Creative Industries
Policy & Evidence Centre
Led by nesta

nesta

# The Art in the Artificial

## AI and the creative industries

**John Davies, Joel Klinger, Juan Mateos-Garcia, Kostas Stathoulopoulos**[1]
June 2020

## Executive summary

**This report aims to quantify levels of activity at the interface of Artificial Intelligence (AI) activity and the creative industries. The UK has a strong position in the creative industries. It has been one of the fastest growing parts of the economy and one of the fastest growing sources of employment in recent years. The UK also has one of the highest shares of the workforce in Europe employed in the creative industries. The importance of the creative industries and AI for the UK's future has been recognised with both areas receiving their own Sector Deals from the government.**

Recent breakthroughs in AI could have a large impact on the creative industries. Digitisation has meant that a high proportion of creative content: images, sound and text is now created, distributed and consumed digitally. Large volumes of creative content is recommended to us on platforms like Spotify and YouTube using AI techniques such as machine learning. Digitisation also means that creative content can be directly analysed using machine learning.

The tools of AI have become more effective and also more accessible with extensive open source software available on sites like GitHub. Recent developments such as Generative Adversarial networks (GANs) and style transfer have enabled new ways to be creative and are influencing artistic activity. The convergence of creative domains driven by wider digital change, such as the use of games technologies in visual effects and architecture, also mean that AI breakthroughs in one area could spread to multiple creative domains.

The UK therefore has an opportunity to exploit the synergies from its strength in both AI and the creative industries. At the same time if AI has a large impact on creative activity, and the UK does not keep pace with developments, then its position in the creative industries could be undermined.

Given the strategic importance of this topic, this report therefore studies levels of activity combining both creative activity and AI in research and commercial applications. To do this it uses data on research from the global scientific research paper repository arXiv, data on UK companies and academic collaborations (Gateway to Research) and data on international startup activity (Crunchbase). This enables international comparisons to be made and trends over time to be examined. The report's findings are as follows:

- **The UK is one of the world's leading research centres in AI, but other countries are rapidly increasing their research capability**. The UK has one of the world's highest levels of AI research outputs on arXiv, the highest in Europe. Between 2015 and 2019 UK publications in AI grew at 365 per cent (A rate of growth that is similar to that of Italy, the US and Czechia). By contrast, among countries with the fastest growing levels of research: Taiwan grew at 1,490 per cent, Sweden at 958 per cent and Japan at 845 per cent over the same period. This partly reflects that these countries had lower levels of research activity to start with, but also shows that the UK's position in AI cannot be taken for granted.
- **The UK has a strong position in AI research in domains potentially relevant for the creative Industries**. The UK has one of the highest levels of publications globally in AI involving media domains that are directly relevant for the creative industries such as image, text, sound. Only the United States, and China have higher levels of publications in these areas. AI research, as opposed to research involving other methods, is increasingly important in these domains.
- **Levels of AI research involving direct creative industry applications are relatively low internationally**. The evidence is that research directly involving AI and creative industry activity is relatively low. The number of papers is in the hundreds internationally, although the UK is one of the more active countries. Low levels of research activity may be due to research of this nature having commercial applications and so being undertaken in secret. It could also reflect low levels of public research work in this activity.
- **There are high levels of research in Generative Adversarial Networks (GANs), a key development in generative AI, and the UK is one of the most active countries**. Although a recent development (2014) there are already high levels of research with several 100s of papers internationally by 2019. UK institutions have published more research on this topic than any other European country.
- **Direct applications of AI in creative industries so far appear to be relatively low**. The levels of activity directly applying AI in the creative industries appears to be relatively limited. It was possible to identify UK companies that involved with AI and creative industry domains in Crunchbase, and university company collaborations involving AI and creative aspects in Gateway to Research, but the absolute level is in the tens of companies and projects. However in Crunchbase, only the US has more companies involved in this activity. This suggests that, at least in so far as it can be identified from text descriptions of companies and research projects, activity in this area is currently at a small scale, although there is evidence that the UK is relatively advanced in this regard relative to other countries.
- **The UK should help encourage activity at the intersection of AI and creative activity**. The levels of research collaboration combining both AI and creative industries is currently low on the evidence of arXiv and the Gateway to Research data. There are also relatively low levels of companies in Crunchbase that combine AI and creative activity. However at the same time, there are high levels of AI research in media such as text, image and sound that are highly relevant for the sector and in new approaches such as GANs and style transfer which are starting to see creative applications. If the UK wants to get the most out of its complementary strengths in AI and the creative industries it should therefore aim to support a higher level of commercial activity and academic research collaboration combining both in future.

Creative Industries
Policy & Evidence Centre
Led by nesta

# The Art in the Artificial

## AI and the creative industries

### Contents

1

# Introduction

Artificial Intelligence (AI) is the goal of developing machines with intelligent capabilities, such as pattern recognition, conversation and, ultimately, general intelligence itself. The field of AI is over 60 years old and has seen a series of breakthroughs in recent years with improvements in machines' ability in computer vision, in understanding language and in playing games like chess and Go.[2, 3]

This progress has been driven by improvements in AI techniques, in particular an approach known as deep learning, combined with the large volumes of data that digitisation has created in many areas. Deep learning is a technique that involves modifying a long-standing AI tool called a neural network (a network which processes data through its layers). It has been found that significant improvements in network performance can be achieved by increasing the number of layers and changing the network structure: for example refinements of filters in network layers have improved performance in image recognition tasks.

Following these developments, like many countries worldwide, the UK government has identified AI as a key technology for the future.[4] AI is considered likely to transform many parts of the economy including the creative industries, an area of competitive strength for the UK. The UK has one of the highest shares of employment in the creative industries in Europe and it has been one of the economy's fastest growing sectors in recent years.[5] Both AI and the creative industries have Sector Deals in the UK's industrial strategy recognising their importance for the UK's economic future.[6]

Given the creative industries' high levels of digitisation and potential for recent breakthroughs in AI to affect creative supply chains it is likely that AI will be important to the sector beyond its existing role in content recommendation algorithms and targeted advertising.[7] The growing convergence in media through the blurring of boundaries between computer games, films and virtual/augmented reality, and through online distribution, mean that the new developments in AI could have wide-reaching effects in multiple creative domains.

The UK is therefore well placed to capitalise on the interface between two areas of existing strength. Conversely, it faces a risk that if it does not do this, and AI is very disruptive to the sector, then its strong position in the creative industries will be undermined.

However, it is challenging to know what is happening at the intersection of AI and creative activity. Work on AI internationally has accelerated in recent years with far greater levels of research and many new developments. It is therefore hard to get an overview of what is happening. An additional challenge is that the creative industries are a complicated and interrelated set of areas which makes tracking the impact of technological change in them difficult. Most data sources about AI and/or business activity are also not labelled in a way that makes it possible to understand levels of activity in AI that are relevant for the creative industries.

To address these challenges this report uses a range of datasets. In particular, we analyse global activity in AI research and business activity related to the creative industries. We are interested in understanding levels of activity and their evolution, composition (in terms of different creative sub-sectors), geography and specific technologies and applications being developed. In particular the report analyses:

### 1. Research: What are the levels of research on:

- **AI that is related to creative industry domains**. The report identifies the levels of AI research potentially relevant for creative industry domains and more direct applications in creative activities.
- **New AI techniques that are opening up creative possibilities**. The report examines research on the recent developments of Generative Adversarial Networks (GANs) and style transfer that are widely considered to have enabled new creative possibilities. These are already finding applications in art sold at auction, music videos and design.

This is analysed using data from arXiv the world's largest open access repository of pre-print scientific research papers and a platform frequently used by AI researchers to share their work. The analysis is done internationally and also for the UK to enable comparisons.

### 2. Applications: What are the UK's levels of application of AI in the creative industries in terms of:

- **Collaborations between academic and industry partners that involve both AI and creative industry elements**. This is analysed using Gateway to Research, the portal for information on publicly funded research in the UK.
- **Companies combining both AI and creative activity in their operations**. This is analysed using data from Crunchbase, a global database of technology companies.

### The report structure is as follows:

- Section 2 describes the data that is analysed.
- Section 3 examines the levels of AI research relating to the creative industries and compares UK activity in this area with other countries.
- Section 4 assesses the levels of UK business activity that can be identified as involving both AI and creative dimensions.
- Section 5 concludes.

2

# Data sources

We use three datasets in our analysis: arXiv, Gateway to Research and Crunchbase.

## arXiv

arXiv is a repository of preprints widely used in the mathematical sciences such as mathematics, physics, engineering, and also by AI researchers with arXiv computer science created in 1993.[8] After collecting 1.5 million papers from arXiv's Application Programming Interface (API), we enriched them with information about the institutional affiliations of their authors and their location, and identified AI-related papers in them based on the text of their abstracts (All documents published on arXiv are required to provide an abstract in English). This resulted in 82,217 AI papers being identified.

Everything in arXiv was merged with Microsoft Academic Graph to identify publications that had been formally published in academic journals and also in conference proceedings as these are particularly important in computer science. Only arXiv publications which had also been published in this manner were included to ensure that the publications used in the analysis had undergone some level of scrutiny.

We analyse the computer science section of arXiv – as we do not expect to find significant levels of activity related to the creative industries and AI in other sections. There is a lot of research in the Humanities related to the creative industries which, as it is not included in arXiv, does not feature in the analysis. For example this means we do not, as philosophy is not on arXiv, analyse philosophical research on AI and creativity, or relating AI to ethics which might be relevant.

## Gateway to Research

Gateway to Research (GtR) is a website developed by the UK Research and Innovation funding agency (UKRI) to allow access to information about publicly funded research in the UK. GtR provides information about publications, people, organisations, collaborations, outcomes and classifications relating to research projects. It includes information on project names, abstracts, the amount awarded and how long the project ran for.

## Crunchbase

Crunchbase is a global technology directory commonly used in the analysis of entrepreneurship and technology ecosystems. It contains the description of companies, metadata about them such as their incorporation date, the amount of funding that they have received and their location.

## Limitations of the data

These datasets are highly relevant and comprehensive, however there are limitations to what they cover. The abstract or text descriptions that we are analysing may not always be a complete picture of projects or companies involved. ArXiv is a comprehensive dataset on the AI research that has been undertaken which is public. In focussing on arXiv computer science we may however be omitting research that is relevant, but in another domain such as arXiv physics. Gateway to Research will capture publicly funded collaborations between companies and universities, but it will not include details of any private work on AI technologies between the two. In the case of Crunchbase, as the focus is on startups, it will not include established companies and their activities in AI. In practice the details of a lot of work on AI technologies will due to commercial confidentiality never be completely public. Sources of information that could, to some extent, be used examine this kind of activity in future include data from the trade press or the code sharing website GitHub. In the case of arXiv and Crunchbase there are also likely to be sample selection effects due to variations in their use or coverage across different countries.

## The treatment of ICT and Computer programming in the analysis

The official UK creative industries definition includes Computer programming activities (The Standard Industrial Code 62.01). Purely theoretical AI research aside, effectively all practical implementation of AI would be a creative industry activity on this basis.[9] However, studying all of AI would, by definition, be a report on AI and distinguishing the direct use of AI to deal with computer programming issues from other applications of AI is not straightforward. For the purposes of this analysis we therefore do not explicitly analyse AI in general or its applications in ICT.

3

# Research on AI relevant to the creative industries

UK research institutions are among the world's largest producers of publications on AI, publishing more than any other European country. Although the UK's AI research levels have been growing rapidly this is also true of many countries around the world.

Figure 3.1 shows the number of AI publications on arXiv for selected countries, where the country is measured in terms of the country that authors' institutions are located in.[10] This approach, which is adopted for all international comparisons means that a paper that had two co-authors, one from a US institution and one from a UK institution, would therefore be counted twice, once for the UK and once for the US. If multiple co-authors come from the same institution then each author counts once to the country count. UK publication levels are the highest in Europe below the US and China.

**Figure 3.1: Number of AI publications by country over time**

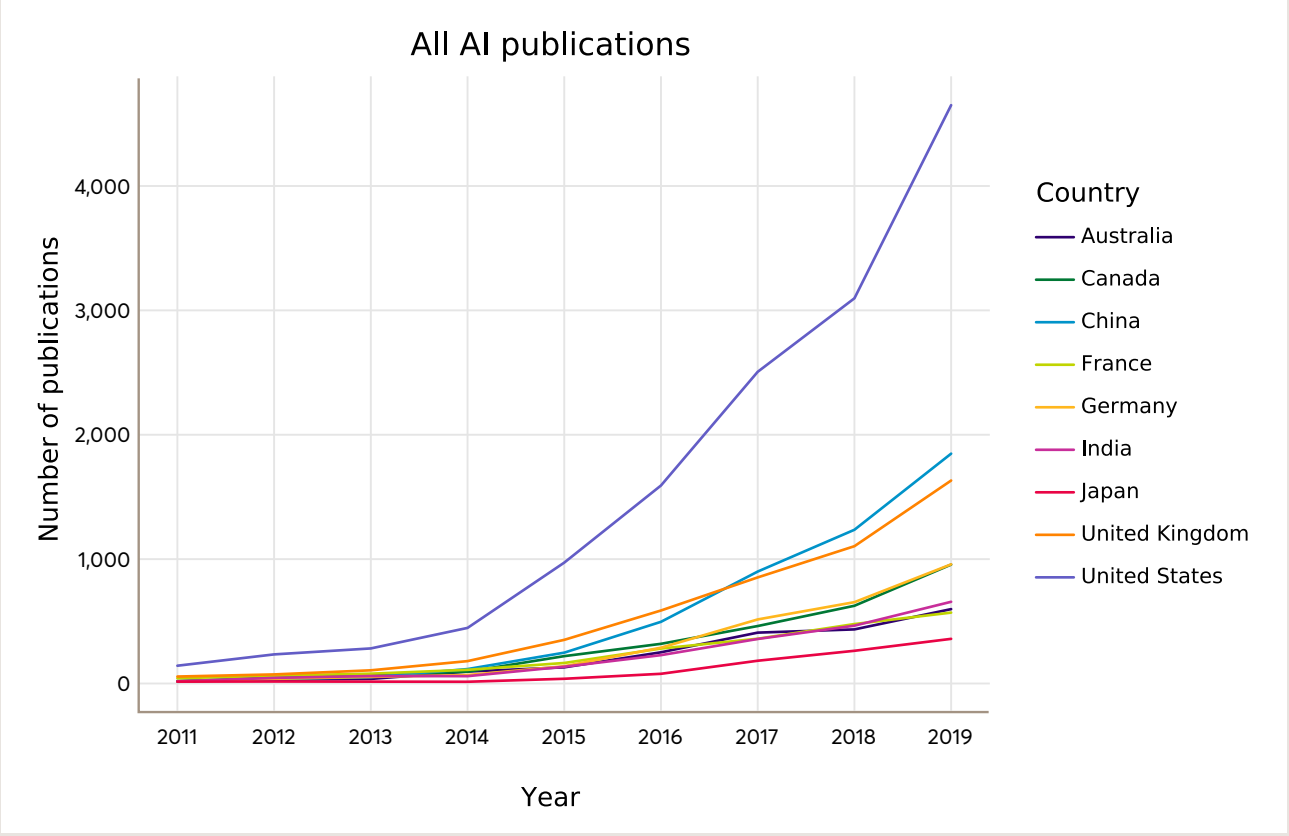

The chart measures research within academic institutions. However for authors based in global multinationals, such as Google, it has not been possible to give a clear country label. Publications involving these organisations have therefore been removed from the country analysis. As the chart does not include the research done by the large tech companies headquartered in the US and China the totals for these countries will be understated as when these organisations publish their work it is frequently to arXiv.[11]

Research globally is growing rapidly as shown in Figure 3.2 which displays the growth rate in publications between 2019 and 2015 for countries that had significant levels of publications which we take here to be over 300 publications in total on arXiv. The growth rate in AI publications in the UK is lower than for most of the countries shown. Between 2015 and 2019, UK publications in AI grew at 365 per cent (A rate of growth that is similar to that of Italy, the US and Czechia), while by contrast among the fastest growing countries Taiwan grew at 1,490 per cent, Sweden at 958 per cent and Japan at 845 per cent. This partly reflects the fact that the UK already has more publications in this topic, but if these trends continue other countries will catch up.

## Figure 3.2: Growth rates in AI publications 2019-2015

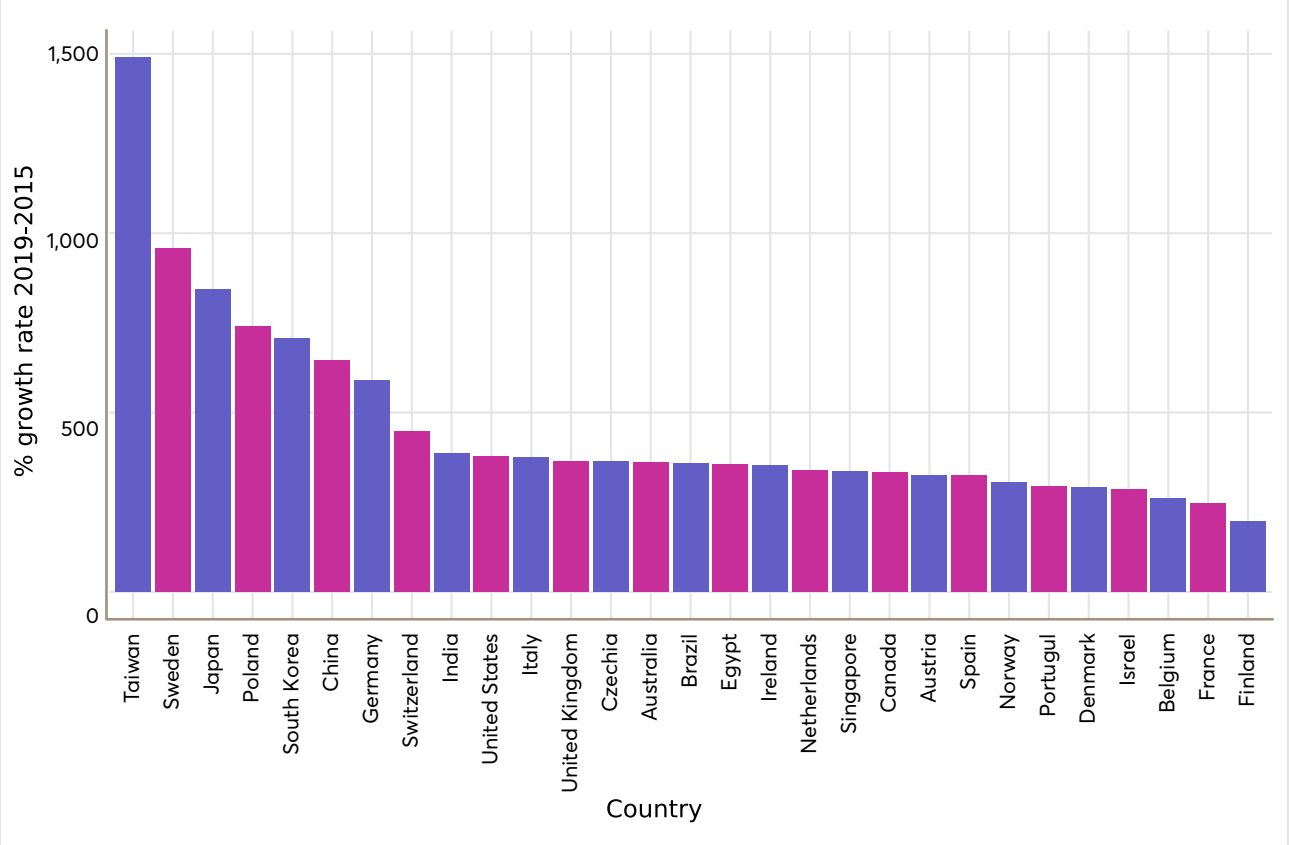

The diagram below illustrates the relationships between different kinds of AI activity and applications of AI in the creative industries.

### Figure 3.3: AI and Creative industry activity and applications

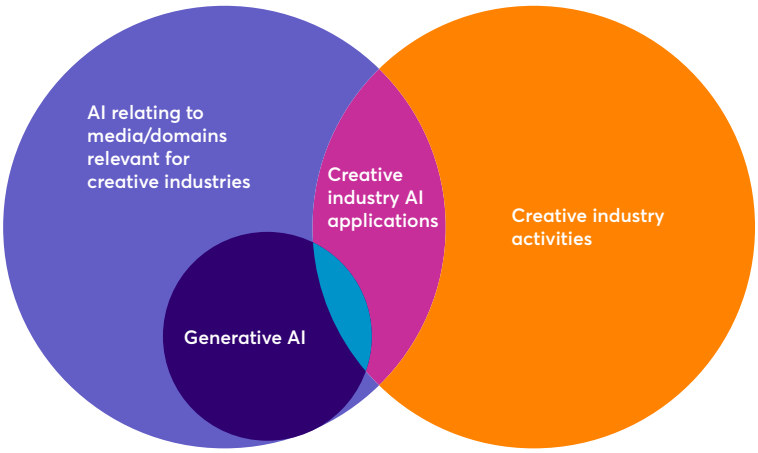


In focussing on AI research relating to the creative industries we examine:

1. AI involving media/domains relevant to the creative industries e.g. relating to images, audio and text. The raw material of the creative industries: language, imagery and sound can all be digitised and converted into data. As machine learning algorithms are often applied to these, AI research of this nature could have future implications for the creative industries, although it may not necessarily be directly analysing creative content.
2. AI research that is being directly applied to a creative industry-related activity, for example using AI to animate a character in a computer game or to generate a particular kind of visual effect in an image.
3. AI that can be considered in some sense as being creative/generative and so could have a direct impact on the creative industries, but which may have wider applications. The two examples of this that we examine are Generative Adversarial Networks (GANs) and style transfer.

In this section we attempt to estimate the UK's position in research that falls in these three categories in turn. In the next section we attempt to estimate levels of activity in applications.

## 3.1 Research on AI related to the creative industries

### Identification of AI

AI is a huge area of research activity ranging from mathematical research into AI algorithms, to building hardware, to attempts to produce software that can reason or have a conversation and direct applications in areas like self driving cars or automatic assistants. To identify papers relating to AI we used the approach in Stathoulopoulos and Mateos-Garcia (2019).[12] An initial list of keywords, namely Artificial Intelligence, Machine Learning, Deep Learning and Data Science was used with a trained word embedding (word2vec) model to find semantically similar terms. The 250 most similar terms of each keyword were retrieved and the process repeated, with the 50 most similar terms of each token on the expanded query list collected. Lastly, tokens with an Inverse Document Frequency (IDF) weight lower than the 5th percentile or higher than the 95th percentile of the IDF frequency distribution were removed.[13] From this related keywords were identified. These were then searched for in the processed publication abstracts and those that contained at least one of the keywords were labelled as AI.

### 3.1.1 AI research in domains relevant for the creative industries

To identify AI research papers in domains relevant for the creative industries the arXiv paper abstracts were matched to a set of keywords representing these domains and papers selected where there was at least one match with the abstract. The keywords that were used to do this are shown in Table 3.1 below. The keywords were obtained by selecting a set of core words representing these domains and then supplementing with words identified as similar to them using a word2vec model trained on the arXiv data.[14] The result was manually inspected and words removed that we considered should not be included as they did not clearly relate to the domains.

**Table 3.1: Keywords used to identify research papers from broad domains**

| Activity | Keywords |
|---|---|
| Image based | Film, image, photo, photography, television, video, videos |
| Sound based | Audio, writing, music, musical, song, songs, sound |
| Text based | Printing, book, e-book, book, journalism, newspaper, text |

Figure 3.4 shows the number of papers that were retrieved corresponding to these three domains.

**Figure 3.4: Number of publications in different domains**

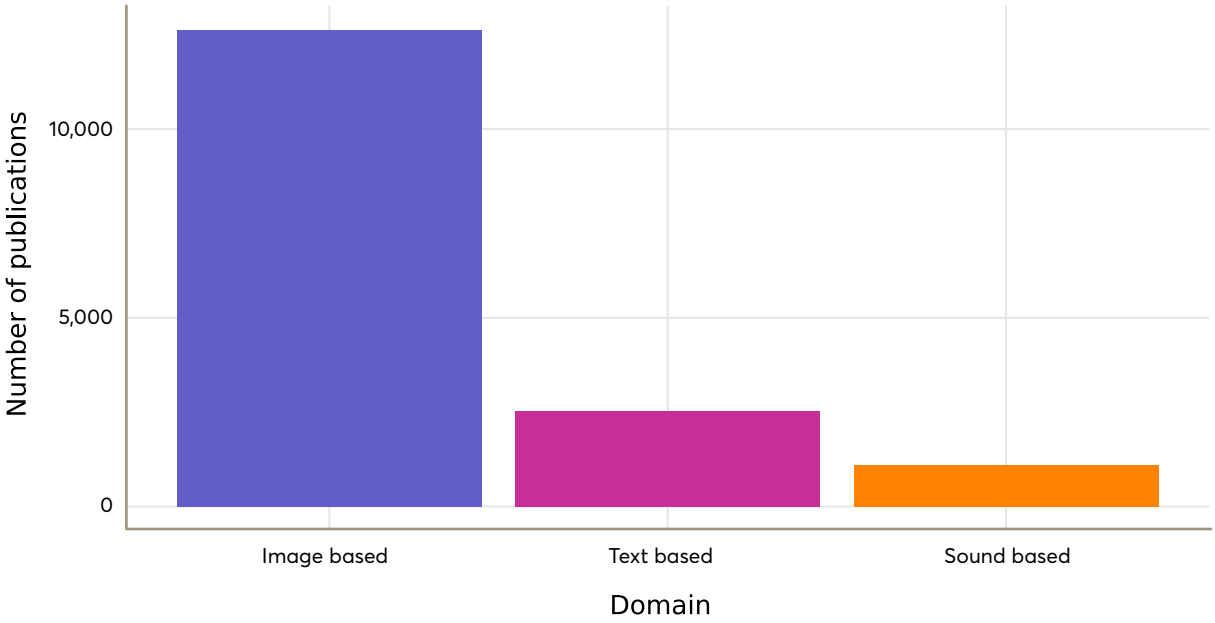


To give some examples of the activities the papers involved:

- Using neural networks to identify breast cancer from imagery.[15]
- Identifying arabic text within photographs.[16]
- Predicting individuals emotions based on images and text shared on social media.[17]
- Filtering out voices from background noise with neural networks.[18]
- Producing a realistic animation of a human face so that it looks as if it is speaking.[19]
- Obtaining semantic understanding from audio signals.[20]
- Understanding the position someone is taking to a topic on Twitter.[21]
- Sentiment analysis of text using a new neural network architecture.[22]

As Figure 3.4 shows, there are very high levels of research relating to image based data, with over 10,000 papers identified. By comparison, at least in so far as can be identified from papers' abstracts using our keywords there was significantly less work on text and sound data.[23] The many publications relating to images reflect the very wide range of issues and applications in this area. For example the multiple applications of image processing in photography, filming; the use of image recognition to identify faces, handwriting, cars, number-plates, medical conditions from images etc.[24] As shown in Figure 3.5 something under half of the papers relating to images were found to involve AI, with a lower proportion for sound based and text based research.

## Figure 3.5: The proportion of publications that involve AI

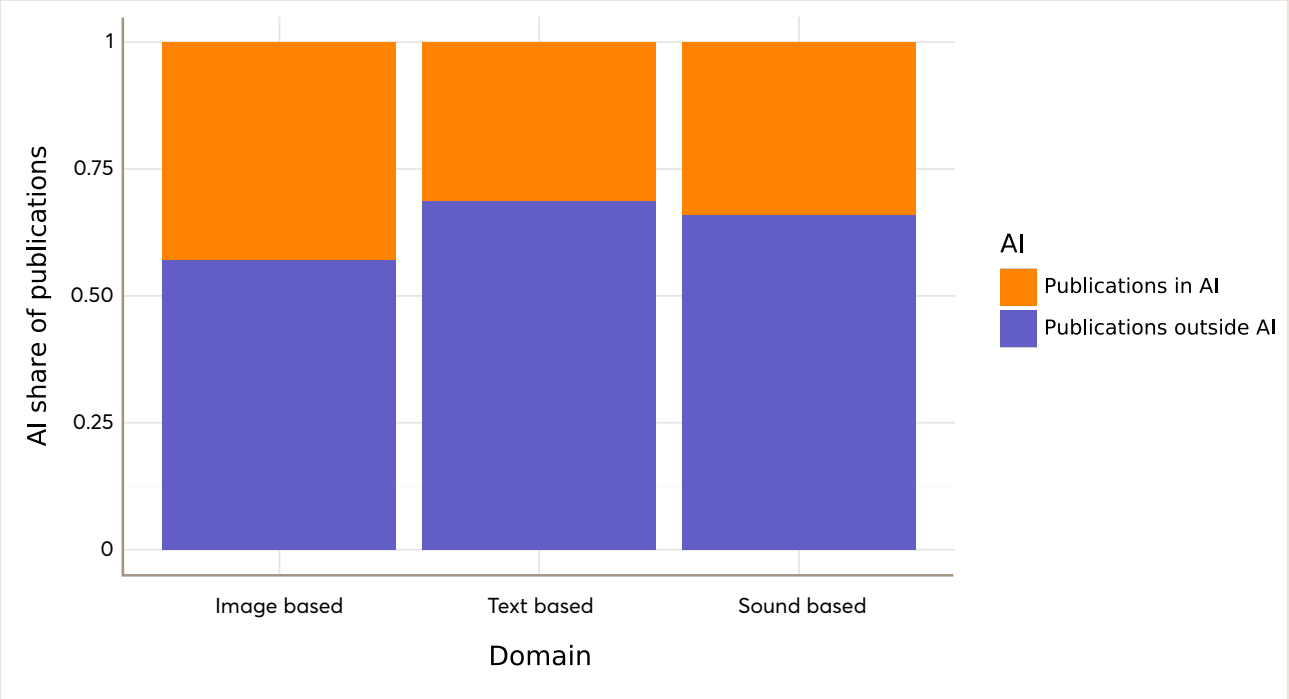


If we combine all these papers together and see how the share of AI of publications in this area has been growing over time. Figure 3.6 shows the relative growth in the number of publications that involved AI. The rapid growth in AI research in recent years is clearly visible with it constituting a much higher proportion of publications on image, text and sound analysis in recent years.

## Figure 3.6: Levels of AI and non-AI research on image, text and sound data

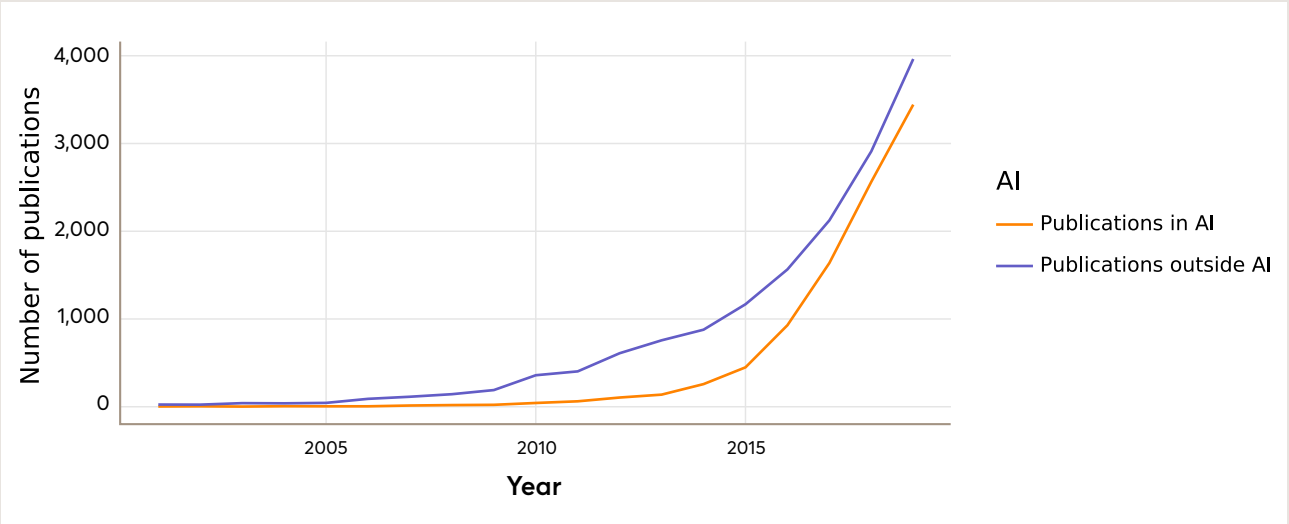

Splitting this into the three categories, we can see the increase in the uptake of AI methods across three domains. Figure 3.7 shows AI publications increasing first and fastest in image based research, with the growing interest in AI spreading to text and sound based analysis later.

### Figure 3.7: Growth in AI methods for image, sound and text based data

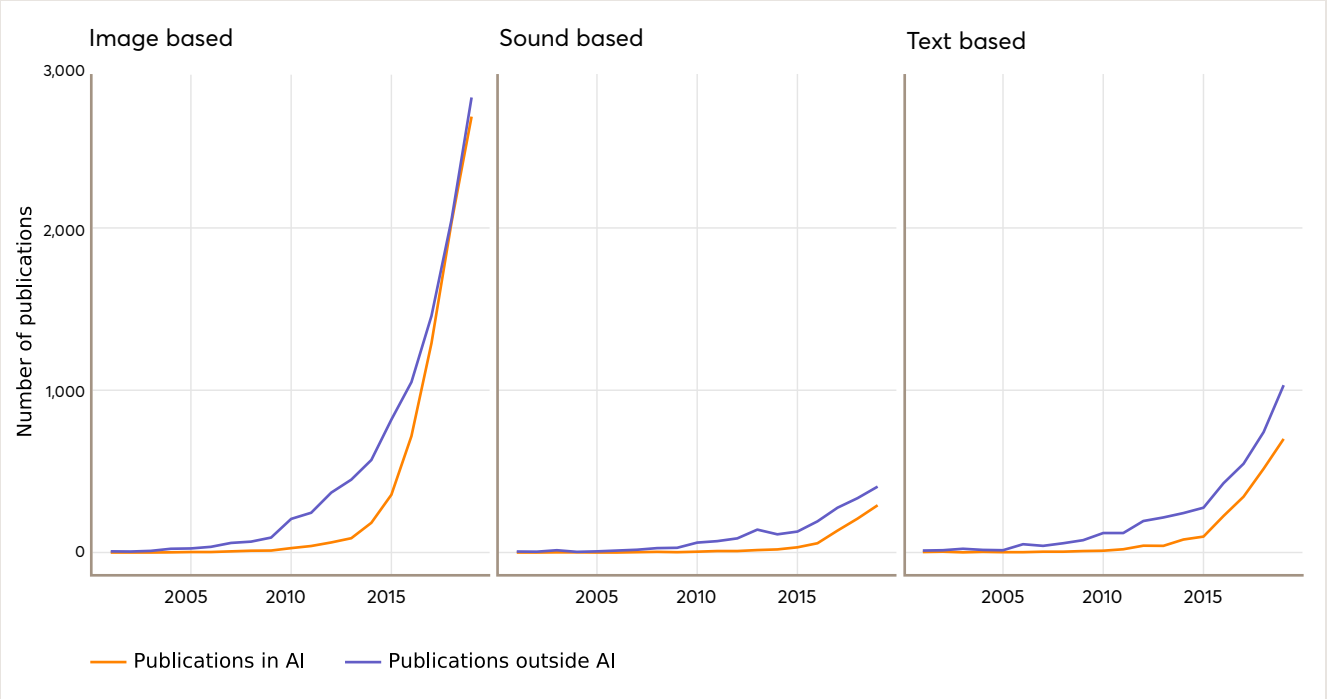


### Table 3.2: AI publications in image, sound and text based media by country

| | Country | Image based | Sound based | Text based | All three | % of all papers |
|---|---|---|---|---|---|---|
| 1 | United States | 3,319 | 199 | 749 | 4,041 | 24.9 |
| 2 | China | 1,590 | 59 | 257 | 1,805 | 31.6 |
| 3 | United Kingdom | 1,138 | 116 | 270 | 1,437 | 25.1 |
| 4 | Germany | 666 | 51 | 139 | 823 | 25.8 |
| 5 | Canada | 658 | 47 | 141 | 818 | 25.2 |
| 6 | Australia | 573 | 33 | 115 | 684 | 29.5 |
| 7 | India | 474 | 26 | 133 | 594 | 25.3 |
| 8 | Switzerland | 461 | 33 | 104 | 570 | 28.6 |
| 9 | France | 449 | 33 | 85 | 549 | 22.2 |
| 10 | Ireland | 412 | 32 | 119 | 538 | 21.6 |

### Figure 3.8: Geographic distribution of research

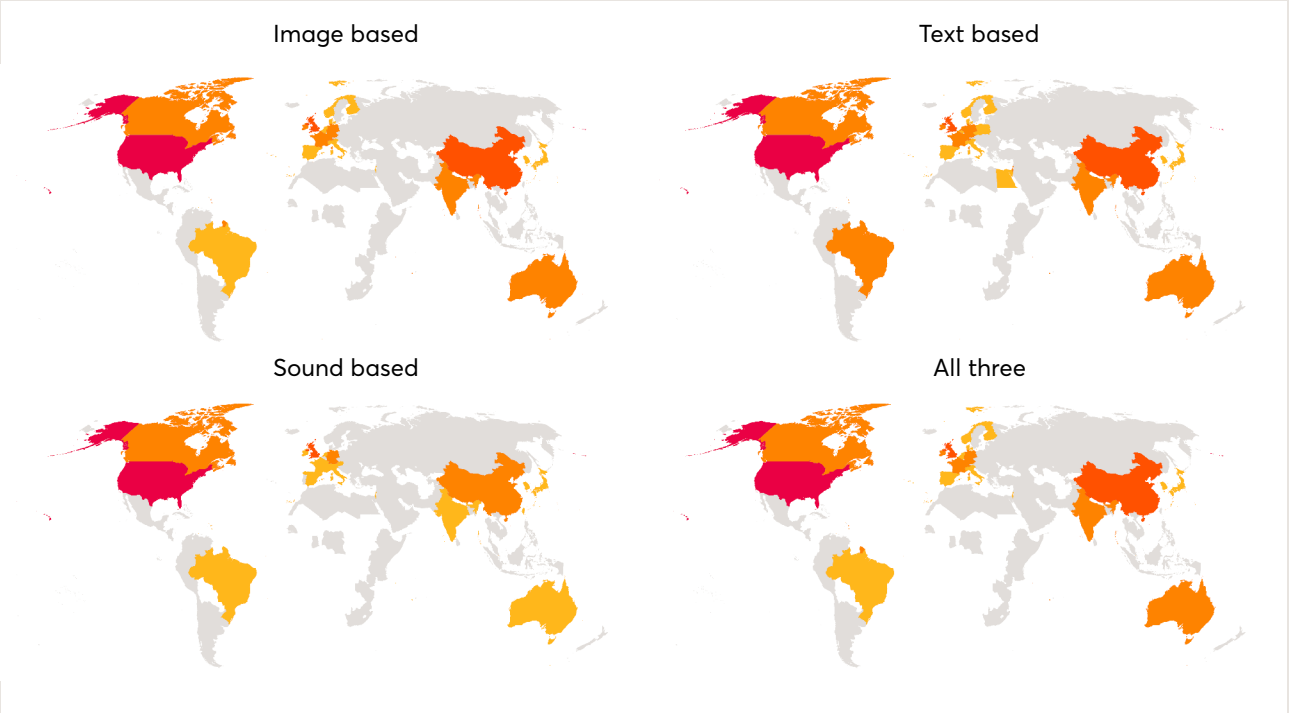


If we look at the international publication levels in AI (Table 3.2) we see that the UK has the most publications in Europe in all three areas and globally is behind only the US and China. This research is not necessarily directly related to creative activity, but as it involves media central to the sector it suggests UK strength in relevant AI research. We now turn to examine AI research that is more directly related to creative activity.

### 3.1.2: Research involving AI technologies were there are direct applications to the creative industries

The analysis of AI research in the previous section focuses on media that, while relevant for creative industries, is quite general and does not necessarily correspond to creative industry activity. To address this we now focus on research that is likely to be more directly linked to creative activity. To do this we examine areas that are quite specific and so are less likely to raise false positives. The areas that we analyse are:

- Advertising.
- Computer games.
- Virtual reality/Augmented reality.
- Artistic activity.
- Museums.

We also examine technical activity that is highly, although not exclusively, used in creative industries such as photogrammetry (reconstructing 3D shapes from multiple photographs) and motion capture (recording human movement) which are widely used in visual effects and computer games.[25]

The basis for the selection of these categories is that they were considered to be likely to be discrete and relatively well defined. Two important areas which are perhaps harder to analyse are design and fashion. As these terms can be used very generically e.g. 'experimental design', or 'fashion' as in trend, including them was considered likely to lead to significant likelihood of false positives. Another issue is that they could be included due to featuring in research quite indirectly, for example there is extensive research identifying people in photographs, and clothing plays a part in this, but this is not necessarily related to fashion. Exploring these two key creative domains in more detail is an area for further work.

As with the previous analysis we start with a set of core terms which we then expand through the use of a word embeddings model. The resulting key words that we matched to the paper abstracts are given in the table below.

### Table 3.3: Keywords used to identify research papers relating to specific creative activity

| | |
|---|---|
| Advertising | Ad impressions, advertisement, advertisements, advertisers, advertising, adverts, online advertising |
| Computer Games | Computer games, mmorpg, Nintendo, Playstation, video games |
| Immersive Reality (VR and AR) | Virtual reality, augmented reality, immersive reality, mixed reality, Magic Leap, Oculus Rift |
| Mocap, photogrammetry and other technology related terms which are used extensively in visual effects and computer games. | 3D modelling, camera tracking, compositing, matchmoving, matting, motion capture, photogrammetry |
| Artistic | Artistic, artworks, painting, paintings |
| Museums | Exhibition, museum |

Note some variations on these terms are excluded for readability.

Figure 3.9 shows the total number of publications whose abstracts matched the keywords in these domains using these specific areas of creative activity (a sensitivity analysis of this is included in the Appendix). The number is much smaller than the wider range of general work on image, text and sound data. VR/AR has the highest levels of activity, followed by advertising. The high levels of research in VR/AR relative to computer games probably reflects that the area is newer commercially and technically and so there is more academic research on the topic. The number for advertising would have been much higher if we had included other keywords such as recommender or social network/social media, which are potentially relevant. Indeed, when in the appendix we as a robustness test broaden the number of words to identify each of the categories, advertising increases the most suggesting that the figure may be understating levels of research related to advertising.

## Figure 3.9: The number of publications by creative industry category in arXiv

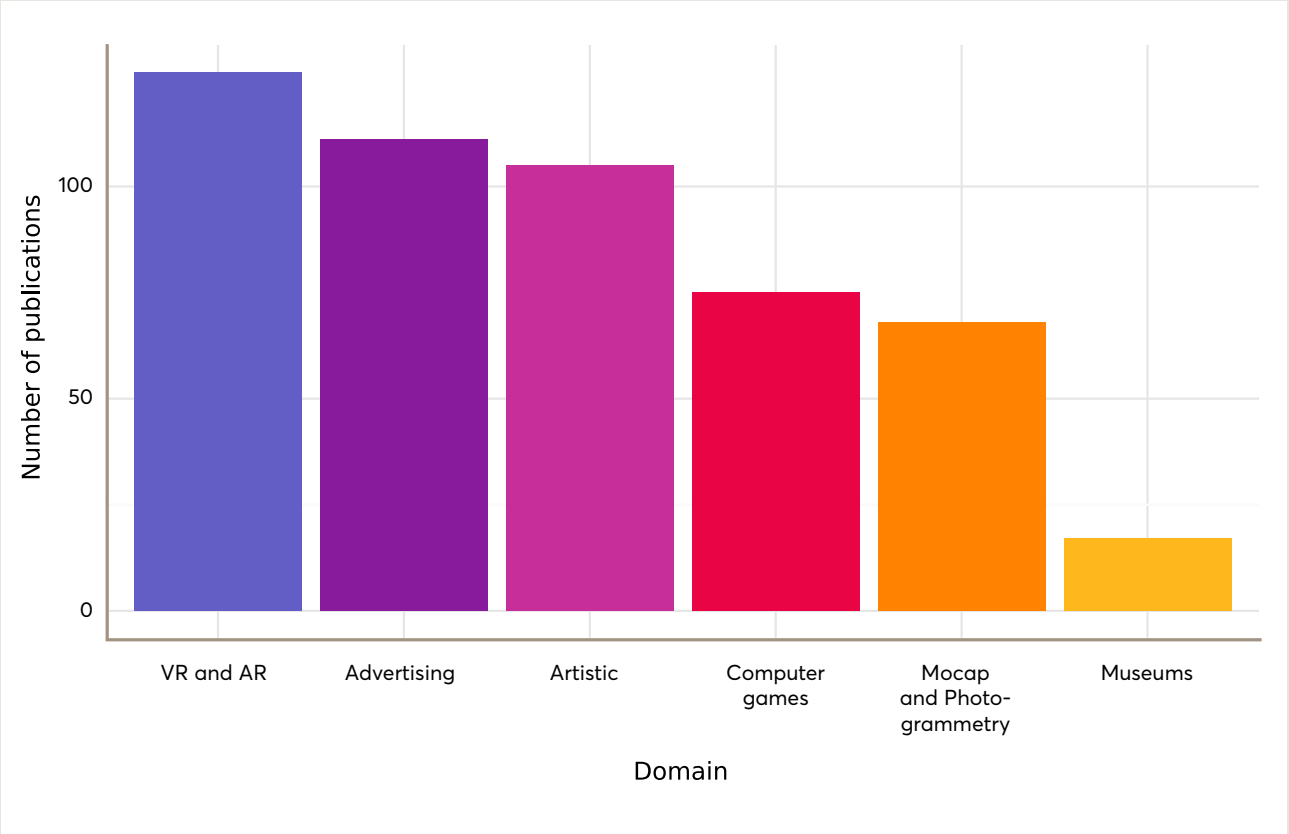


The low levels of research overall is perhaps to be expected given that as many of these areas are commercial a lot of research on them will happen inside companies and be trade secrets without any public record.

Conversely the objective of academic research is not necessarily to have specific practical applications, but often to improve the performance of general techniques. For example one of the applications of recent AI breakthroughs is in data compression; an issue which occurs in many areas where data is available, including the creative industries. It is therefore not surprising that there are fewer direct research applications.

Figure 3.10 below shows the AI share of publications relating to these creative domains. In general the share that AI constitutes of publications is lower than in the broad categories of image and text data.

## Figure 3.10: Share of AI and non-AI publications in applied creative industry research on arXiv

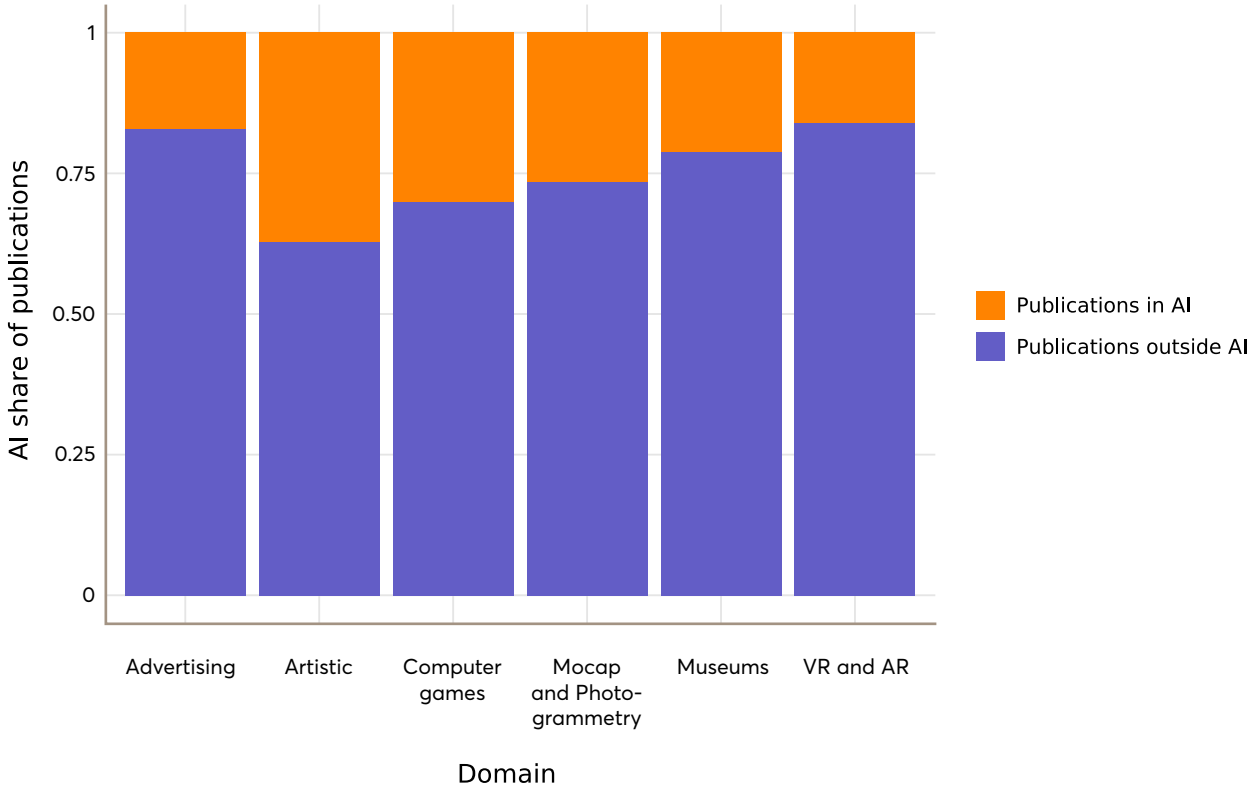


To give some examples from the papers involving AI and how it is being used in these creative domains. AI is being:

- Used to optimise revenues from mobile gaming.[26]
- Tested and developed in computer game environments.[27]
- Used to piece together fragments of objects in museum collections.[28]
- Used to create more realistic avatars in virtual reality.[29]
- Used to develop measures of paintings' visual similarity.[30]
- Used to help develop better approaches to motion capture.[31]

If we then split the AI publications in the domains down by country, we see that the UK has among the highest number of publications in these areas in Europe, although the number of publications involved overall is small. Relative to Germany the UK has more publications on advertising, and computer games, and fewer in the area of mocap/photogrammetry.

### Table 3.4: Number of research publications by country on creative related AI activity

| Country | Advertising | Artistic | Computer games | Mocap and photo-grammetry | Museums | VR and AR | All types | % of all papers |
|---|---|---|---|---|---|---|---|---|
| United States | 32 | 21 | 20 | 22 | 2 | 20 | 117 | 0.7 |
| China | 14 | 11 | 3 | 5 | 0 | 11 | 44 | 0.8 |
| United Kingdom | 4 | 6 | 8 | 8 | 0 | 6 | 32 | 0.6 |
| Canada | 2 | 5 | 6 | 6 | 0 | 3 | 21 | 0.6 |
| Germany | 1 | 2 | 1 | 9 | 0 | 7 | 20 | 0.6 |
| Italy | 4 | 2 | 5 | 5 | 1 | 1 | 18 | 1.1 |
| Australia | 8 | 4 | 1 | 4 | 0 | 0 | 17 | 0.7 |
| Switzer-land | 3 | 2 | 1 | 6 | 0 | 3 | 14 | 0.7 |
| France | 0 | 1 | 3 | 3 | 0 | 4 | 11 | 0.4 |
| Ireland | 1 | 2 | 2 | 2 | 0 | 2 | 9 | 0.4 |

### 3.1.3 Research activity in new AI technologies that are considered to be creative

Recent technological developments mean that AI is likely to become more directly involved in the development of creative content and have a more visible influence.[32] Two recent techniques widely considered to be important for the creative industries as they have provided new creative tools are Generative Adversarial Networks (GANs) and style transfer.

GANs are a technique which can be used to automatically generate realistic imagery.[33] They involve two neural networks competing against one another, one of the networks is trying to create fake versions of a specific image, such as a human face, while the other network is trying to detect the fakes from the real images. Artworks generated by GANs have already sold for over $400,000 at auction resulting in a debate on the role of AI in the future of art.[34] GANs have also received attention through their role in producing films and images where the features of a person are superimposed on another in so called 'DeepFakes'. It has been suggested that these techniques could affect visual effects pipelines in future, automating labour intensive tasks such as creating the faces of background characters in films.[35] The company icon8 has made 100,000 face images created by GANs available for use royalty free and it is conceivable that developments like this could have an effect on the market for stock photography in future.[36]

Style transfer is a technique which allows a style from one area to be transferred to another, for example taking the style from an impressionist painting and transferring it to another object.[37] The technique can also be applied to both image and sound data. Although only a few years old the technique has already been used in music videos and designing virtual reality experiences.[38] An example of an artistic application is Gene Kogan's Cubist Mirror where a viewer standing in front of a screen finds their recorded image alternatively transformed into the styles of different artists such as Klimt, Munch and Hokusai.

Figure 3.11 shows that there are far more publications on GANs than on style transfer internationally. This was identified in terms of articles that mentioned GANs and style transfer in their abstract. The US and China have the highest levels of publications overall on GANS, while the UK has the highest levels of publications on GANs in Europe.

It is likely that the higher levels of interest in GANs is because they are a more radical and complicated change in neural network architecture than style transfer. As they create imagery directly as opposed to transferring a style they are also arguably more creative.

### Figure 3.11: International publications on GANs and Style transfer

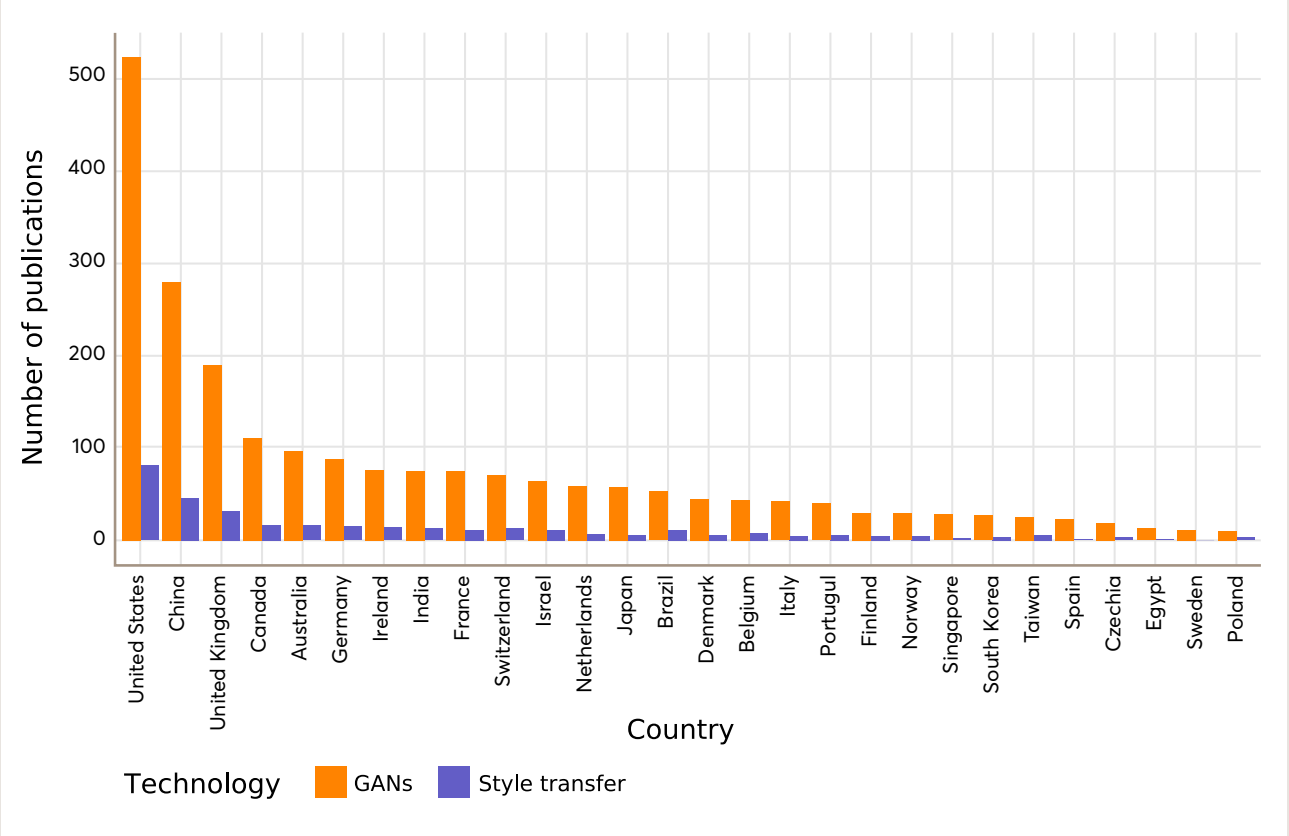


Examples of applications in the research papers include using GANs to:

- Fill in missing sequences in Magnetic Resonance Imaging (MRI).[39]
- Forecast 3D human motion given a set of observed 3D skeleton poses.[40]
- Remove blur caused by camera shake.[41]
- Apply styles to anime cartoons.[42]

Examples of style transfer applications in the research include using it as a:

- Design tool to generate creative digital content.[43]
- Way to swap faces in images.[44]
- Means to transfer styles in polyphonic music recordings.[45]
- Way to change styles in real-time.[46]

4

# Applications of AI in the creative industries

Having analysed research activity relating to creative domains this section estimates levels of application of AI technology in creative industries. It firstly looks at partnerships between companies and universities that involve both, and then looks at creative industry companies that indicate they are using AI.

## 4.1 Understanding academic and external research partnerships related to AI and the creative industries

To understand partnerships between academic research institutions and companies that involved both creative industries and AI we selected Gateway to Research (GtR) projects that had both academic and non-academic partners and identified projects whose text description indicated that they involved both AI and creative elements. This was based on matching the keywords set out in Mateos-Garcia (2018).[47] Projects that did not have a named company attached to them were excluded. This resulted in a relatively small number of 43 projects being identified out of 1389 projects involving AI and 995 involving creative activity. The great majority of projects (35, 81 per cent) were research grants. This is more than the share of research grants in GtR as a whole (51 per cent).

### Table 4.1: The split in AI and creative projects for GtR projects involving non-academic partners

| | | Project categorised as AI | |
|---|---|---|---|
| | | False | True |
| Project categorised as creative | False | 36,132 | 1,346 |
| | True | 952 | 43 |

If we break these projects down by sector we obtain Table 4.2 below.

## Table 4.2: The split in AI and creative projects for GtR projects involving non-academic partners

| AI identified | Museums, galleries and libraries | Film, TV, video, radio and photography | Design | Architecture | Publishing | Advertising and marketing | Crafts | Music, performing and visual arts | Creative industries |
|---|---|---|---|---|---|---|---|---|---|
| No | 301 | 211 | 382 | 116 | 73 | 56 | 14 | 302 | 1,346 |
| Yes | 10 | 11 | 10 | 1 | 11 | 4 | 0 | 12 | 43 |
| AI share | 3% | 5% | 3% | 1% | 15% | 7% | 0% | 4% | 3% |

The 43 projects identified are relatively evenly distributed across the sub-sectors with the exception of crafts, advertising and marketing, and architecture which have few or no projects. Where companies from more than one sector participate in a project then the sector will be appearing twice.

## Figure 4.1: Number of AI-related projects on Gateway to Research by creative category

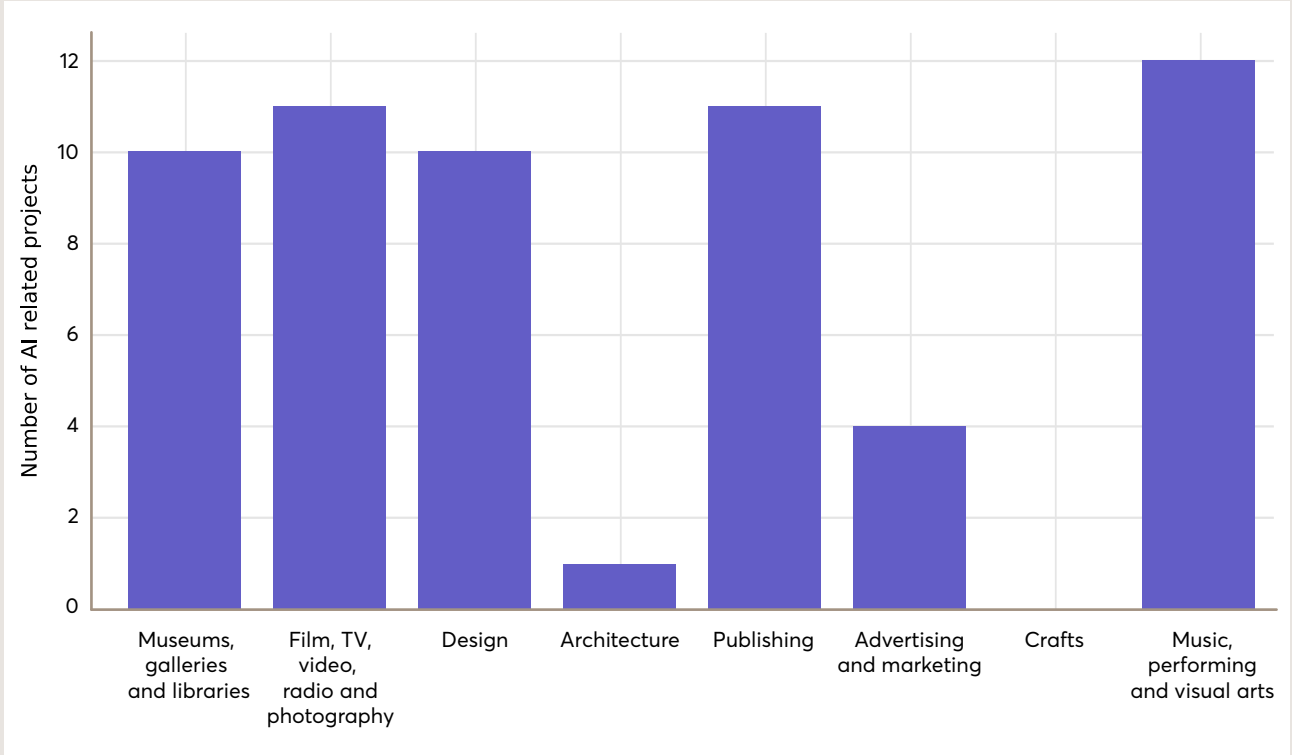

As some examples of what the identified projects were doing, projects included:

- Looking at ways to develop rankings in search.[48]
- Work on producing real time visualisations.[49]
- A research network to support live coding (i.e. where artwork and music is generated live by programming).[50]
- A research network to explore social justice in the digital economy.[51]
- Working on making public spaces safer (a project which makes use of AI).[52]

In general the levels of AI in the projects overall appeared to be lower than that in the arXiv papers as they typically contain multiple areas, rather than being a purely AI-focussed research paper.

In GtR overall there has been a steady growth in the number of projects with corporate partners that had either AI or creative aspects over time. However, the number of projects that combine both AI and creative elements has been relatively stable and at low levels.

**Figure 4.2: The number of AI and creative projects on Gateway to Research over time**

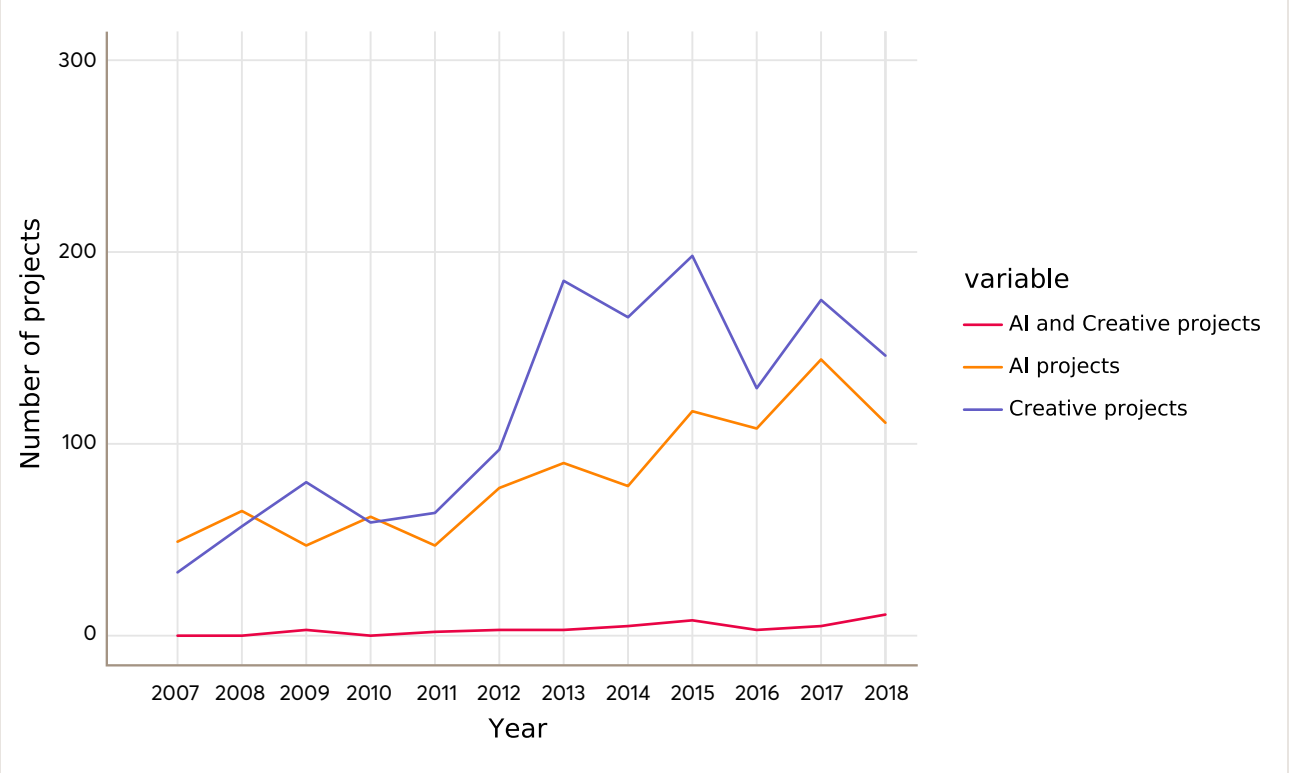

## 4.2 Companies engagement with AI and the creative industries on Crunchbase

To identify companies engaged in activity relating to both AI and the creative industries we analysed data from the Crunchbase database.

To classify companies into creative industries sub-sectors, we have analysed the co-occurrence of sub-sector labels in companies using a community detection algorithm allowing us to identify higher level groups. This takes us from an initial list of 742 categories in the existing Crunchbase industry taxonomy to a smaller set of 129, which we manually assign to creative industries categories using the tags in the table below.

### Table 4.3: Tags for Crunchbase sector clustering

| Creative sector | Crunchbase tags |
|---|---|
| Advertising | Advertising, marketing, brand marketing, personal branding, cause marketing, digital marketing, seo, social media management, content marketing, social media marketing, sem, google, reputation, semantic search, social media advertising, facebook, ad network, advertising platforms,app marketing, local advertising, mobile advertising, video advertising, ad targeting, ad exchange, ad retargeting, affiliate marketing, ad server, dsp, outdoor advertising, app discovery. |
| Architecture | Architecture, building maintenance, construction, diy, furniture, home décor, interior design, civil engineering, home improvement, home renovation, home services, infrastructure, landscaping, home and garden, housekeeping service, janitorial service. |
| Crafts | 3d printing, 3d technology, cad, printing. |
| Design | Creative agency, graphic design, human computer interaction, ux design, usability testing, web design, web development, product design, product management, product search. |
| Film, video, tv | Animation, digital entertainment, digital media, film, film production, media and entertainment, motion capture, tv, tv production, film distribution, video, video editing, video on demand, video streaming image recognition, photography, file sharing, photo editing, photo sharing, visual search, archiving service, celebrity. |
| Games, immersive | Esports, gamification, gaming, pc games, serious games, video games, casual games, console games, gambling, mmo, games, online games, Nintendo, playstation, xbox, augmented reality, simulation, virtual reality, virtual world. |
| Music, performing arts | Art, handmade, museums and historical sites, performing arts, collectibles, comics, audio, music, audiobooks, independent music, music streaming, podcast, internet radio, music education, music label, music venues, musical instruments, event management, events, theatre ticketing, nightclubs, nightlife, trade shows, concerts, event promotion, wedding. |
| Publishing | News, publishing, broadcasting, journalism, social news, blogging platforms, content, content creators, content discovery, content delivery, network, content syndication, contests, ebooks, freemium, subscription service,reading apps. |

Having done this, we have created a set of what we call 'pure' creative companies which only have tags in one creative sub-sector, and trained a supervised machine learning model to identify words in their descriptions that are highly predictive of that category. We then use the model to predict the probability that any company in the Crunchbase dataset – including companies containing mixed sector categories – are in a creative sub-sector. We assign companies with a very high probability (> 0.99) of being in a creative sub-sector to that category.

Using this approach we identified 61,282 creative industry companies internationally in the Crunchbase data (12 per cent of the total). Out of these, 635 companies (1 per cent of the total) are also classified as specialising in AI. The share of AI companies in the Crunchbase population overall is somewhat higher (2 per cent). We identify AI companies in the data based on the presence of relevant keywords in their descriptions.[53]

It is important to note that our definition of AI only captures companies that mention AI and related terms in their description. One prominent example of a company that uses AI, but which does not refer to it in its description, is Google. Although this means we could be underestimating the total number of companies developing or deploying AI, this issue would apply to AI-related work undertaken in other sub-sectors too. Also the share of AI-related companies located in the creative industries is only underestimated insofar as creative businesses are more likely to be doing work on AI which they chose not to mention in their company description.

As Figure 4.3 shows, there are some important differences in the proportion of AI-related companies in the different creative sub-sectors according to the Crunchbase data: AI is used more in games and immersive, advertising and marketing, and less used in other creative sub-sectors. The low share for architecture may reflect that we have a broad definition that is including wider construction and infrastructure activity, thus increasing the number of companies that have less of a digital focus.

Some examples of AI companies in different creative sub-sectors include:

- Advertising and Marketing: advert targeting and analysis of customer behaviour.
- Radio, Film and TV: Use of AI for video compression.
- Music: An AI-powered talent discovery platform.
- Games and Immersive: Development of AI agents for video games.
- Architecture: AI-aided design tools for the residential sector.

The range of examples illustrates, as before, the applicability of AI technologies across multiple creative sub-sectors and functions.

## Figure 4.3: Proportion of projects in creative domains involving AI

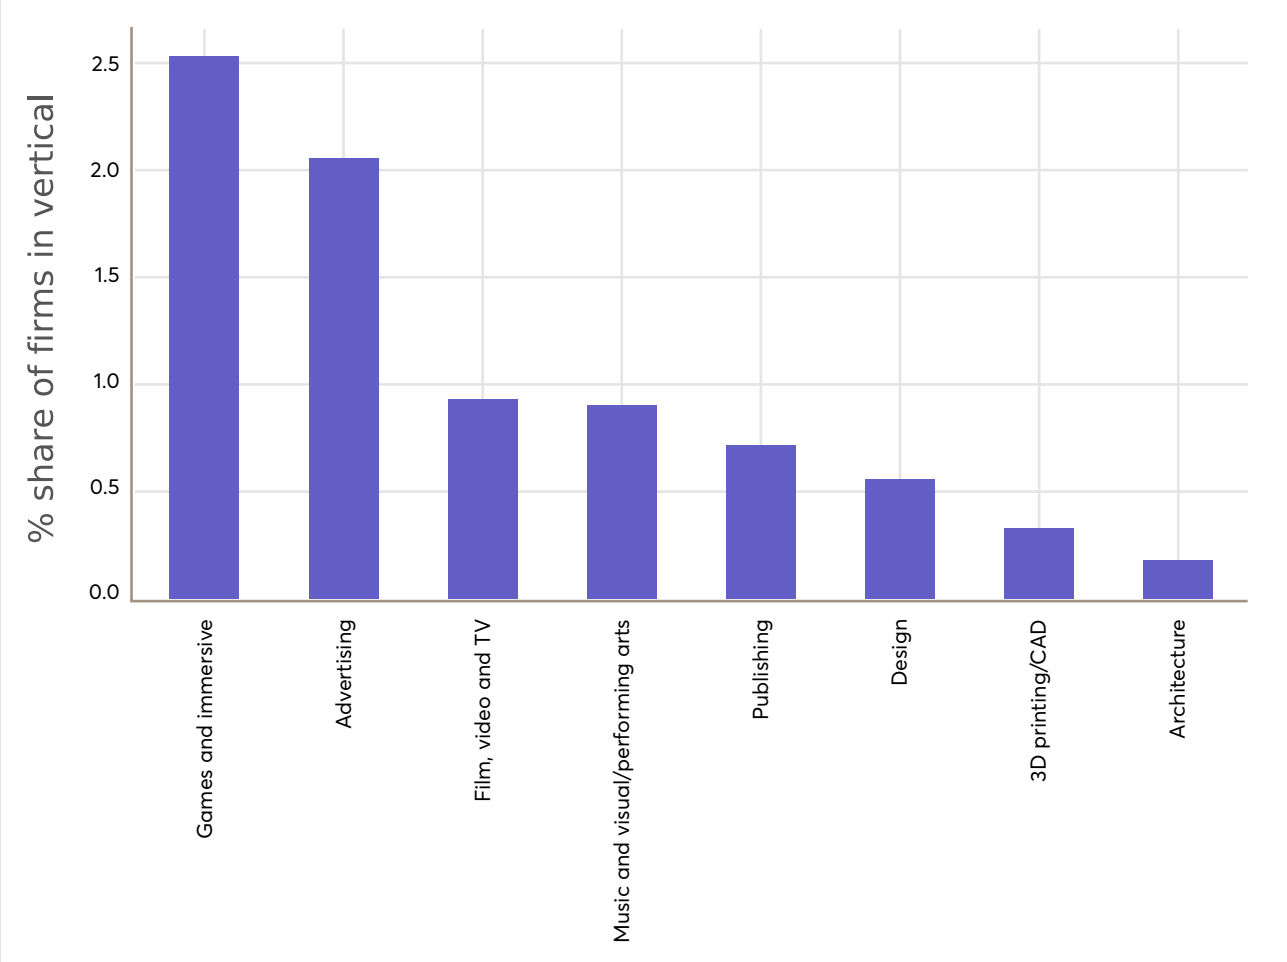


Figure 4.4 shows the number of AI companies in creative sectors by foundation date. As in the arXiv data, there is a significant increase of activity in the mid 2010s after the arrival of deep learning. Intriguingly, the levels of activity have declined after a peak in 2015, although this pattern is also visible in the broader Crunchbase dataset of all categories, consistent with the idea of a decline in levels of entrepreneurial activity in the USA and other countries. One possible explanation for this is investors switching to being keener in investing in more established companies, it could also be an effect specific to Crunchbase.[54] 2018 covers a full year, however perhaps there has been a delay in getting more recent companies added to the database at the time the data analysed was obtained (2019) which has reduced the numbers observed.

## Figure 4.4: The number of AI and creative companies by year

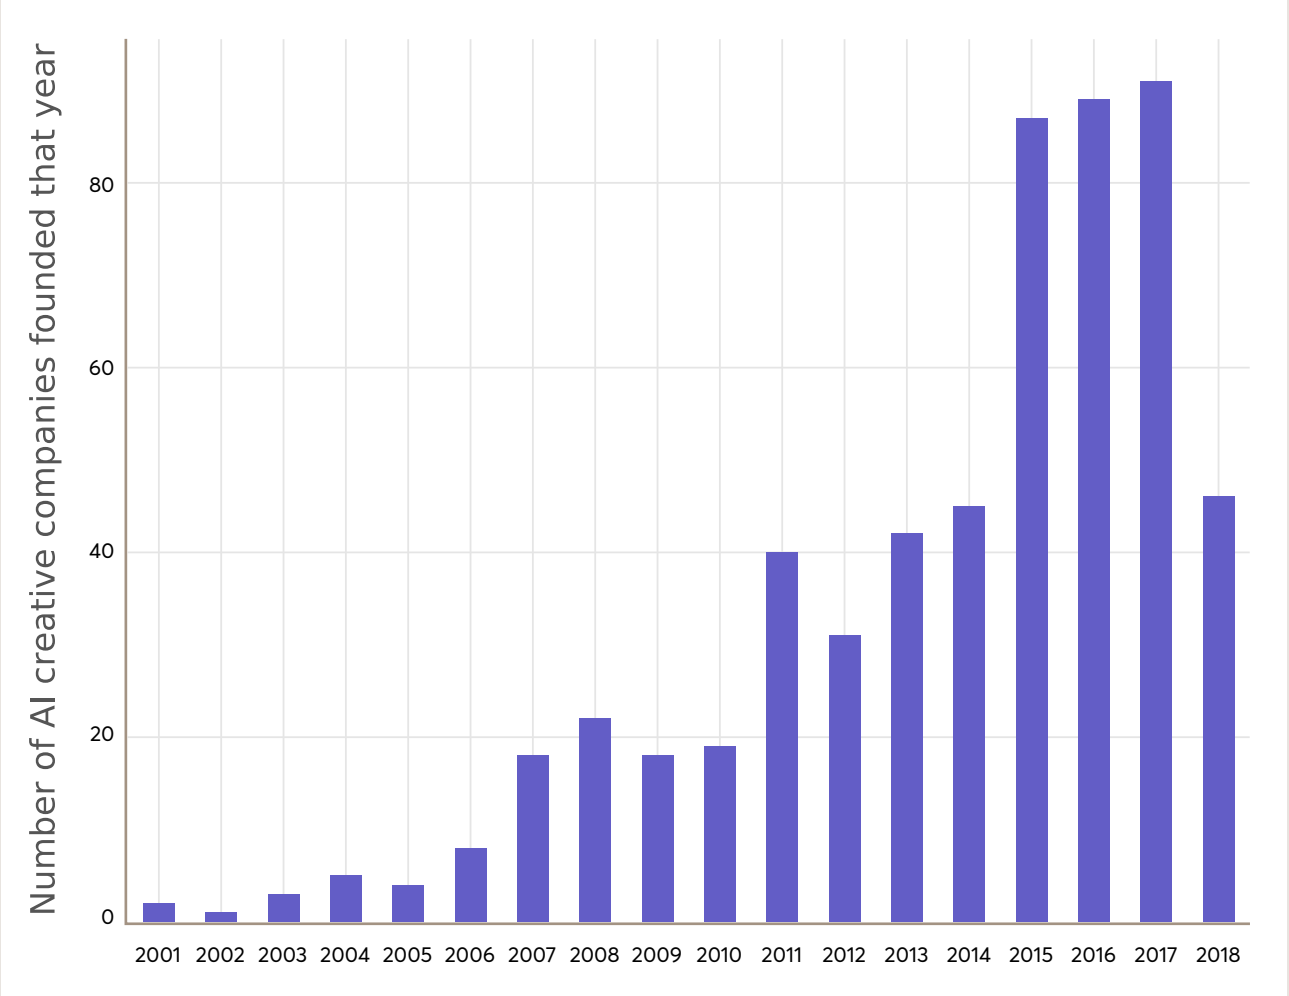


Figure 4.5 focuses on the sub-sectoral distribution of AI activity by year. It shows a diversification of creative industries AI activity away from advertising, which dominated the application of AI systems until the early 2010s, and the rapid growth in AI activity around video games and immersive technology. This pattern of a declining share for advertising in recent years was also be found in the arXiv data.

## Figure 4.5: The international distribution of AI companies in the different creative sub-sectors

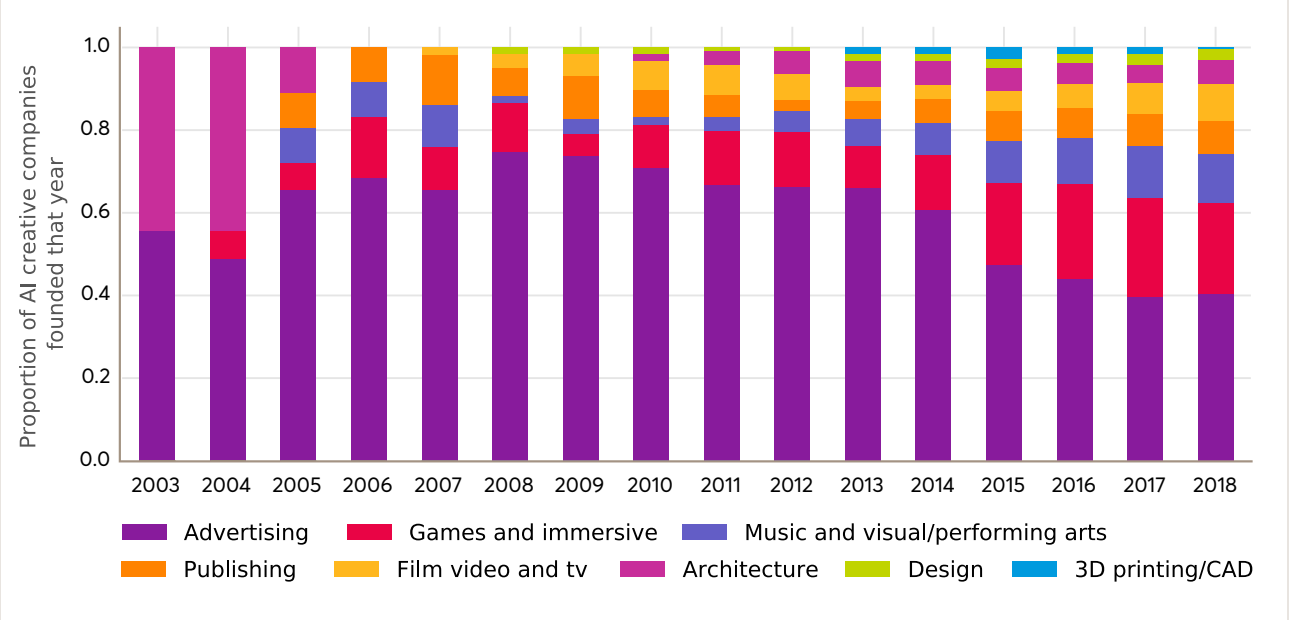


Figure 4.6 presents the international distribution of AI companies in the different creative sub-sectors. The US has the most companies 39 per cent of the companies, followed by the UK with 8 per cent of the companies and then India and Israel. The low level of this type of activity in China shown should be treated cautiously as Crunchbase is perceived to have lower coverage in China. Levels of activity across countries may also reflect variations in coverage more generally.

## Figure 4.6: The international distribution of AI companies in the different creative sub-sectors

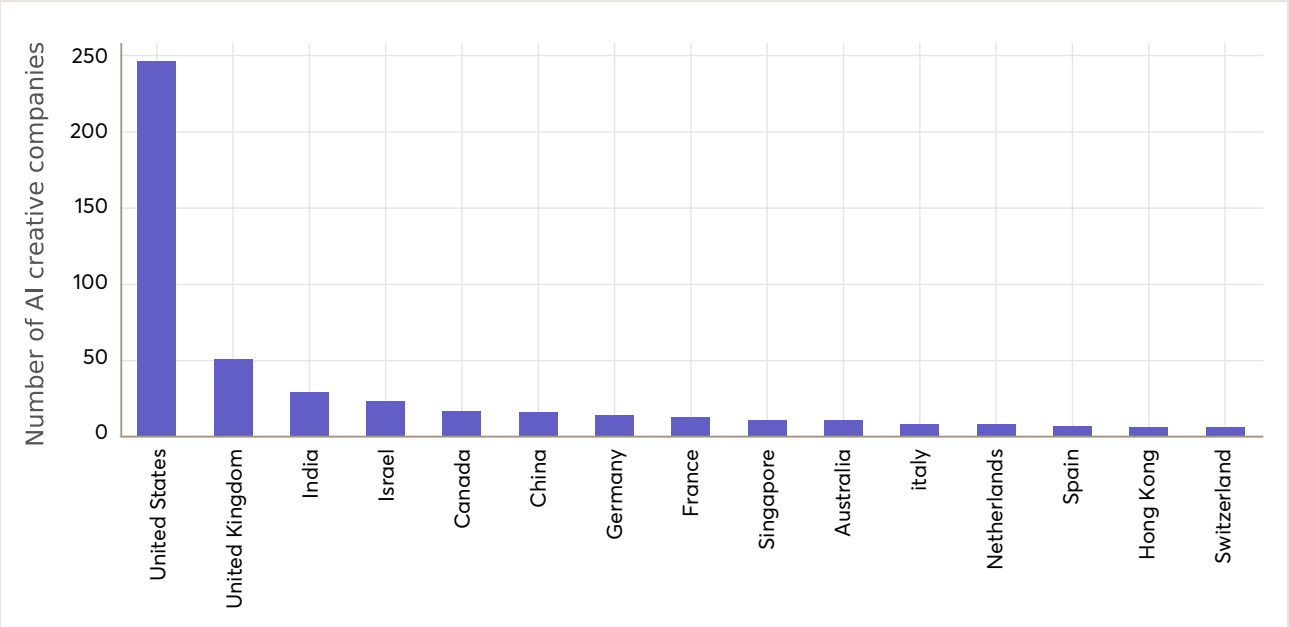

5

# Conclusions

**There are high levels of UK AI research in areas potentially relevant for creative industries, such as on image, text and sound data**. The UK is one of the world's leading producers of AI research overall. In AI research involving media relevant to the creative industries (sound, text and image data) the UK has the highest level of research activity in Europe.

**Levels of AI research with direct creative industry applications are relatively low internationally**. The evidence is that specific research directly involving creative activity and AI is relatively low. The number of papers is in the hundreds internationally, although the UK is one of the more active countries. Areas where comparatively higher levels of AI directly focussed on the creative industries have been identified are advertising, games, VR/AR and techniques like motion capture and photogrammetry.

**The UK is well positioned to take advantage of the role of AI in the creative industries, but other countries are rapidly increasing their levels of AI research activity**. Given the UK's research strength in AI combined with having one of the highest shares of creative industry employment in the European Union, the UK is well positioned to take a leading role in this area.[55] However there are no grounds for complacency with the research showing growing international competition in AI as more countries prioritise the area. Between 2015 and 2019 UK publications in AI grew at 365 per cent (A rate of growth that is similar to Italy, the US and Czechia), while by contrast among the fastest growing countries Taiwan grew at 1,490 per cent, Sweden at 958 per cent and Japan at 845 per cent. Although this growth is starting from lower levels of activity it shows how rapidly research in AI is growing around the world.

**There is evidence that direct applications of AI in creative industries are so far relatively low**. In terms of practical applications of AI in the creative industries, we are able to identify companies that appear to be involved with AI and creative industry domains in Crunchbase, and projects combining universities and organisations working on projects with AI and creative elements in Gateway to Research. However, the levels of this activity is still relatively low (in the tens of projects and companies in the UK). However in Crunchbase, only the US has more companies involved in this activity. This suggests that, at least in so far as it can be identified from text descriptions of companies and research projects, activity in this area is currently at a small scale, although there is evidence that the UK is relatively advanced in this regard relative to other countries.

## Policy recommendations

- The levels of research collaboration combining both AI and creative industries appears to be low on the evidence of arXiv and the Gateway to Research data. There also appear to be relatively low levels of companies in Crunchbase that combine AI and creative activity. At the same time, there are high levels of AI research in media such as text, image and sound that are highly relevant for the sector and in new approaches such as GAN and style transfer which are starting to see creative applications. If the UK wants to get the most out of its complementary strengths in AI and the creative industries it should work to support higher levels of commercial activity and academic research collaboration, combining both in future.

- There is evidence that there are extensive technical skills in AI in UK research institutions as demonstrated by research publications, but less evidence that these skills are being used in applications in creative industries. This indicates an opportunity for skill development by building partnerships between companies and academic institutions, perhaps through creative companies funding PhD studentships.

- There is limited information on the adoption of AI in creative businesses. Other data sources that might be used to examine this include looking at references to AI adoption in the trade press or other more indirect measures such as commits on GitHub from creative companies. Neither of these will fully capture the use of AI inside creative companies and this could be addressed with more detailed qualitative and survey work.

- One of the challenges of developing policy on AI is that the rate of pace in research on the topic makes it hard to keep track of what is happening. A benefit of taking a data-led approach to analysing the arXiv is that it is possible to start to identify AI techniques that are associated with particular domains as they emerge in research fields. This information may be common knowledge within academic research communities, but will often be largely unknown to policymakers and the wider business community outside. Using data to track the domains where technology like AI is being applied in academic research should be useful for both policy and businesses wishing to keep up with developments in AI relevant for the creative industries and other sectors.

6

# Appendix

The Word2vec embedding model used in the arXiv analysis involves mapping the words in the documents that are being analysed into a mathematical space (an embedding). In this space words are represented as points in the space, and the relative proximity of the words reflects their similarity. This allows us to find words that, the embedding space suggests, are similar in semantic meaning in the documents to the key words we initially chose. This helps us to identify other relevant documents beyond those based on our initial word selection. The words' location in the space is obtained from the weights of an intermediate layer of a neural network that is trying to predict, from the text the model is trained on, a word in the text from the words that surround it. This is known as a Continuous Bag of Words (CBOW) model.

The similarity between two words can be measured in terms of the distance between the words in the space using a measure called cosine similarity. This is defined mathematically as the dot product of the two vectors in the space that represent the words, divided by the product of the lengths of the two vectors. The value of this corresponds to a measure of how similar the words' meaning is. A measure of -1 means that the words are dissimilar and the measure of 1 corresponds to the cosine similarity of the word (and its associated vector) with itself i.e. it is identical. The closer to 1 the measure is the more similar the words are.

In the analysis we set the threshold of 0.85 for words similar to our initial chosen set of key words to identify each area to supplement them, and remove words that on inspection appear not to be associated with the areas we are interested in examining. In this section, as a robustness test, we see how sensitive the findings are to choosing a lower cosine similarity threshold of 0.75 to identify the categories i.e. a threshold where more words similar to our original keywords are used to identify the categories, but where the threshold for similarity is weaker. As this increases the number of words that can be used to identify the categories, this will increase the number of publications captured, although it also increases the likelihood of false positives.The findings from this analysis are shown below.

## Figure A1: Number of publications in different domains (lower cosine similarity of 0.75 used)

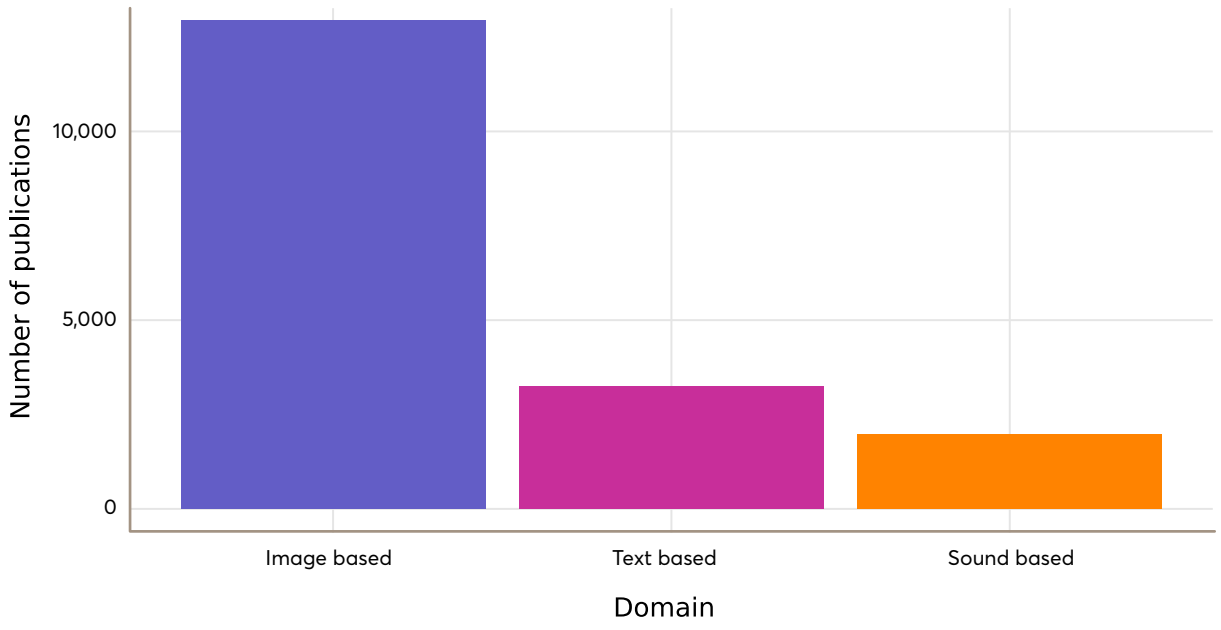


The number of papers identified increases, however, the overall finding of there being many more image based papers identified than text and sound based papers remains the same.

When we examine those papers which were also identified as involving AI by country (Table A1) the finding of the UK's relatively high ranking is retained.

## Table A1: AI publications in image, sound and text based media by country (lower cosine similarity of 0.75 used)

| | Country | Image based | Sound based | Text based | All three | % of all papers |
|---|---|---|---|---|---|---|
| 1 | United States | 3,341 | 377 | 850 | 4,282 | 26.4 |
| 2 | China | 1,596 | 106 | 298 | 1,874 | 32.8 |
| 3 | United Kingdom | 1,148 | 195 | 307 | 1,530 | 26.7 |
| 4 | Germany | 666 | 102 | 194 | 911 | 28.6 |
| 5 | Canada | 665 | 98 | 155 | 871 | 26.9 |
| 6 | Australia | 578 | 51 | 133 | 718 | 31.0 |
| 7 | India | 475 | 61 | 145 | 630 | 26.8 |
| 8 | Switzerland | 464 | 63 | 116 | 599 | 30.1 |
| 9 | France | 450 | 67 | 110 | 596 | 24.1 |
| 10 | Ireland | 414 | 60 | 123 | 560 | 22.5 |

On the more specific creative categories, if we apply the lower 0.75 similarity threshold, then we obtain the figure below for overall paper counts.

**Figure A2: Number of research publications by creative industry category activity (lower cosine similarity of 0.75 used)**

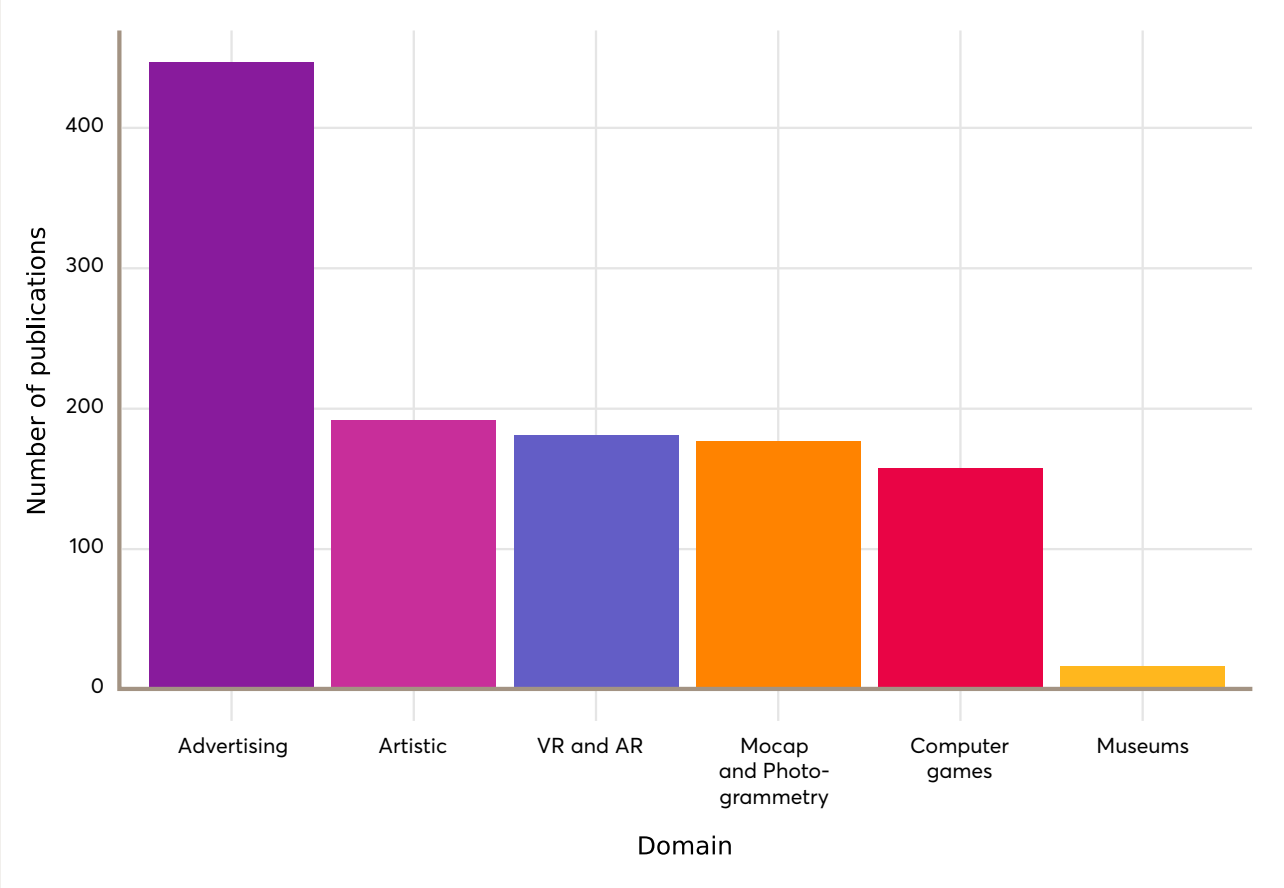


The effect of the lower threshold does in percentage terms, with the exception of museums, increase the number of papers for each category significantly, but as the original number of papers at the higher threshold was not that large previously, the absolute number of papers is still relatively low. The main exception is advertising where there is a very significant increase in the number of papers identified. This is probably due to advertising being closely related to large parts of the digital economy which our original analysis on a narrower criterion was not picking up. This suggests that the number of advertising research papers is much larger when this aspect of advertising is taken into account. When we examine those papers which were identified as involving AI by country (Table A2) the finding of the UK's relatively high ranking is still found.

## Table A2: Number of research publications by country on creative related AI activity (lower cosine similarity of 0.75 used)

| Country | Advertising | Artistic | Computer games | Mocap and photo-grammetry | Museums | VR and AR | All types | % of all papers |
|---|---|---|---|---|---|---|---|---|
| United States | 121 | 41 | 40 | 36 | 2 | 40 | 274 | 1.7 |
| China | 29 | 21 | 10 | 11 | 0 | 13 | 81 | 1.4 |
| United Kingdom | 25 | 12 | 13 | 16 | 0 | 11 | 76 | 1.3 |
| Australia | 18 | 8 | 6 | 5 | 0 | 3 | 40 | 1.7 |
| Canada | 12 | 6 | 9 | 7 | 0 | 6 | 39 | 1.2 |
| Italy | 10 | 6 | 6 | 10 | 1 | 2 | 34 | 2.0 |
| Germany | 4 | 2 | 2 | 14 | 0 | 9 | 29 | 0.9 |
| Switzer-land | 9 | 3 | 5 | 8 | 0 | 3 | 25 | 1.3 |
| France | 9 | 1 | 3 | 5 | 0 | 6 | 24 | 1.0 |
| India | 12 | 1 | 1 | 3 | 0 | 4 | 21 | 0.9 |

7

# References


Bakhshi, H., Frey, C. and Osborne, M. (2015), 'Creativity vs Robots. The creative economy and the future of employment', Nesta.

Bakhshi, H. and Lomas, E. (2017), 'Defining R&D for the creative industries', Nesta.

Chollet, F. (2017), 'Deep Learning with Python', Manning.

Christie's (2018), 'Is artificial intelligence set to become art's next medium?', Christie's.

Covington, P., Adams, J. and Sargin, E. (2016), 'Deep Neural Networks for YouTube Recommendations', Google.

Creative Industries Council (CIC) in association with Tech London Advocates (2018), 'Createch resource book', CIC.

Davies, J. and Ward Dyer, G. (2019), 'The relationship between artistic activities and digital technology development', European Parliament.

Davies, J. (2016), 'State of the Art: analysing where art meets technology using social media data', Nesta.

Davies, J. (2018), 'Freeing our minds: Three ways digital technology could radically change art', Nesta.

Easton, E. and Djumalieva, J. (2019), 'Creativity and the future of skills', Creative Industries Policy and Evidence Centre.

The Economist (2017), 'The Battle of the Brains Google leads in the race to dominate artificial intelligence'.

Eliot, L. (2017), 'Creative AI landscape', Medium.

Eliot, L. and Zhegin, P. (2017), 'How to take AI far beyond gaming', Venturebeat.com.

Eliot, L. (2016), 'Creative machines, humans and cyborgs at Retune 2016', Fad magazine.

DCMS (2019), 'DCMS Sector Economic Estimates Methodology'.

Gatys, L., Ecker, A. and Bethge, M. (2016), 'Image Style Transfer Using Convolutional Neural Networks', Conference on Computer Vision and Pattern Recognition.

Geron, A. (2017), 'Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems', O'Reilly.

Goodfellow, I. (2016), 'NIPS 2016 Tutorial: Generative Adversarial Networks'.

Goodfellow, I., Bengio, Y. and Courville, A. (2017), 'Deep Learning', MIT press for a summary of recent developments.

HM Government (2018), 'Industrial Strategy Building a Britain fit for the future'.

Hall, W. and Pesenti, J. (2018), 'Growing the artificial intelligence Industry in the UK', UK Government.

Inco8 (2019), 'Generated Photos: 100,000 Free AI Generated Faces for Your Designs', Icon8 blog.

Klinger, J., Mateos-Garcia, J. and Stathoulopoulos, K. (2018), 'Deep Learning, Deep Change? Mapping the Development of the Artificial Intelligence General Purpose Technology', Nesta.

Langr, J. and Bok, V. (2019), 'GANs in Action: Deep learning with Generative Adversarial Networks', Manning.

Marchesi, M. (2018), 'Happy Thoughts: Style Transfer and the Impact of AI on Creativity', Happy Finish.

Martel, F. (2019), 'Cultural Policies: Mapping a Field in Reinvention', ZCCE.

Mateos-Garcia, J. (2018), 'Mapping Research & Innovation Missions', Nesta.

Pangburn, DJ. (2018), 'How Aphex Twin's 'T69 Collapse' video used a neural network for hallucinatory visuals', Fastcompany.

Nathan, M., Pratt, A. and Rincon-Aznar, A. (2015), 'Creative Economy Employment in the European Union and the United Kingdom: A Comparative Analysis', Nesta.

Prince, S. (2012), 'Computer Vision. Models, Learning and Inference', Cambridge University Press.

Radke, R. (2013), 'Computer Vision for Visual Effects', Cambridge University Press.

Rehurek, R. and Sojka, P. (2010), 'Software framework for topic modelling with large corpora', The Lrec 2010 Workshop On New Challenges For Nlp Frameworks, 45-50.

Serpentine (2020), 'Future Art Ecosystems'.

Seymour, M. (2018), 'AI for VFX at SIGGRAPH Part 1', fxguide.

Stathoulopoulos, K. and Mateos-Garcia, J. (2019), 'Gender Diversity in AI Research', Nesta.

Ward Dyer, G. (2017), 'With AI emerging as a game-changer for the creative sector, the winner of the 2018 Turner Prize could be an AI and artist duo', Nesta.

8

# Endnotes

1. The report and analysis was produced by John Davies and Juan Mateos-Garcia. The data collection infrastructure to collect the data that is analysed was developed by Joel Klinger. The approach to defining AI that is used was developed by Kostas Stathoulopoulos.
2. The term AI originates from John McCarthy in 1955 when he organised the first conference in the field. McCarthy, J., Minsky, M., Rochester, N. and Shannon, C. (1955), 'A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence'.
3. See the introduction of Goodfellow, I., Bengio, Y. and Courville, A. (2017), 'Deep Learning', MIT press for a summary of recent developments.
4. Hall, W. and Pesenti, J. (2018), 'Growing the artificial intelligence Industry in the UK', UK Government.
5. Nathan, M., Pratt, A. and Rincon-Aznar, A. (2015), 'Creative Economy Employment in the European Union and the United Kingdom: A Comparative Analysis', Nesta. DCMS, 'Creative industries economic estimates'.
6. HM Government (2018), 'Industrial Strategy Building a Britain fit for the future'.
7. Covington, P., Adams, J. and Sargin, E. (2016), 'Deep Neural Networks for YouTube Recommendations', Google.
8. arXiv computer science landing page. https://arxiv.org/archive/cs
9. DCMS (2019), 'DCMS Sector Economic Estimates Methodology'.
10. Klinger, J. and Mateos-Garcia, J. and Stathoulopoulos, K. (2018), 'Deep Learning, Deep Change? Mapping the Development of the Artificial Intelligence General Purpose Technology'. Available at SSRN: https://ssrn.com/abstract=3233463 or http://dx.doi.org/10.2139/ssrn.3233463
11. The Economist (2017), 'The Battle of the Brains Google leads in the race to dominate artificial intelligence'.
12. Stathoulopoulos, K. and Mateos-Garcia, J. (2019), 'Gender Diversity in AI Research', Nesta. https://media.nesta.org.uk/documents/Gender_Diversity_in_AI_Research.pdf
13. IDF is log (1/ (the number of documents that the term appears in/the number of documents in the dataset)). It adjusts for the fact that some terms are quite generic i.e. appear in many documents and so have limited information value by downweighting them.
14. To do this a measure of cosine similarity above 0.85 was used.
15. Rao, S. (2018), 'MITOS-RCNN: A Novel Approach to Mitotic Figure Detection in Breast Cancer', arXiv:1807.01788.
16. Bin Ahmed (2017), Deep Learning based Isolated Arabic Scene Character Recognition, arXiv:1704.06821.
17. Hu, A. and Flaxman, S. (2018) 'Multimodal Sentiment Analysis To Explore the Structure of Emotions', arXiv:1805.10205.
18. Gabbay et al. (2018), 'Seeing Through Noise: Visually Driven Speaker Separation and Enhancement', arXiv:1708.06767.
19. Cudeiro et al. (2019), 'Capture, Learning, and Synthesis of 3D Speaking Styles', arXiv:1905.03079.
20. Haghani et al. (2018), 'From Audio to Semantics: Approaches to end-to-end spoken language understanding', arXiv:1809.09190.
21. Dey et al. (2018),'Topical Stance Detection for Twitter: A Two-Phase LSTM Model Using Attention', arXiv:1801.03032.
22. Wen et al. (2016), 'Learning text representation using recurrent convolutional neural network with highway layers', arXiv:1606.06905.
23. We highlight that in all the analysis that follows, we are not allocating papers exclusively to a single creative sub-sector.
24. Prince, S. (2012), 'Computer Vision. Models, Learning and Inference', Cambridge University Press.
25. For a discussion on the techniques that are used in visual effects see Radke, R. (2013), 'Computer Vision for Visual Effects', Cambridge University Press.
26. Liu et al. (2019), 'A Semi-Supervised and Inductive Embedding Model for Churn Prediction of Large-Scale Mobile Games', arXiv:1901.06247.
27. Guo et al. (2016), 'Deep Learning for Reward Design to Improve Monte Carlo Tree Search in ATARI Games', arXiv:1604.07095.
28. Paumard, M., Picard, D. and Tabia, H. (2018), 'Jigsaw Puzzle Solving Using Local Feature Co-Occurrences in Deep Neural Networks', arXiv:1807.03155.
29. Thies et al. (2016), 'FaceVR: Real-Time Facial Reenactment and Eye Gaze Control in Virtual Reality', arXiv:1610.0315.

30. Saleh, B. and Elgammal, A. (2015), 'Large-scale Classification of Fine-Art Paintings: Learning The Right Metric on The Right Feature', arXiv:1505.00855.

31. Balazia, M. and Sojka, P. (2017), 'An Evaluation Framework and Database for Gait Recognition Methods', arXiv:1701.00995.

32. Davies, J. and Ward Dyer, G. (2019), 'The relationship between artistic activities and digital technology development', European Parliament.

33. Goodfellow, I. (2016), 'NIPS 2016 Tutorial: Generative Adversarial Networks'.

34. Christie's (2018), 'Is artificial intelligence set to become art's next medium?', Christie's.

35. Seymour, M. (2018), 'AI for VFX at SIGGRAPH Part 1', fxguide.

36. Inco8 (2019), 'Generated Photos: 100,000 Free AI Generated Faces for Your Designs', Icon8 blog.

37. Gatys, L., Ecker, A. and Bethge, M. (2016), 'Image Style Transfer Using Convolutional Neural Networks', Conference on Computer Vision and Pattern Recognition.

38. Marchesi, M. (2018), 'Happy Thoughts: Style Transfer and the Impact of AI on Creativity', Happy Finish. Pangburn, DJ. (2018), 'How Aphex Twin's 'T69 Collapse' video used a neural network for hallucinatory visuals', Fastcompany.

39. Sharma, A. and Hamarneh, G. (2019), 'Missing MRI Pulse Sequence Synthesis using Multi-Modal Generative Adversarial Network', arXiv:1904.12200.

40. Hernandez Ruiz et al. (2019), 'Human Motion Prediction via Spatio-Temporal Inpainting', arXiv:1812.05478.

41. Ramakrishnan et al. (2017), 'Deep Generative Filter for Motion Deblurring', arXiv:1709.03481.

42. Zhang et al. (2017), 'Style Transfer for Anime Sketches with Enhanced Residual U-net and Auxiliary Classifier GAN', Xiv:1706.03319.

43. So, C. (2018), 'A Pragmatic AI Approach to Creating Artistic Visual Variations by Neural Style Transfer', arXiv:1805.10852.

44. Korshunova et al. (2016), 'Fast Face-swap Using Convolutional Neural Networks', arXiv:1611.09577.

45. Lu et al. (2018), Play as You Like: Timbre-enhanced Multi-modal Music Style Transfer, arXiv:1811.12214.

46. Ghiasi, G. (2017), Exploring the structure of a real-time, arbitrary neural artistic stylization network', arXiv:1705.06830.

47. Mateos-Garcia, J. (2018), 'Mapping Research & Innovation Missions', Nesta.

48. Algorithms on rank aggregation for preference orderings

49. Real-time Visuals for Performance, Gaming, Installation, and Electronic Environments

50. Live Coding Network

51. EPSRC NetworkPlus on Social Justice through the Digital Economy

52. Sandpit: VoiceYourView – Making Public Spaces Safer

53. More specifically, this includes terms such as 'data science', 'machine learning', 'deep learning', 'artificial intelligence', 'neural network', 'ai ', 'natural language processing' and 'text mining'.

54. Spindler, J. (2018), 'Decline in Seed Investments and why it matters?', AI seed. Basul, A. (2019), 'KPMG report shows majority of VC investment spent on larger scaleups', UKTN.

55. Nathan, M., Pratt, A. and Rincon-Aznar, A. (2015), 'Creative Economy Employment in the European Union and the United Kingdom: A Comparative Analysis', Nesta. DCMS, 'Creative industries economic estimates'.

## Acknowledgements

Thanks go to **Hasan Bakhshi**, **Eliza Easton**, **Mick Grierson**, **Jeremy Silver** and **Bruce Tether** for their comments on the document. Thanks also go to the Creative Economies Centre of Zurich University of the Arts (ZCCE) who supported work that informed this research at an early stage.

**The Creative Industries Policy and Evidence Centre (PEC) works to support the growth of the UK's Creative Industries through the production of independent and authoritative evidence and policy advice.**

Led by Nesta and funded by the Arts and Humanities Research Council as part of the UK Government's Industrial Strategy, the Centre comprises of a consortium of universities from across the UK (Birmingham; Cardiff; Edinburgh; Glasgow; Work Foundation at Lancaster University; LSE; Manchester; Newcastle; Sussex; Ulster). The PEC works with a diverse range of industry partners including the Creative Industries Federation.

For more details visit **www.pec.ac.uk** and **@CreativePEC**

If you'd like this publication in an alternative format such as Braille, or large print, please contact us at: **enquiries@pec.ac.uk**

**Creative Industries**
**Policy & Evidence Centre**

Led by **nesta**

Creative Industries Policy and Evidence Centre (PEC)
58 Victoria Embankment
London EC4Y 0DS

+44 (0)20 7438 2500

enquiries@pec.ac.uk

@CreativePEC

www.pec.ac.uk

9 781913 095192 >