbors & Naïve Bayes Classifier

- Topics: KNN, Distance Metrics, Curse of Dimensionality, Naïve Bayes, Probability Distributions.
- Ethical Focus: Ensuring fair classification in Naïve Bayes models, avoiding misleading assumptions.
- Lab: Implementing KNN and Naïve Bayes with social media sentiment data.
- Quiz: Classification techniques and probability concepts.

Week 8: Unsupervised Learning: K-Means Clustering & PCA

- Topics: K-Means, Elbow Method, PCA, Dimensionality Reduction.
- Ethical Focus: Sustainability in AI and energy efficiency in model training.
- Lab: Applying PCA for dimensionality reduction on real-world datasets.
- Quiz: Clustering methods and PCA applications.

Week 9: Deep Learning - Neural Networks & CNNs

- Topics: Artificial Neural Networks (ANN), Activation Functions, Convolutional Neural Networks (CNN), Backpropagation.
- Ethical Focus: The societal impact of AI automation on job displacement.
- Lab: Implementing a CNN for image classification (e.g., MNIST dataset).

- Quiz: Neural networks and convolutional architectures.

Week 10: Reinforcement Learning & Ethical AI

- Topics: Markov Decision Processes, Q-Learning, Policy-Based RL, and Deep Q Networks.
- Ethical Focus: AI decision-making in autonomous systems and robotics.
- Lab: Implementing reinforcement learning models (e.g., game simulations).
- Quiz: Reinforcement learning concepts.

Week 11: AI Ethics, Safety, and Societal Impact

- Topics: Interpretability, Explainability, Bias in AI, Accountability in AI Development.
- Ethical Focus: Preventing harmful biases in machine learning and ensuring ethical AI deployment.
- Lab: Bias detection in AI models, evaluating fairness in predictions.
- Quiz: AI ethics and bias mitigation strategies.

Pathways

The I-CAP Fellows partnered with Computer Science faculty to develop a structured AI curriculum that expands access to AI education at the community college level. This collaboration led to the creation of certificates and an Associate of Science (AS) degree in AI, designed to provide students with both foundational knowledge and hands-on experience. These pathways ensure that learners—regardless of their prior technical background—can progressively build AI expertise, positioning them for workforce entry or seamless transfer to four-year institutions.

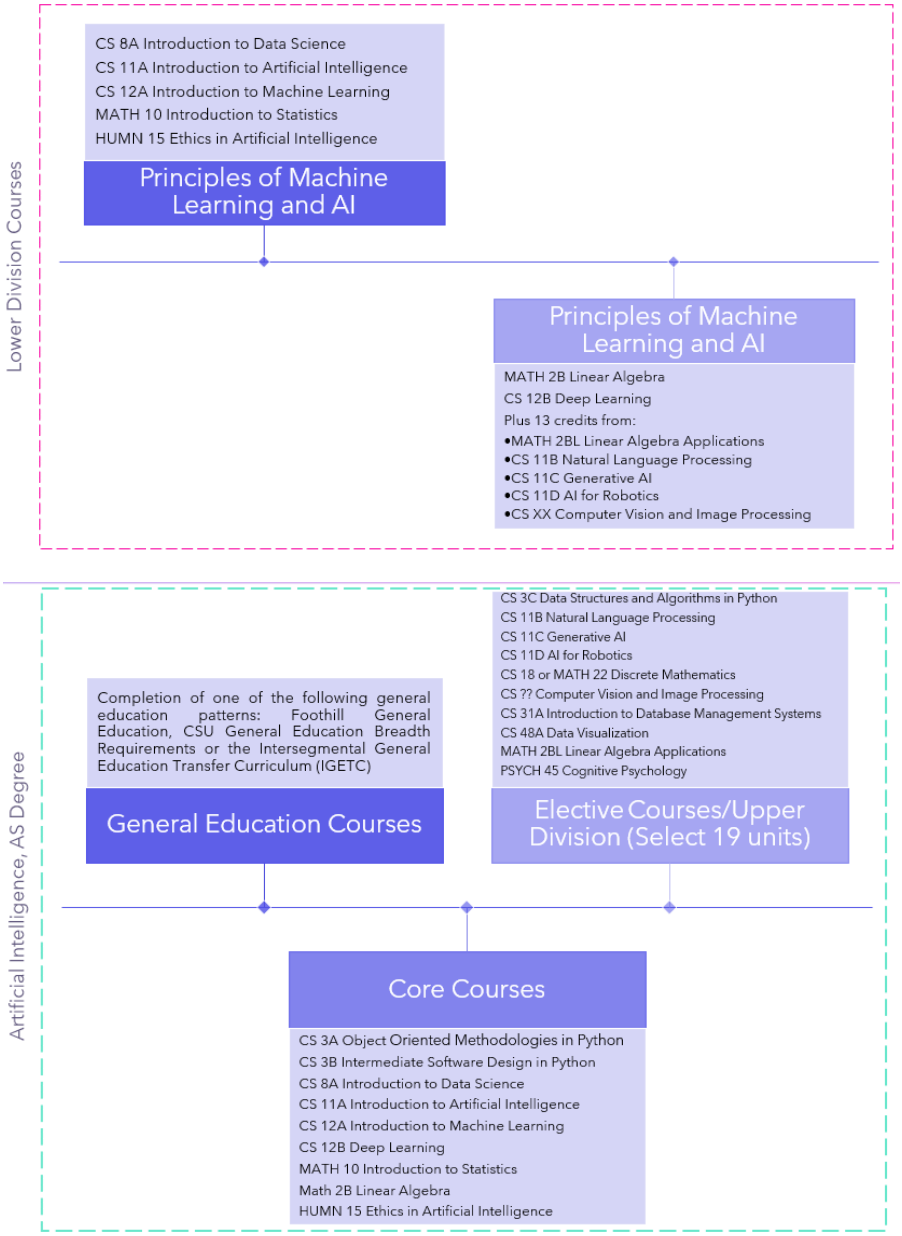

Figure 6. Pathways: Courses for Principles of Machine Learning and AI; Artificial Intelligence AS Degree

Exploring the Development of a Bachelor's Program

Building on the development of stackable certificates and the AS degree in AI, the curriculum team recognized that while these programs provide critical entry points into the field, they alone would not be sufficient to meet industry demands. Employer surveys and labor market data revealed that 94% of AI-related job postings in the Bay Area require at least a bachelor's degree (Table 3) [1]. This underscored a structural gap in available educational pathways, highlighting the need for more advanced credentials to support long-term career mobility. In response, the team expanded its vision beyond short-term certifications and the AS degree, and began to work to develop a comprehensive Bachelor's in AI for Business that blends technical expertise with the strategic and ethical considerations essential for leadership in AI-driven industries.

Education Level	Job Postings	% of Total
High school or GED	2,608	3%
Associate degree	1,714	2%
Bachelor's degree & higher	71,448	94%

Table 3: Education Requirements for Artificial Intelligence for Business Occupations in the Bay Region

Bachelor's in AI for Business

The demand for AI in business is accelerating, with companies across industries integrating AI to enhance decision-making, streamline operations, and create personalized customer experiences. However, while traditional AI degrees emphasize coding, algorithms, and deep learning models, they often overlook the strategic and operational contexts in which AI is applied.

To address this gap, the proposed Bachelor's in AI for Business program is designed to equip graduates with technical proficiency and the ability to implement AI solutions in corporate settings. By combining AI expertise with business acumen, this program prepares students to develop AI models, drive innovation, optimize business processes, and make data-informed strategic decisions. Graduates will be positioned for high-impact roles where understanding market dynamics, customer behavior, and the financial implications of AI adoption is as critical as technical mastery.

Companies today are not just looking for AI specialists who can build models—they need professionals who can leverage AI to generate revenue, improve efficiencies, and enhance user engagement. This dual competency makes graduates of this program significantly more competitive in the job market compared to those with a purely technical AI background. As illustrated in Figure 7, AI-driven business roles

such as Data Analyst, Business Analyst, Product Manager, and Customer Insights Analyst require a blend of AI/ML expertise and industry-specific knowledge, making this interdisciplinary program a key career accelerator.

Fig. 7: AI for Business Career Pathways

To develop this pathway, the curriculum team collaborated with Foothill College's Business Department to design course syllabi that integrate AI and business principles. Key courses include Introduction to Artificial Intelligence, Introduction to Machine Learning, Deep Learning, Natural Language Processing, Large Language Models (LLMs), Advanced Data, MLOps and DevOps, Advanced Algorithms, Generative AI, and Explainable & Interpretable AI.

In addition, the team proposed two new courses—Emerging Technologies and Business with AI Integration—alongside an expanded Business Data Analytics curriculum to further strengthen students' ability to apply AI in practical business scenarios. Figure 8 outlines the extended pathway beyond the Associate Degree, showcasing AI and business-aligned coursework that prepares students for leadership in AI-powered industries.

Figure 8. BS Pathway: The pathway is divided to General Courses (not included in the degree), Lower Division, and Upper Division Courses

Recommendations

This joint effort between I-CAP Fellows, community college faculty, and administrators yielded valuable insights, which are presented below as recommendations for other community colleges seeking to design and implement AI and ML curricula.

1. Aligning AI Education with Workforce Needs

- Identify gaps between workforce development-ready candidates and existing curriculum to create syllabi that address practical, real-world demands.
- Integrate industry case studies and partner with AI-driven companies for project-based learning. For example, Foothill College could collaborate with Silicon Valley firms to offer AI capstone projects, allowing students to work on real-world AI problems in fields such as healthcare, finance, and smart city planning.
- Create feedback loops from industry partners and graduates to refine and update the curriculum to meet changing market demands, particularly in the landscape of California's leadership in AI.

2. Curriculum Design and Flexibility

- Sequence courses to remain flexible and aligned with new technological developments and workforce needs.
- Integrate design thinking into the curriculum to equip students with the skills needed to identify novel applications of AI across industries.
- Facilitate collaboration between educational experts, industry professionals with applied experience, and research and development leaders. This can be achieved through joint curriculum design, industry-embedded faculty programs, and real-world AI projects
- Offer modular learning pathways allows students to gain specialized skills in specific AI areas (e.g., machine learning, natural language processing) in addition to a full degree programs.

3. Accessibility, Inclusivity, and Ethical Implementation

- Proactively design AI programs that are accessible to students from all socioeconomic backgrounds, genders, and geographic regions by offering flexible learning options, targeted outreach initiatives, and financial support mechanisms.
- Embed ethical AI considerations—such as bias mitigation, transparency, and fairness—into every stage of curriculum development to ensure students

understand the broader societal impact of AI.

- Encourage cross-disciplinary approaches, including collaboration with fields like ethics, business, and engineering, to enhance the breadth and applicability of AI education.

References

[1] Bay Region Center of Excellence for Labor Market Research. (2024). *Labor Market Analysis for Artificial Intelligence for Business Occupations: Foothill College.* Foothill College.

[2] Silicon Valley Leadership Group. (n.d.). *Education & Workforce.* Retrieved January 27, 2025, from https://www.svlg.org/education-workforce/

[3] Crown Counseling. (n.d.). *Community College vs. University Statistics.* Retrieved January 27, 2025, from https://crowncounseling.com/statistics/community-college-vs-university/#:~:text=Overall%20Confidence,confidence%20in%20four%2Dyear%20institutions.

[4] California State Auditor. (2023). *California community colleges: A review of enrollment declines and efforts to improve student outcomes.* Retrieved January 27, 2025, from https://www.auditor.ca.gov/reports/2023-123/

[5] Gartner. 2024. *Gartner Predicts 80% of AI Workforce Will Need Upskilling by 2027.*

[6] LinkedIn. 2023. *LinkedIn 2023 Emerging Jobs Report.*

[7] Stanford Institute for Human-Centered Artificial Intelligence (HAI). *Research on AI Ethics, Fairness, and Inclusivity.*

[8] AI Now Institute. *Research Reports on Social Impacts of AI.*

[9] LinkedIn's 2023 Emerging Jobs Report: The demand for AI skills

[10] World Economic Forum. 2025. *Future of Jobs Report 2025: Insight Report – January 2025*

[11] Microsoft. 2024. *2024 Work Trend Index Annual Report from Microsoft and LinkedIn*

[12] Forbes. 2024. *Companies Are Hiring Fewer College Grads – So Why Bother With School*

[13] Forbes. 2024. *Bridging the AI Talent Gap: Revitalizing The Value Of Education Through Employability*

[14] CSET. 2024. *AI and the Future of Workforce Training*

About the Authors

This work is authored by Akshay Mendon and Hamed Honari, professionals with strong backgrounds in Artificial Intelligence, Machine Learning, and Data Science. Their collective expertise spans Natural Language Processing (NLP), Computer Vision, Generative AI, and AI-driven business applications, making them well-equipped to explore the intersection of AI and its practical implementation in industry. With extensive academic, research, and industry experience, they have worked on bridging AI advancements with real-world solutions across multiple domains, including healthcare, finance, energy, and enterprise AI applications.

Hamed Honari

Hamed received his PhD in Electrical & Computer Engineering from Johns Hopkins University. With a strong background in Artificial Intelligence, Machine Learning, and Statistical Signal/Image Processing, he is passionate about software prototyping, proof-of-concept development, and AI-driven methodology research. His primary focus lies in the application of AI to neuroimaging, computational neuroscience, and interdisciplinary research and development.

Before joining Stanford University, he worked as a Data Scientist at the World Bank Group in Washington, DC, leveraging AI for innovative solutions in various industries. His role involved design thinking research and PoV prototyping, including work in Computer Vision, Generative AI (LLMs), and NLP.

Hamed has a proven track record of applying AI research and problem-solving across diverse domains, including healthcare, financial and public sectors, energy, and interdisciplinary engineering. His ability to integrate AI-driven solutions into real-world applications has helped organizations leverage cutting-edge technology for large-scale impact.

Akshay Mendon

Akshay received his Bachelor's in Engineering in Electronics and Telecommunication from the University of Mumbai and is currently pursuing a Master's in Computer Science at Santa Clara University. As a Graduate Teaching Assistant at Santa Clara University, he has assisted in courses on Pattern Recognition and Data Mining and Artificial Intelligence.

He has multiple book chapters and research papers published in IEEE, Springer Nature, and IGI Global. His primary interests lie in Natural Language Processing (NLP), Computer Vision, and Generative AI, focusing on their real-world applications and advancements.

Beyond academia, Akshay is an Editorial Board Member at Primera Scientific Publication House and has completed multiple internships as a Machine Learning (ML) and Deep Learning (DL) Engineer, where he has worked on AI-driven solutions, data modeling, and applied AI technologies. His expertise bridges technical AI development with practical implementation, ensuring that AI innovations are effectively translated into impactful solutions.

About our Partner

Located in the heart of Silicon Valley, the **Foothill-De Anza Community College District** serves the communities of Cupertino, Los Altos, Los Altos Hills, Mountain View, Palo Alto, Stanford, Sunnyvale, and parts of Saratoga and San Jose (district boundary map / trustee areas map).

For more than 60 years, Foothill-De Anza has demonstrated excellence and innovation in academic programs and student services. As one of the largest community college districts in the United States, Foothill-De Anza provides credit classes for about 50,000 students a year. The colleges are active members of the League for Innovation in the Community College, a national consortium of leading two-year institutions.

The mission of the Foothill–De Anza Community College District is student success and educational excellence. The district and its colleges provide access to affordable, quality educational programs and services that develop a broadly educated and socially responsible community that supports an equitable and just future for California and the global community. Every member of the district contributes to a dynamic instructional and learning environment that fosters student engagement, equal opportunity, and innovation in meeting the various educational and career goals of our diverse students.