AI and the Future of Skills

Understanding the educational and work implications of AI and robotics

AIFS Project Brief Summer 2022

Artificial intelligence (AI) and robotics are major breakthrough technologies that are transforming the economy and society. To understand and anticipate this transformation, policy makers must first understand what these technologies can and cannot do.

OECD
BETTER POLICIES FOR BETTER LIVES

Overview of the project

The OECD launched the Artificial Intelligence and the Future of Skills project to develop a programme that could assess the capabilities of AI and robotics and their impact on education and work.

Considering both expected future development of AI and existing research, AI and the Future of Skills project aims to develop a new set of measures to serve as a foundation for research and policy on how AI and robotics will transform skill demand and educational requirements in the decades ahead. It addresses the following concrete questions:

> **What human capabilities will be too difficult for AI and robotics to reproduce over the next few decades?**
>
> **What education and training will be needed to allow most people to develop some work-related capabilities that are beyond the capabilities of AI and robotics?**

To anticipate and adapt to developments in AI, it is crucial for policy makers to have a solid understanding of what AI can do and how its capabilities and limitations compare to those of people. Understanding this comparison is also important for AI researchers as they work to improve their technologies.

The project builds upon a successful pilot project that offered a novel way to measure AI capabilities and connect them to skills needed for work and education. Using the OECD's Survey of Adult Skills (PIAAC), this study showed that 62% of workers in OECD countries use literacy, numeracy and problem-solving skills on a daily basis at work at a level AI systems were able to perform in 2016 (Elliott, 2017).

These results serve as strong motivator for the project to build valid measures that translate AI and robotics capabilities into terms that are meaningful for policy and illustrate the implications for education and training.

The importance of assessing AI capabilities

A fast evolution of AI

In the 18th and 19th centuries, the development of industry and machine manufacturing led to a great shift of skill demand, causing workers to migrate from agriculture to manufacturing. The 20th century saw a comparable shift from manufacturing to services. Today, AI is poised to bring another revolution, transforming skill demand with a technology that is becoming capable of reproducing a much broader range of human skills than the simpler technologies that brought the earlier revolutions. This is already leading to substantial impacts on work, with more profound impacts yet to come. Almost everyone in today's advanced economies is familiar with systems that use AI in some way, such as online shopping, personalised recommendation, instant search across the Internet's vast store of humanity's knowledge, and cars that sense the other cars around them. But AI systems will soon have much greater capabilities.

Major improvements in AI technology, such as the development of machine learning algorithms, are transforming our daily lives with new and powerful technologies. Machine learning systems that learn from data, identify patterns

or make autonomous decisions are found in many of the technologies we use every day, from smartphones to cars to the homes we live in.

However, advances in AI go well beyond the everyday applications that have been widely discussed the past few years. The first AI and the Future of Skills publication (Vol 1: Capabilities and Assessments) describes recent developments in AI capabilities and how they relate to fundamental human cognition and skills (Forbus , 2021, Chapter 2 in OECD, 2021[1]):

» **Deep Learning** systems are at the core of tasks such as automatic translation, facial and speech recognition and image captioning. They are inspired by human biology, but do not necessarily share the same learning models as humans. For example, the systems require massive data sets for training, whereas humans are often able to learn incrementally from single data points.

» **Knowledge graphs** are used to represent networks of real-world facts – such as events, objects or concepts – as well as relationships between them. Knowledge graphs are applied in systems such as recommendation systems or research tools. Concretely, knowledge graphs are used when searching on Google for a new restaurant or looking for new contacts on LinkedIn.

» **Reasoning** systems in AI are capable of in-depth and fast sequential reasoning. Applications of AI reasoning are found for instance in the search for programming errors and for automatic quality control systems, such as those involved in cars, robots, and banking systems.

Each of these three streams of development in AI can be described as genuine revolutions that are leading to major shifts in our economies. At the same time, however, each of these AI capabilities is restricted in important ways compared to the corresponding capabilities in people. Learning systems in AI have only limited ability to learn concepts and skills, which are essential in human development and education. Knowledge systems in AI are limited in their ability to describe the history of the AI system itself (personal knowledge) and what it knows (meta-knowledge), types of knowledge that are essential to higher-level human thinking. And reasoning systems in AI are limited in their ability to reason about everyday situations (common sense) or to recognise that a situation suggests a problem that needs to be solved.

These are substantial limitations in the capabilities of AI systems that cause AI to fall far short of human capabilities in essential ways. Related to these distinct limitations, AI is also currently unable to integrate these different aspects of intelligence. A fourth revolution related to that integration,

1. Forbus, K. (2021), "Evaluating revolutions in artificial intelligence from a human perspective", in *AI and the Future of Skills, Volume 1: Capabilities and Assessments*, OECD Publishing, Paris, https://doi.org/10.1787/004710fe-en

Integrated Intelligence, is now just beginning and is attempting to build AI systems having agency that can handle complex problems and operations by themselves.

As a result of the current fundamental limitations in AI capabilities, AI can fail in ways that seem surprising. For example, an AI Go player can easily beat the best human players, but then find it impossible to adapt to play on a board of a slightly different size.

This brittleness in performance is a hallmark of current AI systems, which researchers are working to overcome. Being able to properly communicate with users, adapt to unexpected situations and propose creative solutions remain important challenges for AI. As researchers explore and begin to overcome these challenges, AI capabilities will become increasingly general. The implications for work and education will be transformational, as machine intelligence starts to take on more of the characteristics of human intelligence.

From AI capabilities to understanding their impact on education and work

As AI's capabilities expand, it is critical to understand how they will transform work and education. Observers commonly argue that humans will be essential for many roles for a long time into the future because of AI's limitations. But what are the implications resulting from those limitations – now and in the near future – for work and for the education that prepares people for work?

» Will AI have a greater role in some industries than others? How will the pattern of this role shift according to evolving capabilities?

» Will AI have a greater effect on low-skill or high-skill workers? Younger or older workers? Workers with more or less education?

» Which types of education best prepare people to work together with AI and which types of education best prepare people to adjust to new tasks, as older tasks are performed by AI?

» How will each of these answers change over the next decade or two?

The answers to these questions depend on how the capabilities of AI evolve. Today we are accustomed to AI systems that are relatively inflexible: they fail to adapt when circumstances change and mindlessly ignore problems that seem obvious to us. But researchers are working to make AI more flexible and responsive. Building on the work of the 20-year roadmapping project of the Association for the Advancement of Artificial Intelligence (AAAI), the next page box on Emerging AI capabilities illustrates some ways that AI could develop over the next two decades and provides a vignette illustrating one of the ways such increasingly powerful AI might be used.

To better understand AI's impact on societies, it is useful to measure AI capabilities in ways that relate to human capabilities and how they are used in the world of work. Such measures will help us understand the potential for replacing workers. They will also help to understand how current jobs will be transformed and how new jobs will be designed.
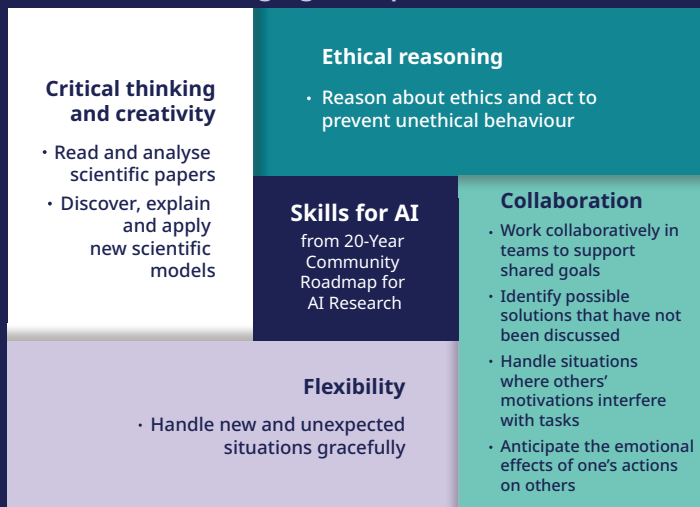
Looking at how jobs are and will be transformed by AI also has direct implications on how education should evolve in response. As the objectives of education change, the processes involved in education – such as learning objectives, types of training or educator requirements – will need to be considered and adapted carefully as well.

What about the future?
Lessons from AAAI U.S. Roadmap

The integration by AI systems of more complex skills such as **collaboration, critical thinking, creativity or ethical reasoning** are in the scope of AI research, and are believed to be feasible by 2040 (Gil and Selman, 2019[1]). Those believes lead experts to imagine hypothetical scenarios of our future lives with AI systems.

Emerging AI Capabilities

Critical thinking and creativity
- Read and analyse scientific papers
- Discover, explain and apply new scientific models

Ethical reasoning
- Reason about ethics and act to prevent unethical behaviour

Skills for AI
from 20-Year Community Roadmap for AI Research

Collaboration
- Work collaboratively in teams to support shared goals
- Identify possible solutions that have not been discussed
- Handle situations where others' motivations interfere with tasks
- Anticipate the emotional effects of one's actions on others

Flexibility
- Handle new and unexpected situations gracefully

Source: Gil, Y., & Selman, B. (2019), pages 24, 27, 29, 48-50, 54 and 68.

1. Gil, Y., & Selman, B. (2019). A 20-year community roadmap for artificial intelligence research in the US. *arXiv preprint arXiv:1908.02624.*

Student Jody

Jody is a middle school student in rural Wyoming who has grown very interested in insects. Jody discusses her classroom lessons with her personal AI tutor, using it as a source of anytime/anywhere learning about insects in her farm and biology in general. The system draws on open educational resources to find interesting questions and topics to discuss, building on its extensive knowledge of what Jody knows in order to better challenge her in productive ways. When her family's crops are affected by an infestation, she sits with her parents and her tutor to read about what the pest could be. She asks the system detailed questions about different species, and narrows it down to three. She works with the system to learn more about the candidate species, then comes up with an experiment to determine which one it is, and discovers the culprit insect. When Jody wonders if her crop pest discovery would be a good start for a science fair project, the system helps her plan it and identify potential roadblocks, as well as ways to work around them. It suggests to her parents that they and she attend some local mentoring events, including a special museum exhibit and a talk by an astronaut. It also helps her find high school and college peers from her county studying biology. Ten years later, Jody is an accomplished veterinarian.

Can AI read?
A pilot assessment using an international education test

An OECD pilot study (Elliott , 2017[1]) illustrated how AI capabilities with respect to the core human skills of literacy and numeracy might be assessed with the OECD's international Survey of Adult Skills test (PIAAC). An update of this work in 2021 suggests how AI capabilities have changed in recent years.

The two studies asked a group of computer scientists whether AI could answer the same questions from the PIAAC test that adults need to answer. The figure below shows that, according to computer experts, in 2016, successful performance of AI on literacy is expected to range from 90% on the easiest questions (at Level 1 and below) to 41% on the most difficult questions (at Level 4 and above). The expected performance of AI systems decreases as the question difficulty increases, as also occurs for humans. In 2021, although the expected AI performance on the easiest questions remained similar, the
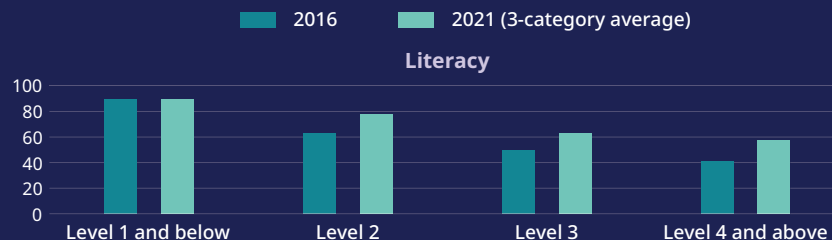
expected performance for the more difficult questions levels increased by about 15 percentage points.

When adults take the PIAAC literacy test, the performance level they are assigned to indicates they can answer two-thirds of the questions of that level correctly. In this sense, AI capabilities in literacy assessed by experts resemble the performance of adults at Level 2 in 2016. In 2021, AI is closer to adults performing at Level 3 in literacy. This improvement in literacy performance of AI is consistent with major developments in natural language

processing over the past few years.

A change of AI performance in literacy from Level 2 to Level 3 has important implications for the portion of literacy-related tasks in the workforce that AI might be able to perform. In OECD countries, roughly 52% of the workforce perform above Level 2 in literacy on the PIAAC test and so would be able to "outperform" state-of-the-art AI in 2016 (Elliott , 2017, p.32[2]). However, only 13% of the workforce has literacy levels greater than Level 3 on PIAAC, the expected level of AI performance in literacy in 2021.

1, 2. Elliott, S.W. (2017), *Computers and the Future of Skill Demand*, OECD Publishing, Paris, http://dx.doi.org/10.1787/9789264284395-en

Assessment of AI capabilities on literacy using PIAAC items

Legend: 2016 ■ 2021 (3-category average) ■

Literacy

Note: Expert ratings of computer capabilities to answer literacy questions, averaged with Maybe=50%, by level of PIAAC question difficulty.

What has been done so far:
Existing taxonomies and experts' recommendations

The first volume from the AI and the Future of Skills project, Capabilities and Assessments, was released in November 2021. This report lays the foundation for a robust methodology to assess the capabilities of AI and compare them with human capabilities. Drawing from computer science and many areas in psychology, the report describes key difficulties in assessing AI and various taxonomies of skills and types of assessments that might be used to describe and measure AI capabilities. The contributing experts each formulated recommendations for the project, providing crucial elements for building a comprehensive approach to the project's goals.

The report illustrates a set of taxonomies of human skills and corresponding tests coming from different domains of research. Taxonomies stemming from cognitive and clinical psychology present tests of cognitive, social, collaborative and emotional abilities. Work in education has developed tests focusing on core human skills such as literacy, numeracy, problem solving and creativity. Industrial and organisational psychology has developed taxonomies classifying occupations by work tasks and required skills and competences. Concrete assessments in the latter domain can be taken from a variety of vocational and occupational tests.

Alongside suggestions of existing taxonomies and tests from different domains, the experts also highlighted fundamental differences between humans and AI. These differences need to be carefully addressed to provide valid measures of AI capabilities and various approaches were proposed to that end.

Educational Research and Innovation

AI and the Future of Skills, Volume 1
CAPABILITIES AND ASSESSMENTS

OECD
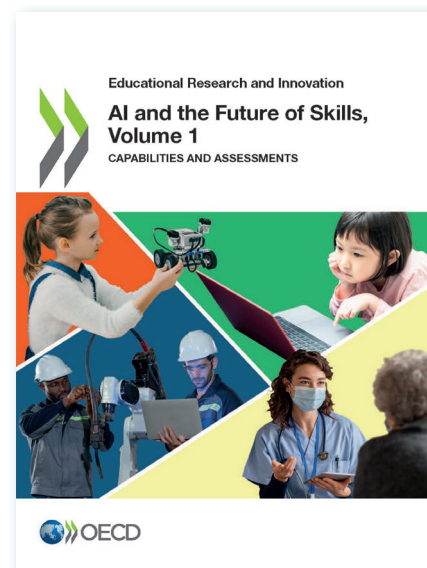
Including measures of basic cognitive skills that most humans have but not AI systems

In the absence of severe disabilities, human adults share many low-level skills that AI systems often do not have, such as abilities to navigate in a complex physical environment, use basic language and understand many basic rules of the world. Although familiar tests in educational or work settings never assess these low-level skills, they are assessed in tests used in the fields of animal cognition and child development, and some researchers are now applying the same approaches to assess the low-level skills of AI systems.

How many bottles are there?

	Humans	AI
Computation skill	Assessed	Assessed
Image recognition	Assumed	Assessed

Using and meaningfully combining multiple assessments to account for fundamental differences between human and AI capabilities

Existing tests for humans may be misleading when used to assess AI. Human tests usually take for granted that the test takers all share general features of human intelligence, such as image processing and language understanding, and the tests often require these capabilities without considering them part of the test. However, when an AI system is given the same tasks, these assumed features can be more difficult than the skills the test intends to measure. For example, a numeracy test question that involves counting bottles in a picture poses a simple problem of quantitative reasoning—the act of counting—but a potentially more difficult problem of image processing to identify the bottles in the picture. Since almost all tasks require complex combinations of capabilities, valid assessments of AI capabilities will require a combination of assessments of isolated skills and assessments of complex tasks involving multiple skills in combination.

Towards a novel approach

Using tests to cover existing skills

Considering both existing taxonomies of human capabilities and the recommendations from experts, the project is developing a novel approach to assessing what AI can do that draws on different contexts and types of tests to provide complementary perspectives on AI.

Covering skills as comprehensively as possible is an essential condition for understanding the impact of AI on the world of education and work because it cannot be assumed that capabilities that appear together in humans will also appear together in AI systems. Evaluating AI performance on occupational and educational tests, as well as using benchmarks and competitions developed specifically for AI, will allow the project to cover both isolated and more complex skills. The occupational and educational tests provide connections to the important applications and implications of AI

systems, with the complexity of skills demanded in the real world, while the benchmarks and competitions that computer scientists use to test and refine their AI systems are specifically shaped to focus on AI capabilities at their current stage of development.

Assessing AI Capabilities

As the figure illustrates, the different tests expected to be used aim to cover a great part of all existing skills that are meaningful for education and work. Although benchmark and competition tests have not yet been systematised and aggregated, they measure AI performance in well-specified tasks, often in relatively narrow skill domains. However, these benchmark and evaluation tests do not cover all human skills and need to be complemented by occupational and educational tests. The former provides a link from AI to job performance (e.g. how to put a wheel on a car), and the latter provides a link from AI to important general skills necessary to live in modern societies, such as literacy and numeracy, as well as the specialised skills developed in formal education in many technical domains, such as the sciences and many professions. But as the figure also shows, the various tests are not fully distinct and can present some
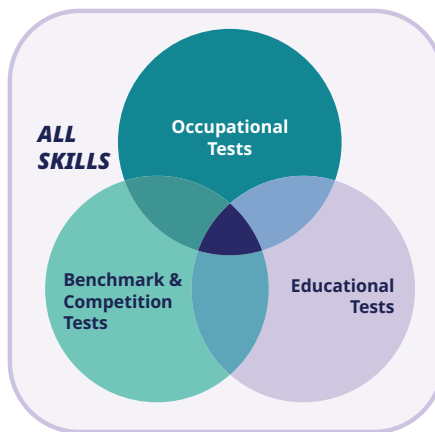
overlaps in the skills they are assessing. For example, many occupational tests measure domain knowledge requiring basic competences acquired in education, such as reading a plan on how to put a wheel on a specific type of car. Similarly, benchmarks and evaluation tests will sometimes use tasks that have been adapted from occupational or educational tests.

The project's long-term objective is to develop a set of indicators that describe AI's capabilities and their implications across the full range of relevant skills, building on existing tests, while making use of the overlaps and complementarities occurring between the different types of tests. The goal is to communicate AI's current capabilities to non-technical audiences to develop a shared understanding of the current limitations of AI systems, the implications of those limitations for the ways that people and AI systems will work together, and a set of indicators that will signal when those AI limitations change.
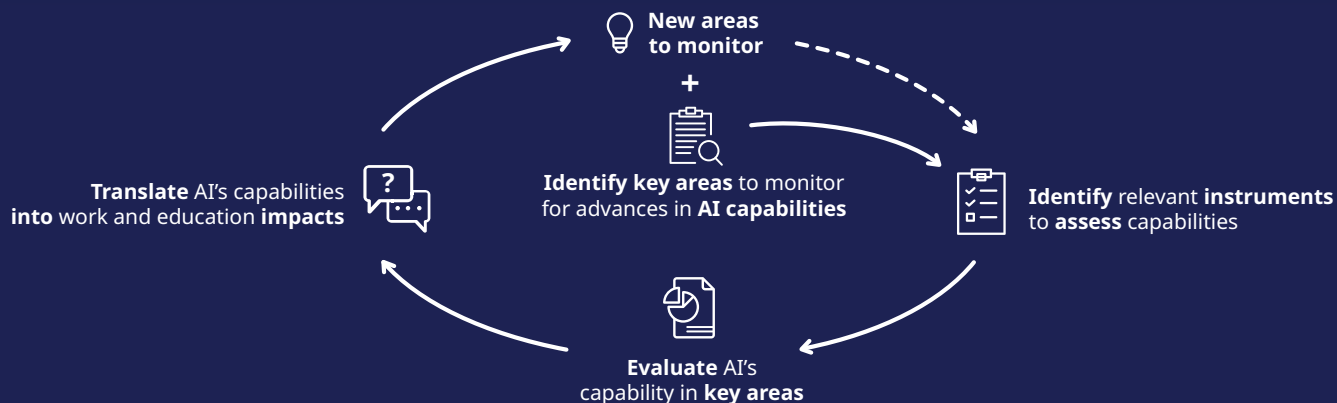
Project methodology

The project is relying on an iterative strategy involving valid, reliable and meaningful measures of AI capabilities.

Step 1 Identifying key areas to monitor for AI advances

AI systems involve many different capabilities, going from basic cognitive skills to increasingly complex, transversal and domain-dependent skills that reflect the requirements of complex contexts. In order to assess AI capabilities and compare them to human ones, the project will build a comprehensive taxonomy of those capabilities — across domains such as language, problem solving, collaboration and movement — and identify key areas to monitor at any time for consequential advances in AI capabilities.

Step 2 Identifying relevant instruments to assess each key area

The project will then identify the types of tests and tasks that can assess AI capabilities in the key areas and provide a way of understanding the relationship between AI and human capabilities in these areas. This selection can draw on a wide range of tests developed in different areas of human psychology and computer application.

New areas to monitor

+

Identify key areas to monitor for advances in **AI capabilities**

Identify relevant **instruments** to **assess** capabilities

Evaluate AI's capability in **key areas**

Translate AI's capabilities **into** work and education **impacts**

Step 3 Evaluating AI's capability in key areas

This crucial step aims at evaluating AI performance on the previously identified tests and tasks assessing the different capabilities. The project is exploring two different ways of doing this: with expert ratings of expected AI performance and with empirical results of current AI systems (usually results available in the research literature). Each of these approaches has strong and weak points; an evaluation approach that combines them can potentially provide the best basis for evaluation.

Step 4 Translating AI's capabilities into work and education impacts

As highlighted by experts in the project's Capabilities and Assessments report, the fundamental differences between human and AI capabilities often make it misleading to compare them in the abstract. For this reason, the final step of the assessment will translate current AI capabilities into their likely implications for the world of work and education. This final step will also provide policy makers with the starting point for considering crucial education, labour and social policies that can help individuals and society anticipate and adapt to new AI capabilities. To make this translation from AI capabilities to their work and education impacts, multiple approaches have been suggested by experts and will be explored. This work will include a process of envisioning the redesign of some work tasks to take advantage of new AI capabilities and identify new ways that humans and AI are likely to work together.

Iterating Refining the key areas and assessment tasks and repeating the assessment and analysis

The previous steps will lead to a first understanding of key AI capabilities, how they compare with humans and how they will affect work and education. The project's ambitions to become an ongoing programme includes iterating and innovating its assessments. Since the pace for AI technological advances is for the moment unlikely to decrease, the project expects that tests in new areas will become relevant over time and tests in some existing areas may become unnecessary if AI capabilities become so advanced that the problems are essentially solved. To track these changes and provide a continually evolving picture of the most important new AI capabilities, the assessment will need to be repeated and refined over time.

Work activities

| | | Review of skill taxonomies and tests | | Experiments with tests and approaches | | Systematic assessment | | Implications for education and work, and reiteration... |

| 2019 | 2020 | 2021 | 2022 | 2023 | 2024 | 2025+ |

Reports

| | | | Technical report Volume 1: Capabilities and assessments | Report: PIAAC update | Technical report Volume 2: Assessment methodology | | Reports on applications (see following box for an example) | |

Key achievements and links:

Reports

» **Project Pilot Study:** *Computers and the Future of Skill Demand* (Stuart, 2017)   October 2017
Link: https://www.oecd-ilibrary.org/education/computers-and-the-future-of-skill-demand_9789264284395-en

» **Methodology report:** *AI and the Future of Skills, Volume 1: Capabilities and Assessments*   November 2021
Link: https://www.oecd-ilibrary.org/education/ai-and-the-future-of-skills-volume-1_5ee71f34-en

Conferences

» **International Conference AI-WIPS**
Session: Benchmarks and competitions: How do they help us evaluate AI?   February 2021 / February 2022
Link: https://oecd-events.org/2022-ai-wips/session/0a762073-4d83-ec11-a507-a04a5e7d20d9

» **Session:** Can AI Read and Count? Assessing Computers with an OECD Test
Link: https://oecd-events.org/2022-ai-wips/session/c876be7d-4d83-ec11-a507-a04a5e7d20d9

What's next:

» 5-year update Pilot Study report   upcoming
» Methodology report: Volume 2   upcoming

What about the future?
A potential application from AAAI scenario

Student Jody
What possible impact of AI capabilities on work and education?

Based on the AAAI Roadmap futurist vignettes on student Jody, AI systems might be able to present strong planning skills, communicate properly with people and have access to databases providing significant knowledge on a large range of domains. In light of these AI capabilities, the AIFS project aims to answer questions like the following.

How can we translate these advances in AI to the adjustments needed for education of middle school students?

To do this translation, we will be looking at different areas and processes involved with education and work, such as for example:

» **Learning objectives**: The set of skills that training is intended to develop.

» **Training curriculum**: The content of training in light of AI capabilities.

» **Human-AI synergy**: The ways humans and AI will work together, based on their respective strengths.

» **Educators' knowledge and competences**: The different requirements for educators.

This document, as well as any data and map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

Centre for Educational Research and Innovation

AI and the Future of Skills

Acknowledgments and Contacts

To learn more about the project background, events and publications, please visit our webpage (https://www.oecd.org/education/ceri/future-of-skills.htm) on OECD's website or contact the team via email (futureofskills@oecd.org).

OECD
BETTER POLICIES FOR BETTER LIVES

Federal Ministry
for Labour and Social Affairs