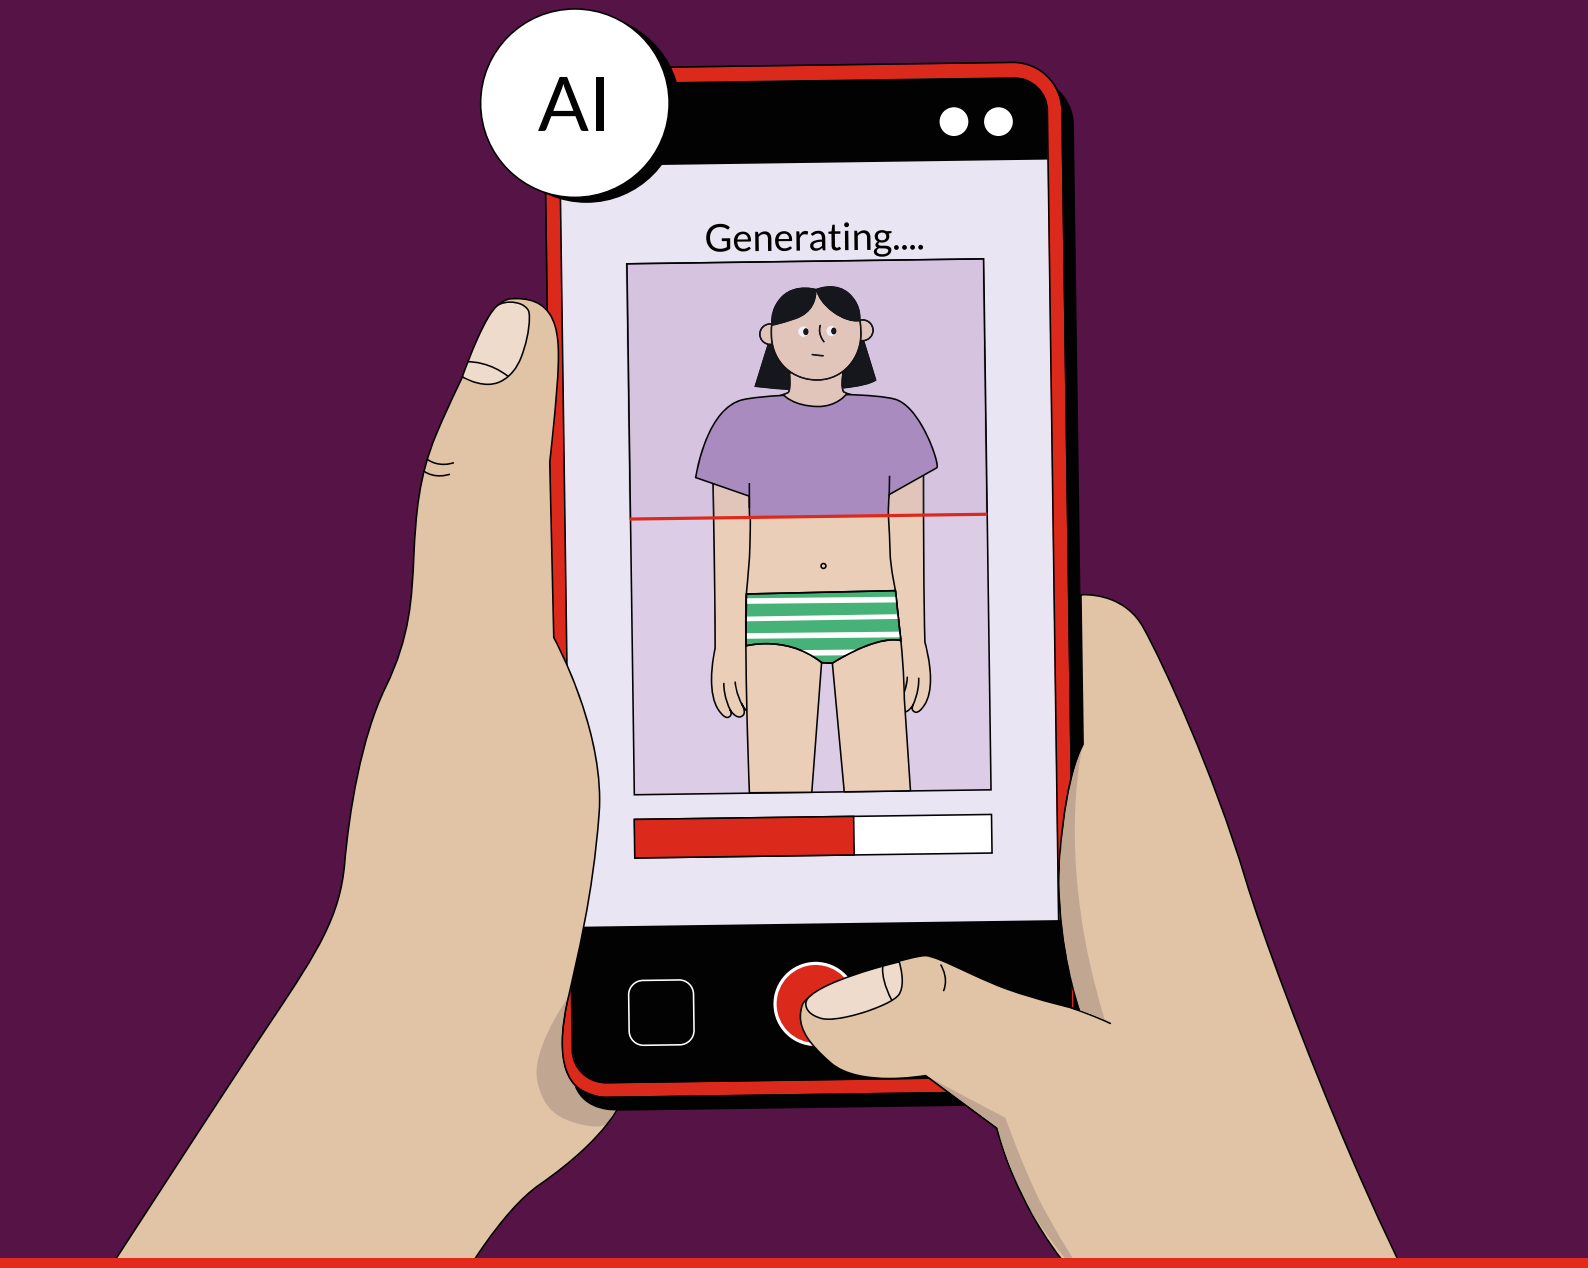


When machines learn too much

Findings from dark web discussions regarding the production of AI-generated child sexual abuse material and the technologies used for that purpose

Save the Children

Contents

Key concepts 3

Introduction 4

1 Children's images on social media and digital footprint 5

2 The use of children's images as source material for AI 6

3 Conducting the study 7

4 Findings 8

4.1 General opinion of participants in dark web forums regarding AI-generated CSAM 8

4.2 Community, praise and enthusiasm 11

4.3 Styles and preferences 14

4.4 Ethics and justification 16

4.5 The Future of AI 21

4.6 The use of AI to depict former victims or well-known individuals 22

4.7 The quality of AI-generated CSAM and related discussions 23

4.8 Discussions related to AI technologies and their use 24

5 Discussion 26

6 Needs for further research 28

7 Recommendations 29

Sources 31

Appendices 32

Key concepts

Artificial intelligence, often abbreviated to AI. Artificial intelligence refers to a computer programme's ability to perform intelligent tasks that mimic the human mind. Examples of such tasks include speech recognition, language translation, modelling and the creation of audiovisual material.

Generative artificial intelligence refers to AI that creates new content, such as text, images and videos, based on prompts provided by the user, utilising generative models and training data. Generative AI utilises deep learning – considered an advanced form of machine learning – to create new content.

The dark web is also known as the 'anonymous web'. It refers to content on the internet that requires special software, settings or actions to access. The aim of the dark web is to make users largely anonymous. Traffic on the dark web is usually encrypted and routed through various servers.

Dark web forum that is focused on child sexual abuse generally refers to a discussion platform specialising in sexual abuse or sexual exploitation of children. The term child sexual exploitation forum or CSE forum is often used.

'Child sexual abuse material' is a term used to describe material in which a minor is subjected to sexual abuse. The term is often referred to by its English abbreviation, CSAM (child sexual abuse material).

'Artificial intelligence-generated child sexual abuse material' is a term referring to material depicting sexual violence against children that has been produced using artificial intelligence applications or platforms. In this report, the term is abbreviated to **'AI-generated CSAM'**, where the acronym CSA stands for 'child sexual abuse'.

Introduction

Generative artificial intelligence (AI) has emerged as one of the most significant technological innovations of the 2020s. It enables the automatic generation of text, images, audio, videos and programme code, which has opened up new opportunities in areas such as education, research, healthcare and the working life. Although generative AI came to the attention of the general public particularly through services such as ChatGPT, its development is based on scientific breakthroughs that have emerged over decades, such as neural networks, deep learning, transformer architecture and large language and image models.

At the same time, the rapid development of this technology has given rise to new societal and security-related issues. A particularly significant concern is the misuse of AI to harm children and young people, and to produce and disseminate harmful digital content. Generative models enable the creation of highly realistic digital content, making it difficult to distinguish between real and manipulated material. As a result, authorities, technology companies, researchers and non-governmental organisations have had to develop new approaches to ensure children's safety in digital environments.

The aim of this report is to strengthen the capacity of professionals working with children to recognise the AI-related dimensions of child sexual abuse and to understand their impact on children's safety and protection. Of particular interest are AI-based technologies that can be used to produce or modify harmful digital content, and their impact on the online criminal environment.

The aim of the report is to explore how participants in dark web forums focusing on child sexual abuse view AI-generated child sexual abuse material (CSAM). Furthermore, the aim is to examine the views expressed on these forums regarding the technologies used to produce such material, as well as how participants assess the potential, limitations and impacts of these technologies.

1 Children's images on social media and digital footprint

It is common for parents and other close ones to share recognisable images of children on platforms such as social media and instant messaging services. Sometimes, images are shared publicly by a representative of a school, sports club, nursery, or similar organisation.

According to some international estimates, parents share 1,500 images or videos of their children online before the child turns five (Unicef). A large part of children have not given their consent for their images to be shared online and, due to their age, they are not always even capable of giving it.

A digital footprint refers to all the data and information that every internet user leaves behind online, for example on social media, websites and services. This includes, amongst other things, all images, videos, comments, likes, live streams, clicks, and purchases. In addition to actual internet users, data is also left behind about those who do not (yet) use the internet themselves, such as children, whose images and information are shared by adults. Sometimes, a child may be on social media themselves or, for example, playing online games, in which case they also create a digital footprint through these activities.

Sharing children's images and other data online means that children may have a digital identity before they even understand what 'digital' means. However, adults need to understand that children have a right to privacy. Social media is a relatively new phenomenon, so it is difficult to predict how making children's private lives public online might affect their safety as they are growing up or when they become adults.

2 The use of children's images as source material for AI

In everyday images, children are doing things that are typical and appropriate for their age. They play, swim, do their homework and sleep. There is a vast amount of these everyday images of real-life situations, as well as images of children taken specifically for advertising purposes, available online. Such everyday images and other data relating to children are used to train various forms of AI. AI can be used to create entirely legal images and videos featuring child characters, but it can also be used to produce CSAM.

Anyone can create sexual material from any everyday image or video of a child in a number of different ways. A child's voice can also be used to create sexual material, if it is available. This often results in fairly convincing sexual material featuring a child (a 'deepfake'). In practice, AI can be used to create a video of virtually any child in which an adult character or 'virtual copy' of an adult subjects the child to sexual acts, such as raping the child.

Using AI, existing AI-generated CSAM that has already been shared online can be endlessly re-edited. The child in the material may have already grown into an adult, but the sexual abuse directed at them online does not cease.

An image or character of an adult can also be age regressed into a child using AI and placed in a sexual situation. The character or virtual copy of any adult can also be placed in a sexual context with a child.

3 Conducting the study

The report aimed to examine two different perspectives relating to AI-generated CSAM. Firstly, the aim was to investigate the nature of discussions taking place on dark web forums that are focused on child sexual abuse regarding AI generated CSAM. In particular, the interest was in how participants perceive AI generated CSAM and whether any recurring themes or clear consensus on the topic can be observed in the discussions. The first research question became: *'How do participants in dark web forums focused on child sexual abuse perceive AI-generated CSAM?'*.

The second aspect of the report focused on AI technologies. In particular, it was of interest what participants in the discussions thought about AI technologies, whether some technologies recurred in the discussions, and to what extent participants are willing to share AI technologies or instructions on how to use them. The second research question became: *'How do participants in dark web forums focused on child sexual abuse perceive the technologies used to produce AI-generated CSAM?'*.

Data collection and search terms

The data to be analysed was collected using software provided by an external IT company during April 2025. The results included the entire discussion, so some of the final results may predate or postdate the search period used, although the focus remains on the selected search period. Only English-language discussions were included in the data for analysis.

The material used in the analysis of this report was collected without the participants in the discussions being aware that their comments would form part of the report. This approach was adopted to ensure that the data required for analysis could be obtained. It is not possible to identify individual participants from the data.

Some of the search terms are based on previous research or reports on the subject, such as Kokolaki and Fragopoulou (2025) and the Internet Watch Foundation (2023 and 2024). In addition, search terms were selected by randomly monitoring discussions on dark web forums and chat rooms and identifying key, frequently recurring terms.

Before the final selection of search terms, a set of search terms was tested by using them to retrieve search results. The results obtained were preliminarily assessed, and the 20 best performing search terms were selected. The search terms are explained in the glossary attached to the report.

The 20 search terms selected for the report: ai, ai generate, ai guide, ai prompt, prompt, lora, deepfake, stable diffusion, stable diffusion sd, flux.1, midjourney, dall-e, diffusion model, fine-tuning, nsfw, checkpoint models, checkpoint, realistic, cfg, and exif.

4 Findings

The data was approached using broad research questions. The various headings emerged as the data was examined and categorised by theme: the headings describe recurring patterns within the data. The decision was made to include, as extensively as possible, recurring patterns and themes emerging from the forums relating to AI-generated CSAM. Not all sections necessarily correspond exactly to the research question. The findings are drawn from five different dark web forums, one of which focuses exclusively on AI-generated material.

The findings section uses direct quotations from discussions on the forums. For ethical reasons, certain parts of the quotations have been redacted. Listed below are the parts redacted from the quotations and what they have been replaced with.

- A well-known female victim of sexual abuse → Jane Doe
- A well-known male victim of sexual abuse → Joe Doe
- A girl/woman known to the public → Jane Doe
- A boy/man known to the public → Joe Doe
- Username of a person participating in discussions → user X
- Forum name → forum
- Links to files → link
- Name of an AI model → model X

4.1 GENERAL OPINION OF PARTICIPANTS IN DARK WEB FORUMS REGARDING AI-GENERATED CSAM

Based on the data as a whole, it can be said that AI-generated CSAM is well-regarded and its potential is recognised. People are aware of its scope for development, its sense of authenticity and the ease with which it can be created, and there are increasing calls for more of it to be produced, with creation taking place at an ever-increasing pace. The comments repeatedly expressed admiration and praise for the material created and its creators. Many commented that the material was better than any real material seen before and the best they had ever seen. Many also found it difficult to distinguish AI-generated material from real material or, alternatively, whilst they knew the material had been AI-generated, they commented that it was, in practice, just like real material.

> ***"Love AI. It can be whatever you want."***
>
> ***"this is hotter than any real porn i've seen!"***
>
> ***"AI is really becoming very close to being as good as if not better than the real thing!"***
>
> ***"I don't know what to say any more. Over the last years AI pictures were getting better and better but that they will get this perfect so soon I never thought."***

> ***"It took me about 3 1/2 minutes to make these 2 images. No Loras used and no editing."***

There was a clear divide in the opinions expressed on the forums, which may be explained by the fact that one of the examined forums focused exclusively on AI-generated material. It can be assumed that this forum attracts people who have a positive attitude towards AI-generated material and are specifically looking for it. On the more general forums, attitudes towards AI-generated material were at times more negative: some felt it was entirely undesirable, whilst others even wrote in an aggressive tone if there was even the slightest suspicion that published material had been generated by AI. On the more general forums, there might be a dedicated channel for AI-generated material, and it was not necessarily tolerated on the forum's other channels.

The data contained recurring comments in which AI-generated and real material were not distinguished from one another. Occasionally, material was posted that mixed real and AI-generated material. Sometimes this was disclosed at the time of posting, but at other times only if someone asked directly.

> ***"We need to separate real content from fictional content. Otherwise we don't know what's real or not and we will all turn insane."***

Some commenters on the general forums wrote that CSAM must be real, as only knowledge of actual child sexual abuse arouses them. Some were of the opinion that AI-generated material is only to the taste of the masses or people with a low IQ.

> ***"Everybody has his own taste. As for me, AI CP is not inspiring for me at all. For excitement and stimulation I need real action by real kids."***
>
> ***"AI seems to be extremely popular among low IQ people from third world countries."***

However, AI-generated material that was posted also received many positive comments from those who were not previously familiar with it or who mainly consumed real material. Some wanted to consume both real and AI-generated material in the future and were surprised by the quality of the AI-generated material.

> ***"Wow, I have to say I am impressed! This is better quality than what I've seen before and I generally just skip on right by anything AI. (...) I would definitely want to see more of this personally."***

AI-generated material was viewed favourably under certain conditions. It was felt that real material was needed to train the AI. AI-generated videos were also often seen as still being so short or otherwise flawed that they sometimes received critical feedback. However, there was firm belief in their future development, and comments highlighted that full-length AI films, for example, would soon be commonplace.

> ***"But we got one dilemma, what about training data? For AI to understand sexual acts and the child anatomy it must learn from real CP, right? To get that 9/10 realism the body parts have to be extremely accurate. I believe we need the real stuff to train models."***

People were also prepared to give up material featuring real children in favour of convincing AI-generated content. This was due, for example, to the fact that AI can create any kind of content and was felt to better be able to fulfil one's own fantasies, or because in the real material, the children looked bored or frightened.

> ***"I'm honestly to the point were I'm willing to give up the real pedo porn and just look at and watch ai generated ones. I know that sounds crazy but iv deleted most of my content to replace it with ai. AI in general can create all of my fantasies, its not perfect but soon it will be. I struggle sometimes wondering if real content is fake and if fake is real."***

Most of the comments that the posted AI-generated material and its creators received were complimentary, even effusive. AI-generated material was seen as a means of realising all fantasies, as AI can produce almost anything. At times, it was also seen as a safer alternative to real material: people were aware of the shortcomings and variability in AI legislation. The enthusiasm was pervasive in many ways. People were excited both by the material itself – its realism and its ability to arouse the viewer – and by the sophistication of AI technologies and their potential to create scenarios that might never be seen in real material. AI was seen as a tool of a new era, capable of fulfilling everyone's fantasies. Some saw AI-generated material as their only option and hoped that it would supplant all real material.

> ***"THAT is the AI Pedo-Dream COME TRUE!!!!!!!!!!!!!!!!!!"***
>
> ***"Thank you SO MUCH!!! And please keep making more of these: Those are exactly what I always hoped AI would lead to when it first came about a few years back!"***
>
> ***"I can't believe my eyes. This is AI. This is fantastic. My endless gratitude to the author for these masterpieces."***
>
> ***"I've been thinking more and more about the fact that I am jerking off more and more to AI CP than real CP and that I am enjoying it even more than the real thing! The customization of situations, locales, models' bodies and facial expressions, etc. are creating the sort of things that we never or almost never see in real CP and the "mistakes' of the technology are getting less and less noticeable with each new generation of AI productions (...)"***

AI is used to create not only images and videos but also stories and virtual partners. Sexual conversations are held with AI, in which the AI is asked to appear as a child of a certain age. It was reported that AI is trained not only on real CSAM but also on everyday material available on social media.

> ***"I did create an AI 12 year old wife so to say using [model x] lol and I do feel so damn immersed with her like the way she talks, chilldish humour etc it's so damn immersive... Like I actually feel like I'm taking care of this imaginery virtual little girl in my head. Perhaps it gives me a glimpse of what it could be like if it was real?"***
>
> ***"Now for faces, Social media postings are open to the training models for the faces. Whatever is uploaded to Facebook, Instagram & whatever is available for training..."***
>
> ***"Yeah i have been on the social media train for 8-10 years now and get most of my material from there from there during the last decade. In the very beginning was great, then things toned down a bit after google got in trouble for helping pedos find similar content - hell goodle did nothing wrong and it's algorithm worked just like it was supposed to. After that died down things are only getting better."***

Overall, there was strong justification for the creation and use of AI-generated material, and it was even seen as a way of protecting children. Comments that AI-generated material does not harm real children were repeated.

> ***"These are not real girls. This is a cartoon, a drawing. Those who look at AI haven't raped living child. It's not necessary when there's AI. It protects your children. It keeps pedofiles safe."***

4.2 COMMUNITY, PRAISE AND ENTHUSIASM

A sense of community was fostered on the forums in many ways. New users often introduced themselves, whilst those posting material described the creation process and the themes involved. Lively discussions arose in the threads about the material, its creation, and various preferences regarding it.

> ***"I am an AI child porn image creator and I have thousands of images I want to share. My interests are from 0-18 male or female (I am into adults too but that is for a different forum) Hardcore, softcore, BDSM, hurtcore, I don't judge. My interests are in all those areas and I create AI art for all tastes. I am posting a selection of what I have and can create on demand (though with close to 110,000 images i done it all (except GORE) let me know if you guys are interested in this and/or if the quality meets your standards."***

Skilled creators of AI-generated material were repeatedly asked to produce more material, to continue existing image series or to create videos. Creators were given plenty of new ideas for themes on which to base their work, and were provided with detailed requests to adapt the material according to the commenters' own fantasies. Requests for guidance on creating material were frequent. The creators were asked many questions, such as how and with what tools they had produced the material, or how they had managed to make a particular detail look so realistic.

> ***"Could you please share instructions how to create such a realistic images? What tools, models, etc do you use?"***

The creators of the content were encouraged and lavishly praised. Those publishing their material for the first time were also thanked and praised. There could be several complimentary and encouraging comments beneath a post. Even critical analyses of the material were often written in a neutral or encouraging tone. Occasionally, attempts were made to quickly silence critical feedback, and those who had offered criticism were scolded for being ungrateful.

> ***"Whenever you started you will get better, keep at it and share with us"***

Feedback on self-created material was also requested and appreciated. Material might be published as though it were an experiment, with others asked to comment on it so that the creator could learn more. Feedback was said to boost the creator's motivation to produce more material.

> ***"I have been busy merging models in an attempt to get [model x] that does photorealistic softcore and hardcore with children. I think I have gotten pretty close, but it's not perfect. I would like to test this model more so if there is interest, I'll take a few requests in this thread. Just describe what you want (age, setting, scene)."***

Creators rarely kept their skills in content creation a secret but rather, sharing them within the community was commonplace and a regular occurrence. Content creators were happy to explain which AI technologies they had used and how they had arrived at their final result. Some brought up in their comments that it is necessary to open up the creative processes to the community. There was a wealth of discussion about content creation, both spontaneously and in response to questions, and progress – such as mastering a new technique – was frequently shared on the forum.

> ***"Personally, I think creators of AI images should, as far as possible, share their methodology. Helping others discover different ways of creating on topic images can only be a good thing."***

The material was analysed in great detail. The details, any potential problems involved in their creation, and the final results were examined at length. A written background story was often added alongside the posted images, videos and/or image series. Stories containing the history and details of the situation depicted served to reinforce the impressions associated with the created material: the poster might write a rather lengthy story, complete with details, to introduce their creations.

Praise and enthusiasm were closely linked to the creation and maintenance of a sense of community on the forums. Referring to the quality of the material and also to the wish for more of it helped to sustain the community. People wanted to follow the community's rules so that the community would remain amicable and, consequently, more likely to survive in the future.

> ***"Keep learning and producing, my friend, and we will continue to consume the content you share with us!"***
>
> ***"Amazing work, you are the reason i keep logging on here, thanks for your hard work (...)"***

> ***"I think AI rape content shouldn't be a problem here and it's a topic of active discussion at the moment with the mods and admin. I am just trying to avoid being made an example of if they decide to not allow it, but I also don't want to have the community pissed off at me for posting that kind of stuff."***

Attempts were often made to solve problems together. Sometimes the issue might be, for example, that there wasn't enough varied, high-quality material available on the open web of a famous child, making it harder to achieve a high-quality result using AI. Participants in the discussion might spend quite a long time trying to resolve the issue and suggest various solutions.

> ***"I'm not sure how to go about this. [user x], any ideas? Interesting challenge. I will look into it! I have a serious need to see more of her."***

The sense of community was also evident in discussions concerning, for example, legal perspectives and the illegality and punishability of the material. In these discussions, an external party seemed to pose a threat to the community, giving rise to 'we' focused speech that served to strengthen the community.

> ***"The more LEA is focusing on fictional abuse, the less is focusing on real sexual interactions between kids and adults, so the more them slutty girls can freely express themselves with us. Meanwhile, what we pedos need to do is keep focusing on the real girls, the ones who matter, the ones made of flesh and bone and feelings and thoughts. Keep promoting new tiktok trends and fashion trends that normalize sluttier behaviours and that make them aspire to life off her body from a young age. Keep working on creating and mantaining livestreaming apps like the late Unico that connect us with them girls in safe spaces and that gave the community so much god-tier content. Keep working on the sexual and flirty aspect of the girls around us. Don't loose focus comrades."***

The sense of community was also strengthened through positive 'we' language, introductions and expressions of gratitude for the community's existence. Comments sometimes included accounts of personal journeys towards AI-generated material and individual ethical views on its use. Participants also spoke of their personal reasons for being on the forum: some were there to share material, some to look for it. Many hoped to strengthen their own AI-related skills by joining the forum, and the community often served as a source of inspiration in this regard. The importance of the community's existence was frequently emphasised. These comments repeated the societal stigma surrounding having a sexual interest in children: participants often felt they had to hide a part of themselves, or that they were not socially acceptable because of their preferences.

> ***"Greetings everyone! I'm glad i found this board, this is what I'm looking for. I started creating ai generated images one year ago, the technology and possibilities are just so exciting! Thank you founder and operators for this clean and fast(!!) forum!"***

> ***"I love these images. Such adorable little girls and I love how clean and professional AI art is compared to a lot of the ametuer IRL images out there. It's more erotic art than porn in a way and I find it very beautiful. I feel like this could very well be the future of CP because of some of those things I mentioned. Also no secrets to be kept and no children to be hurt in the process. I can't help the feelings that I have but I can promote a better way of doing things where no children get hurt. I'm glad we have this community here to explore and share this art form. I guess I went on a tangent there. Thanks for the beautiful images!"***

4.3 STYLES AND PREFERENCES

Forum users' preferences regarding AI-generated CSAM varied. Some wanted material that was as realistic as possible, whilst others wanted the exact opposite. There were repeated comments in the data that AI can be used to create virtually anything.

> ***"However, the thing about AI created images is that you can keep creating and creating all you like, the sky's the limit. If you, like many others, don't have access to an actual child, it can be a good alternative to fulfilling your fantasies."***

AI-generated material was often based on a subject chosen by the creator, which could be either a real child or a fictional character. It was repeatedly mentioned on the forums that users train the AI using images of real children, and that real CSAM featuring real children is used as training and source material for AI-generated material. AI was used to create new material featuring entirely fictional characters, such as Disney princesses, (female) superheroes or comic book characters.

> ***"I love these. Love the cartoony style. I've always wanted to see these characters nude so this is just perfect."***
>
> ***"Some images are based on real girls, mostly for getting the poses I wanted."***
>
> ***"Plenty of real cp already exists. Which could be used as source material for hentai/AI artwork.."***

In the majority of comments, participants wanted the material to be as realistic as possible. On the other hand, whilst hoping for a look that was as photorealistic as possible, participants also wanted it to include unrealistic features. In the discussions, there was a desire for hentai-style or similar material, created using AI and based on real CSAM. The fact that others might not necessarily know what the material is based on was seen as an exciting aspect.

> ***"Realism in AI is important to me."***
>
> ***"If somebody who doesn't know it's AI see it they would believe it is real."***

Forum participants envisaged that, in the future, children appearing in AI-generated material would become objects of 'fandom', much like, for example, adult entertainment stars. For this reason, it was hoped that children appearing repeatedly in AI-generated material would be given names and bios (fictional introductions), so that fans could more easily discuss them and request specific types of material featuring them.

> ***"Wouldn't it be nice to have an AI section for AI porn stars. One topic for a porn star with bio, set names and her/his works? For fans this would be awesome. Right now we see great content but the girls don't have names, it is all mixed up. There are no bios of girls and it is not possible/very difficult by girl. In the real world everybody has favorites, in both adult and child porn. Wouldn't it be great to have characters that fans can follow and even meet in virtual reality?"***

In addition to a name, the children appearing in the material might be given, amongst other things, an age, height and weight, as well as other details concerning their bodies, personalities or sexuality. Some felt that the richness of AI-generated material lay precisely in the fact that it could be as realistic as possible, mimic real-life situations, be cartoon-like, 'plastic' or contain fantasy elements, or anything in between. Many shared that they enjoy a wide variety of styles.

> ***"Just because we can does not mean that everything needs to be photo-realistic now."***
>
> ***"holy crap, these are amazing. I typically prefer photorealism but this set is so dynamic and gripping."***

The details of the material were scrutinised and analysed meticulously. Both the accuracy of AI and, conversely, its errors were highlighted. When it comes to errors, there was strong confidence in AI's development and the creator's growing skills. The characteristics of the children featured in the material – such as skin colour and texture, facial expressions and, in particular, genitalia – were analysed extensively in the comments. Requests to create specific types of material highlighted preferences regarding the kind of activity the material should contain, as well as preferences regarding the children's age, appearance and ethnicity. The requests often mentioned activities that would be explicitly sexual and include elements typically seen in adult entertainment, such as group sex, intercourse or oral sex, or, conversely, children in little or no clothing engaged in seemingly innocent activities, such as having an evening snack. Many requested that the published material could be updated to feature children of a specific age or with specific characteristics.

> ***"Realistic as reasonably possible. Girls about 5 to 8 years old, no babies, caucasian with white or tanned skin, blond or brown hair, nude or in nightgown."***
>
> ***"Also, what's really important that's missed by many A.I. artists is they don't include the expressiveness of the girls. You have them gasping and reacting, which makes the scene sexual. I also like how you have girls from different ethnicities and colours, which is great as sexuality is universal. I'm looking forward to seeing more."***

The activities depicted in the material were discussed from an ethical perspective, but also in terms of personal preferences. The material was analysed meticulously in terms of the activities depicted, and participants described in great detail what had aroused them the most. In this context, the children's facial expressions, in particular, were analysed from many angles, and differences in preferences on this matter came to light: some want to see only happy children enjoying themselves, whilst others hope to see suffering, fearful faces or expressions of disgust or shame. Many were excited by the potential of AI to create bizarre, unrealistic or graphic material without it actually appearing to be violent, e.g. children enjoying BDSM sex. Some comments pointed out that it was not particularly difficult to produce graphic material and that it was even easier than, for example, creating images or videos depicting sexual intercourse.

The vast majority of the discussions centred on girls. It inevitably emerged from the data that the vast majority of children appearing in the material were girls, and a preference for boys was even marginal. There were discussions in which one's own preference for boys was regarded as somewhat embarrassing or as something that makes oneself inferior to others. If boys did appear in the material, they were often requested or described as active participants, because girls were the central focus as objects. The arousal was therefore derived from what was done to the girls, rather than what the boys did.

Similarly, the names of certain girls recurred in the material. They were referred to in such a way that it was assumed to be clear to everyone who was being discussed. Participants hoped for a specific type of material featuring them, or wanted a particular girl to be added to material that already existed. However, without knowing the background, it was difficult to interpret whether this referred to a real child of whom CSAM existed – material that had been further generated using AI – or a fictional child of whom AI-generated CSAM had been produced repeatedly.

> ***"OMG I've been waiting to see her like that for so long!!!!! I;m in love xxx"***
>
> ***"Just wow. I have always wanted to see these girls back to life and now the technology is within reach. I need to do some serious studying!"***

4.4 ETHICS AND JUSTIFICATION

There was a great deal of discussion on the forums about the various ethical and justification related aspects of using CSAM. The use of AI-generated material was often seen as a more ethical alternative to consuming real material. Some forum participants said that AI-generated material was the only option for them. On the other hand, some said that AI-generated material left them cold and did not elicit the desired reaction. Rather, they shared, the material must be filmed in a real-life situation and must show sexual abuse of a real child.

There was much discussion about the ethics of material created solely by AI, without it always being compared to real material. Some felt that AI-generated material must not depict a child's suffering: consequently, only so-called 'romantic' material was, in their view, ethically sound. Some said they had joined the forum specifically because they love children and believe that no harm should be done to them.

> ***"while real hurtcore pics appear just sick to me, this is definitely free-spirited, creative art."***

The material that some people wanted should not show any signs of abuse, and instead the children were to look happy and as if they were enjoying themselves. Such material was also relatively realistic: amongst other things, greater account was taken of a child's age-appropriate physical development, and tools such as the Tanner scale – which describes a child's physical development and the gradual changes of puberty – were used as a guide, whilst any potential unrealistic elements were readily pointed out. This was also justified by the argument that so-called 'hurtcore' – a term referring to material containing violent elements – is precisely the reason why all paedophiles are viewed as evil. Some pointed out that as long as such graphic material is produced and consumed, no one with a sexual interest in children will receive any form of acceptance.

> ***""Hurtcore" is the reason people want to not only completely remove pedophiles from society, but also execute them."***
>
> ***"For me, such content is nothing but garbage and only harms the cause of those who truly love children."***
>
> ***"this may sound weird but if put them in sexual situations i feel like i am raping them somehow even though they do not exist. i enjoy the stuff leading up to sex more than the sex. i am a hopeless romantic .."***
>
> ***"Well this just goes to show that people are right by calling us monsters..AI or not, this is just wrong for a girl (LOVERS) site!"***
>
> ***"I understand it's AI but perhaps we should have a little taste. I've noticed a divide between the child lovers and the hurtcore crowd. As a lover I am disgusted by the thought of taking pleasure from the pain and suffering of children. I am here in the AI part of the internet because I don't want to hurt anyone. Everybody I know in the real world thinks pedophiles should die reguardless of weather they act on it or not. People like to fantisize about "when" society accepts pedos but that's never gonna happen and when people see us drooling over rape and murder and torture of children they are proven right as to why we can't be allowed to exist."***

On the forums, particularly on the forum dedicated solely to AI-generated material, there was discussion about whether 'hurtcore' material should be permitted there at all, and if so, where it should be placed so that those who wished to avoid seeing it would not be exposed to it. There was very detailed and in-depth discussion about the ethics of the material.

> ***"Honestly, the thought of [forum]] getting a hurtcore section gives me the creeps. Personally, I have no interest in enjoying violent sexual imagery. However, I see the point of keeping it isolated and locked away if it's not removed completely."***
>
> ***"So it would be nice for me if images like that were confined to subforms of this site so I could easily skip them."***

Some found the idea of any graphic material being present on the forum unbearable. Some were prepared to allow the graphic material, provided it was separated from other content into its own channel, whilst others suggested this in a neutral tone and identified benefits for the community: explicit material generated by AI was seen as a lesser evil than forums where real explicit material is shared. Allowing graphic material produced by AI in its own section was thus seen as a way of staying away from forums containing real material, which might then lead to the forums no longer existing in the future as they would have no user base. Similarly, it was thought that allowing graphic material produced by AI might potentially prevent people from harming children.

> ***"I would like you to consider that allowing AI rape images here could have potentially positive effects as well. Some people are always going to seek out hurtcore content. If they can consume AI content and that content is plentiful and varied, they're less likely to keep searching for the real thing. It could potentially lead to real hurtcore sharing communities seeing less activity and possibly shutting down. That means less people sharing images and videos of real children being hurt. I also have a theory that having access to AI rape images could potentially keep someone from deciding to rape a child since they can't "scratch their itch" any other way."***

The perspective of the authorities and the law was frequently viewed on the forums as follows: if those who produce or possess AI-generated material are held accountable, real children are put at risk. According to the logic, the authorities are "driving" offenders with a sexual interest in children to prey on real children. Furthermore, focusing on AI-generated material diverts resources away from investigating crimes against real children. Criminalising and punishing AI-generated material was seen as an absurd attempt to "save fake children" and was often met with derision. On the other hand, it was believed that the authorities were interested in AI-generated material and its criminalisation because it included graphic and violent material. There were occasionally lengthy discussions on the forums regarding the legal aspects and media coverage of the issue.

> ***"This crackdown is a bit of a concern as it means that AI CP is no longer a safe alternative."***
>
> ***"So now they are trying to save fake children, that's interesting."***
>
> ***"so they're now trying to push the boundary for pedos to go and try to get real kids (...) if you push a cyber pedo into a corner, he will start to hunt in the real world"***
>
> ***"People who like such AI images are a reflection of their own real life values. Extreme sick content is a great excuse for LEA to be far more motivated to shut down this great site."***

Among other things, it was suggested on the forum that 'hurtcore' should only be created in the hentai style or a similarly unrealistic style, so that it would be as clear as possible to everyone that the material was not real. On the forums, real and AI-generated material were repeatedly confused, and users were not always fully aware of which was which. On the other hand, there was also a deliberate desire to create material where it is unclear whether it is real or not.

> ***"It might be worth considering limiting hurtcore content to only hentai and semi-real/3d only, to ensure there is no question that the content being shown is not real."***
>
> ***"When you look at child porn rape you don't have any strong evidences that it's fake. With real content no one will confirm that it was fake, at least convincing. And AI arts comes really close to real content in quality."***

Some of the forum participants believed that, precisely because the material was created using AI, it could depict anything, as no real child was suffering– in which case even 'hurtcore' would be entirely permissible. Overall, a major part of the justification was that AI-generated materialwas not seen as harmful to children, unlike real material. Graphic material was justified on grounds such as democracy, freedom of expression, the fact that even ordinary films, games and adult entertainment contain various forms of violence, and the argument that AI-generated 'hurtcore' helps prevent real children from becoming victims. The graphic material was also defended on the grounds that it keeps users content and away from more brutal and potentially real material.

> ***"No real child is being hurt, sure we share a certain idea however there's also these violent fictional movies, stories and games out there and they're all fictional as well."***
>
> ***"The fact it's (hurtcore) a turnoff for some users doesn't mean it should be banned as we are completely in the art territory."***

Furthermore, AI-generated material was viewed as art. Users wanted it to be a channel and form of free self-expression available for anyone who wished to use it. Sexuality and a sexual interest in children, as an aspect of it, were seen as part of humanity and self-realisation. Comments highlighted that havinga sexual interest in children is not a choice, and that if the associated needs could be satisfied with AI-generated material, this would spare real children. Even when it was known that AI-generated material was based on real material, it was viewed as an artificial product which was not seen to harm the children behind it. Rather, its sole impact was considered to be on the viewing audience. It was also argued that AI does not store the images used as the basis for the generated material, such as an image of a real child.

> ***"Even if the content is generated based on an actual person, it is still only a fake representation and the only people being effected by it are the pedos here. Satisfying a basic human drive without any actual people involved."***
>
> ***"I did not choose to be a pedo, it is simply a fact my life. Sex is one of the most basic driving forces of humans. This force is sufficiently satisfied here by contributions of 100% fake AI generated art. (...) Taking action to prevent the sharing of this 100% fake material increases the chances of REAL children being harmed. I am satisfied by this material to the point of paying less attention to REAL children."***

> ***"When models are trained on dataset of girls for example, once they are done the model does not store any pics of any real girls or anything for that matter. I know this becaue people have accused me of using 'the base photos'. It then does it's best to approximate what you asked for in the text prompt and created an entirely new unique rendering of it. So there will never be a case where you create a girl and one day you'll run into said girl because she does not exist."***

Some participants in the discussion pointed out that banning graphic AI-generated material on forums would be hypocritical, as all the material there is harmful and illegal. There was discussion on the forums about whether users would give up real material if AI-generated material were to be legalised. The responses were mixed. Many, however, commented that they were already moving in the direction of using almost exclusively AI-generated material.

> ***"if I have the choice, then yes, i would choose a.i. over real cp. It has more potential in the long run."***

On the forums, AI was sometimes seen as responsible for the material instead of a human: after all, it is the AI that ultimately creates the material, not the human. The human merely issues a command to the AI, which the AI then carries out. AI-generated material was also compared to drawn material, in which the human contribution was seen as much greater and more significant than in AI-generated material.

> ***"The point I am making is that Lolicon is willfully drawn to look a certain way, by the artist, the human producing the art. This is the complete oposite of Ai art. The artist is not the person producing the CP, it is the Ai. All you did was type a bunch of words. It's like a glorified notepad. Is it illegal to write stuff down? No (at least not as far as I could find, drawn yes, but text no) So what would you be getting arrested for? The ai should be the one on trial. This is the reason why they say that the Ai holders are not responsible for what people do. Ai is random noise that gets cleaned up, not by you, but by the ai itself. You have no control over what it does, it does it itself."***

The use of the material was justified by a variety of views, such as the claim that all men feel sexual desire towards minors, and that all women have rape fantasies. Beliefs also emerged on the forums suggesting that viewing CSAM is quite common – it is simply done in secret from others.

> ***"Not to be crude, but I firmly believe that CP is like eating boogers - EVERYONE does it, but no one will openly admit it."***
>
> ***"Any straight male who says he's never been attracted to a girl "under age" it's a liar."***
>
> ***"Most women have fantasies of being raped. It doesn't mean they are actually going to go out and try to make it happen (most of them). Women want to be taken, it is inconvenient for those who want life to be neat and clean and proper, but for those who actually have women in bed and pay attention to their needs, the reality of what turns them on doesn't fit with what is "supposed to" be "correct"."***

Furthermore, viewing AI-generated images – such as those created using nudifying apps – becoming aroused by them and using them as masturbation material was considered quite common and, according to the participants, does not harm anyone. On the contrary, it was seen as part of (a boy's) normal development. Discussions took place on the forums regarding whether the use of deepfake images should be punishable. The penalties were not seen as being based on valid grounds, and this gave rise to a debate in which there was a desire to punish the victim of the violence rather than the perpetrator.

> ***"They wanted to ruin his life because of him jerking off, in private, affecting nobody, with fake AI generated pictures. This speaks to the mentality of these people. They want to cause real life harm to a boy, for virtual mental crimes, and for his attraction to young girls, wich is perfectly normal. If it were me in his place, i would mess with this girl by telling her, that even thought the AI pictures were destroyed, i still jerk off thinking of her. I bet the idea alone would be incredibly anoying to her."***

4.5 THE FUTURE OF AI

The comments repeatedly reflected a belief in a bright future for AI: it is believed that AI is constantly improving, and whatever is not yet possible will certainly be so very soon. Dreams of virtual child companions that are as lifelike as possible also recurred in the comments. The improvement of AI was seen as inevitable, and keeping up with it was at times difficult even now.

> ***"Thank you [user x] and the other A.I. cp video pioneers. Some people can't imagine what is coming. Not only will we get new girls, there will also be new videos with classic cp superstars and cp videos with Hollywood girls. Bad lighting, bad camera work and fat ugly male participants will become a thing of the past. No limits baby."***

Comments emphasising the potential of AI recurred. AI seemed to bring a sense of control over the consumption of material: anyone can create it with whatever characteristics they wish, and there is no longer any need to wait for new material whilst uncertain about its quality or release time. These comments also echoed the idea that AI-generated material protects children, too, and that children would, therefore, no longer be subjected to violence in the future.

> ***"This is more than just a new tool. It represents the dawn of a whole era for pedos. No longer will we be bound by real life and wait for some one to post and be concerned about what happened to the child. We can create our new reality where we decide what we want to see and know that no actual child was involved."***
>
> ***"Realize this a new frontier a new era. We are getting closer and close to being able to produce our own kiddie porn with the ages and settings and action we want."***

Some commenters hoped that AI and its development would help spread CSAM more widely across the open web and to a larger group of people. It was not necessarily believed that legislation would keep pace with the development of AI in the future, either.

> ***"Whenever one site is taken down, another appears. It's been going on for years. The authorities cannot win. The odds are in our favor. And thanks to AI, it's getting easier and easier to convert people, and the prolific nature of kids with smartphones and constant exposure to ever-loosening morals up on the clear web means that things are moving in our favor. The law is always playing catch-up with technology."***

AI-generated material posted on forums was often analysed in great detail, and these discussions frequently included reflections on how the development of AI would resolve the problems currently being faced. Overall, AI and its development were regarded as a significant milestone, and it was compared, for example, to the significant role photography once played. These comments also emphasised that it can be difficult to distinguish AI-generated material from genuine material. On the other hand, this was also something people hoped for: AI was seen as developing to such a high standard that it would, in the future, always be impossible to distinguish material produced by it from photographs.

> ***"Well its only going to keep improving, there is no forcing the cat back in the bag. Someone posted some short vid clips, these are just a taste of what is coming. And don't forget VR with tactile suits."***
>
> ***"This is really is extraordinary. Some are very extremely close to perfect. If you were to show this to someone who didn't know, I think they would not able to tell they are fake. And with the new processors and new software coming out, there will be no difference at all. It's hard to believe the capabilities they have now. You do really great work."***

4.6 THE USE OF AI TO DEPICT FORMER VICTIMS OR WELL-KNOWN INDIVIDUALS

There was a great deal of discussion on the forums about children who had already been subjected to sexual abuse, and of whom CSAM had been shared online. Some of these children/victims were referred to on the forums as so-called 'CSAM superstars', and CSAM of them was widely used and shared. The discussions highlighted how AI can be used to produce new CSAM featuring these individuals.

The quote below is a typical example of a discussion amongst several forum participants concerning a well-known victim of sexual abuse and the production of new material featuring her.

> ***"I'd love to see some AI of my favorite Jane Doe A undressed also Jane Doe B"***
>
> ***"User X posted some deepfakes including some with Jane Doe A. See the Deepfake Videos post in the Other forms of AI or mixed forum."***
>
> ***"User Y has a pretty good Jane Doe A loras, they are floating around and still available on TimelessAI."***

Similar examples were found for several individuals who were known to have been victims of sexual abuse as minors.

> ***"The last of this series, Jane Doe. link"***
>
> ***"can you please please upload the Jane Doe-Lora ?!?! its so amazing <3!!"***
>
> ***"Fantastic lora! Could you make one for Jane Doe for hunyuan?"***
>
> ***"Could you post the LoRA for Jane Doe?"***
>
> ***"I trained a LoRA on Jane Doe recently and used it to generate this video of her dancing"***
>
> ***"This was trained on images which contained the upper body of Jane Doe, several images went below the bellybutton to capture its position. The age of Jane Doe ranged from about 6 to 10, there is some age diversity in this. It was trained on all kinds of facial expressions, some less attractive, some funny. Hopefully to capture the essence of Jane Doe."***

In some of the discussions, an interest emerged in creating material featuring well-known individuals, such as children appearing in films or TV series, the children of public figures, or children who are otherwise in the public eye. To some extent, there was also interest in creating material featuring a well-known person who has now grown into an adult, depicting them as a child.

> ***"Not sure if anyone went to go see the [X] movie but I couldn't keep my eyes of Joe Doe. I'd love to deepfake some porn of him, anyone know how I can start?"***
>
> ***"Hello everyone. That would also be a question from me, would it be possible to take nude AI pictures of Jane Doe. When she was still a minor. Or are there real leaks of her online?"***
>
> ***"On the advice of fellow member stuffins, I will be splitting each of my child actress posts into the LoRA and the image set. You can find my experiment images here. link"***
>
> ***"Sometimes I get my ideas for new LORAS from random girls I find in swimsuit and panty ads online. These girls deserve to get showcased too, not just the normal lolita stars: link"***

4.7 THE QUALITY OF AI-GENERATED CSAM AND RELATED DISCUSSIONS

There was a great deal of discussion on the forums regarding the quality of AI-generated CSAM and how to improve it. In relation to quality, the discussions frequently touched upon, amongst other things, the use of different AI models and their settings.

> ***"The last one even have the rounded corners that Omegle had. The overlays can be used with the image overlay node from Efficiency nodes for ComfyUI."***
>
> ***"Until the next generation models comes out, such as the cascaded of conceptual ones, I fear the diffusion based outputs would keep this cartoonish touch."***

> ***"I would say the absolute top level of realism would be to add pictures to existing sets that never existed. Like using AI to extend a set to include more pictures. i am sad I don't have the skill-set to attempt such thing."***

In discussions regarding the quality of AI-generated CSAM, the focus was often on the functionality of various lora (low ranking adaptation) models, as well as their settings and usage to improve quality.

> ***"Here's a few pics using your lora. My feedback: 1. seems to do quite well with the vaginas, and the hip proportions, but seems to have an issue that can cause the body to become long sometimes. It doesn't seem very friendly with other loras. 2. I think the lora is best used at a strength of 0.7 or less, because it has too strong an effect at 1.0. 3. My Custom SDXL/Pony model is heavily biased toward 4yo girls and under, and this lora helped me to get images of girls looking older in the 5-12 range. Could use some fine-tuning, but its definitely on the right track and works pretty well."***
>
> ***"Impressive, what beautiful videos. Seeing that AI/software can generate a video with a single image, I think the LoRas "character" thing will soon become meaningless, especially given the pace of advancements and improvements."***

Some of the participants in the discussion were particularly enthusiastic about how the development of AI models will enable much higher-quality AI-generated CSAM in the future.

> ***"Soon anything we can imagine can be created"***
>
> ***"We have a hot future in our hobby"***
>
> ***"If the result with images right now is brutal, I can't even imagine what the video generation will be like in a few years. Crazy, the future is here."***

4.8 DISCUSSIONS RELATED TO AI TECHNOLOGIES AND THEIR USE

There was a significant amount of technical discussion on the forums regarding AI technologies. There was a wide range in the discussions about technology, often depending on the questioner's own level of expertise regarding the use of AI. Some of the questions and answers related to basic-level use, whilst others delved into very specific technical details.

Generally speaking, participants who were still relatively inexperienced in using AI often asked very broad, non-specific questions, and might ask about several different things at once.

> ***"Hi everyone. Recently I've seen some really sweet and lovely AI-generated pics and I'd love to know how to create pics like that myself. For a start - what software & hardware is needed? I guess that most of the pics aren't created on Chat GPT. Or am I wrong and they are all created on publicly available sites? There are loads of pics/scenes I want to create. Please tell me what to do."***

> ***"What negatives do you use? What model? What checkpoint? (I haven't tried mixing checkpoints yet) What lora? (I am still trying to figure these out)"***

Some of the participants were much more familiar with AI technologies, and their expertise was considerably more advanced. In the discussions, these individuals often focused on more detailed aspects of AI technology. They were keen to improve the process of creating material or the quality or quantity of AI-generated CSAM.

> ***"The technology is exploding with new models coming online regularly. Some of the best generation models are model X and model Y but both are only paid online. This makes them dangerous and often impossible for ontopic images due to excessive censorship of ideas."***
>
> ***"Could you share the dataset for this lora please? I would like to have a go at training Model x once I get my new GPU."***
>
> ***"To those who are training Flux lora can you tell about your process? I imagine different process depending on goals."***

There was also discussion about which AI model software options are the best and what specific characteristics might be associated with different AI models.

> ***"I know about model x, but what other good (realistic) models are there? is there any better than model y?"***
>
> ***"model x imo is best for most realistic i have it working well with custom loras for nude and hard realatorg"***
>
> ***"I like Chilloutmix because it is suitable to generate east Asian people that I belong to."***

Many participants were happy to share freely available software and models, either developed by themselves or by other participants, and asked to be given specific models for their own use.

> ***"I can offer these LORA models to reduce the age for model x and model y. link"***
>
> ***"I'm desperately looking for a model for model x. Does anyone know where I can find?"***

A wide variety of prompt lists for using different models and generating material were shared in the discussions. The lists contained both positive and negative prompts. Prompts are, in general, a significant part of the use of many image generators. For this reason, all examples relating to them have been deliberately omitted from this report, so that they are not utilised in the creation of AI-generated CSAM.

5 Discussion

Technological development today is rapid and has transformed our everyday lives in many ways. AI is considered innovative and useful, and much is expected of it. At the same time, however, it has enabled the development of new forms of child sexual abuse. AI allows AI-generated CSAM to be produced easily, quickly and on a massive scale. The AI technologies used to create CSAM are accessible to anyone and are easy to use. Safeguards are often inadequate and easy to circumvent. In Finland, the Criminal Code defines an image depicting a child in a sexual manner as an image or a visual recording that depicts a child in a sexual manner, either factually or realistically. The possession and distribution of realistic AI-generated CSAM may, therefore, be illegal.

The findings of this report show that AI-generated CSAM has its user base, which is dedicated to producing, searching for, further developing, and consuming the material. There was a strong sense of community on the CSA forums on the dark web. The data showed that comments were thought through and were often long and detailed. Other forum users were praised, encouraged, spoken to in a friendly manner, and warmly welcomed to the forum. Feedback on AI-generated CSAM was provided readily and unsolicited, and other users were consulted on both ethical issues and matters relating to the creation of the material.

To some extent, the sense of community fostered and perpetuated an environment in which child sexual abuse was repeatedly justified using various arguments, and in which there was a great deal of enthusiasm for the subject. Different preferences regarding the quality and detail of the material were expressed without restraint, and in turn these often served as sources of shared enthusiasm. It was repeatedly felt that, when it came to AI-generated material, the only limit was one's imagination. The community created a solid foundation for the proliferation of the material and the normalisation of child sexual abuse.

Throughout the data, there was a strong belief in the potential of AI and in the fact that it is constantly improving. Some forum participants also firmly believed that AI-generated CSAM is the answer to ethical concerns and will ensure that real children are no longer subjected to sexual abuse. A recurring argument on the forums was that AI-generated material prevents sexual offences against real children. However, AI-generated CSAM is repeatedly produced from images and videos featuring real children – both everyday ones and those depicting sexual abuse. It always sexualises the child and does not protect children from sexual abuse.

It is also possible that AI-generated CSAM may perpetuate or even reinforce the preferences of a person with a sexual interest in children and lower the threshold for seeking out CSAM: a recurring theme in the data was that participants were constantly and diligently searching for material to train the AI and produce more CSAM, and that they were prepared to abandon genuine material even entirely, as AI can, in practice, produce virtually any kind of material. Even material produced by AI based on real material was not viewed as harmful to children in the forum discussions.

People who wish to consume AI-generated CSAM or create it themselves are constantly looking for ways to circumvent the safety measures built into AI models. The data revealed that legislation relating to AI and its shortcomings were well recognised and, partly for this reason, AI-generated CSAM was perceived as safer than real material. At the same time, many forum participants felt that AI-generated material was the best they had ever seen – better than real material – and wanted to see more of it and learn how to create it themselves. Advice and instructions on both creating the material and circumventing security measures were shared repeatedly.

One in five children encounters sexual abuse online. Among girls, the figure is as high as one in three. According to research, the majority of deepfake material features girls and women. In the data, the proportion of girls as victims of sexual abuse was pronounced. The vast majority of discussions concerned girls – boys might appear in the material as active participants, but rarely as targets of sexual abuse, compared to girls. In any case, AI currently allows for the mass production of CSAM of any child whose images or videos are available online. All digital environments should be safe for all children. AI-generated CSAM, the ease with which it can be created and found, and its vast scale normalise child sexual abuse, and its volume has been steadily increasing. It cannot be regarded as a milder version than real material that helps keep real children safe – quite the opposite.

Technological development must not take place in isolation from ethical values or responsibility. It must not mean an explosive increase in child sexual abuse or an expansion of the range of methods available to perpetrate it, nor is it the role of AI to lower the threshold for violence. In AI-generated material, children who have experienced sexual abuse are victimised again and again as images and videos are created and further processed from existing real material. Some of the children featured in AI-generated material have already grown into adults, yet the sexual abuse they endured continues online. We need more robust measures to tackle this issue and louder discussion on the ethics of technology and safety of digital environments, as well as on how society can keep pace with technological developments and protect children from sexual abuse – and what happens when legislation, safety measures and monitoring prove, time and again, to be inadequate.

And what happens if machines learn too much, and what might such a development mean for the future? Until recently, large language models were regarded in public debate as the key manifestation of AI development, and attention was mainly focused on their functioning, potential and the challenges associated with them. However, the development of AI is no longer limited to language models, but is advancing simultaneously across several different areas. For example, the rapid development of large image models and their potential integration with other advancing technologies, such as extended reality, raises the question of what kinds of new and even more versatile AI-based solutions will emerge in the future. At the same time, it is worth considering whether such technologies might also increase the risk of serious abuses, such as facilitating or enabling child sexual abuse in digital environments. This raises the key question of how technological development can be guided and regulated so that such threats do not become part of our future.

6 Needs for further research

Several themes emerged from the findings of the report, on which further research is needed in order to develop the prevention of child sexual abuse to be more effective. From the perspective of protecting children, it is essential to gain a better understanding of whether the consumption and production of AI-generated CSAM reinforce cognitive and social mechanisms that perpetuate the sexualisation of children or lower the moral barriers to child sexual abuse. It would be particularly important to investigate how CSAM enabled by generative AI – which is, in practice, limitless and customisable to individual users – influences users' perceptions of the acceptability of such material, as well as thought patterns that justify sexual interest in children.

Furthermore, it would be worthwhile to investigate how the perceived 'victimlessness' of the material and the associated justifying discourses may reinforce the objectification of children or perpetuate ways of thinking that normalise the sexualisation of children. Another key area for further research is the potential link between AI-generated CSAM and contact offences. From the perspective of protecting children, it is important to determine whether the use of such material may reduce or increase the risk of sexual offences against children.

Future research should also examine minors as users of AI-generated CSAM. It would be particularly relevant to investigate how early exposure to such material affects an individual's attitudes, beliefs and sexual interests in the long term. At the same time, it would be important to examine whether use during adolescence is linked to thought patterns that perpetuate the sexualisation of children or to a later risk of child sexual abuse.

In addition to these themes, more research is needed on effective, long-acting intervention methods. It would be particularly important to determine which types of interventions are most effective in situations where an individual regularly consumes large quantities of AI-generated CSAM. In the future, it would also be worthwhile investigating how interventions can influence the thought patterns that sustain the use of the material, beliefs that downplay the harm caused, and cognitive processes that normalise the sexualisation of children.

7 Recommendations

Up-to-date legislation, regulations and guidelines, and their enforcement

Society should act in accordance with children's rights, which requires **adequate legislation, its implementation and monitoring** to ensure, for example, that children's safety and privacy are realised.

Up-to-date and comprehensive legislation is needed to protect children from AI-facilitated sexual abuse. When drafting and updating legislation, the impact on the consistent realisation of children's rights must be assessed. In addition to new legislation, the enforcement of existing legislation is required. For example, the ban on so-called 'nudifying' apps, which came into force as part of the update to the EU's AI Act (the Digital Omnibus on AI), must be urgently enforced.

One key form of source material for AI-generated CSAM is everyday images of children taken from the open web and social media. It is essential that legislation guarantees **children the widest possible rights to control their own digital footprint, including retroactively**. Based on the findings of the report, children's growing digital footprint appears to be one factor increasing the risk of them becoming victims of sexual abuse.

Giving children greater control over their digital footprint

In light of the findings of the report, children's growing digital footprint appears to be one of the factors increasing the risk of them becoming victims of sexual abuse. However, children's ability to control their own digital footprint is limited. It is, therefore, important that parents and guardians take into account the child's right to privacy and safety in their own actions, for example when planning to share something in digital environments. Parents and guardians should carefully consider what sort of images they post of children online. It is not in a child's best interests for their images – particularly intimate ones – to be potentially visible to millions of people online.

Children's growing digital footprint also poses challenges for other actors working with children, such as youth and sports organisations. Children have the right to be visible online, but at the same time we must be aware that the more images of children are shared on the open web in the context of leisure or sporting activities, the more images are available to criminal actors for the creation of AI-generated CSAM. There should be clear guidelines on how children's images are used and published in the activities.

The ‘safety by design’ approach to the heart of AI innovation

‘Safety by design’ refers to the principle whereby safety is built into a service or system under development right from the design stage, rather than safety being addressed only retrospectively when problems arise.

The ‘safety by design’ approach should be at the heart of AI innovation. From the outset, AI innovation should include an assessment of whether a potential innovation is dangerous and whether it is, in fact, of such a nature that it should not be brought to market.

As technology develops, there is a need for better built-in technical safeguards that cannot be circumvented. The safeguards should prevent the creation, via AI, of CSAM or material that sexualises children within applications or platforms.

Technology companies must ensure that no AI-generated or real CSAM is shared on their services. Companies and platforms should also clearly stipulate that the sharing of intimate images of children, images where children are in little clothing, or other images that violate children’s privacy, is not permitted on their services. They should provide clear guidance to their users on this matter and, where necessary, take action if images are shared.

Media and AI literacy education for all

The continuous development and widespread adoption of AI require that citizens possess sufficient AI literacy to be able to interact safely with and critically evaluate material generated by AI. This requires AI literacy education that is available to all children on an equal basis nationwide, which can be implemented, for example, as part of broader media education. Children must have sufficient skills to assess whether material may have been produced by AI and what that implies.

Professionals working with children play a key role in enhancing children’s AI literacy, for example as part of broader media education. Guardians and parents also have a role to play in enhancing children’s AI literacy as part of their normal parenting responsibilities. Professionals, parents and guardians, therefore, need up-to-date information and tools.

High-quality age and developmentally appropriate sexual education for all

AI-facilitated sexual abuse also creates new needs for sexual education. All children should have equal access to evidence-based sexual education appropriate to their age and stage of development, which includes topics relating to AI-facilitated sexual abuse. These topics include, for example, what AI-facilitated sexual abuse means and how to protect one’s own and others’ boundaries in relation to it.

Sources

Internet Watch Foundation. 2023. *How AI is being abused to create child sexual abuse imagery*. Available at: https://www.iwf.org.uk/media/q4zll2ya/iwf-ai-csam-report_public-oct23v1.pdf

Internet Watch Foundation. 2024. *What has changed in the AI CSAM landscape?* Available at: https://www.iwf.org.uk/media/opkpmx5q/iwf-ai-csam-report_update-public-jul24v11.pdf

Kokolaki, E. & Fragopoulou, P. 2025. *Unveiling AI's Threats to Child Protection: Regulatory efforts to Criminalise AI-Generated CSAM and Emerging Children's Rights Violations*. Available at: https://www.researchgate.net/publication/389548465_Unveiling_AI's_Threats_to_Child_Protection_Regulatory_efforts_to_Criminalize_AI-Generated_CSAM_and_Emerging_Children's_Rights_Violations

Media Education Association. *Infographic: Digital Footprint*. Available at: https://mediakasvatus.fi/materiaali/infograafi-digitaalinen-jalanjalki/

Save the Children Finland. *A Better Digital Life for the Whole Family – A Guide for Parents*. Available at: https://www.pelastakaalapset.fi/ajankohtaista/parempaa-digiarkea-koko-perheelle-opas-vanhemmille/

UNICEF. *Sharing content about your child*. Available at: https://www.unicef.org.au/sharing-childrens-images-online?srsltid=AfmBOoqUW_1TCOC7MpbNYaCmDCBOwBR6vByzS1qO-OHwi6e0K71R79g_

Appendices

APPENDIX 1: GLOSSARY

AI (artificial intelligence): artificial intelligence

AI CP (artificial intelligence child porn): a term used on dark web forums by which participants refer to material depicting sexual violence against children. It is important to remember that the word 'porn' should not be used in relation to sexual violence against children.

AI-generated: created using artificial intelligence

AI guide: an artificial intelligence guide

AI prompt: prompt are used to tell AI what user wants them to do

Checkpoint model: a saved snapshot of an AI model's training state – including its learned weights, biases and optimiser configurations – at a specific point in time

CSAM (child sexual abuse material): material depicting sexual violence against children

Dall-e: artificial intelligence software developed by OpenAI

Dark web (also known as the deep web): content on the internet that requires special software, settings or actions to access. The aim of the dark web is to make users largely anonymous.

Deepfake: material that is created by AI and it's purpose is to show person in a situation that has not happened in real life.

Diffusion model: diffusion models are generative models that 'diffuse' training data with random noise and then learn to reverse the diffusion process to generate new images

Fine-tuning: fine-tuning refers to the process whereby a ready-made, pre-trained base model (such as a language or image model) is modified and further trained using new, more specific data.

Flux.1: AI software developed by Black Forest Lab

Hentai: Japanese-style pornographic comics

Hurtcore: material depicting sexual violence against children involving graphic violence

LEA (law enforcement agency): a term referring to the authorities

Lolicon: a term referring to erotic or pornographic anime and manga-style images depicting young or childlike girls, and the entire genre

Lora (low-rank adaptation): a fine-tuning model

Midjourney: AI software developed by Midjourney

Nudify app: an AI app or platform that can be used to turn everyday images into sexually suggestive and/or nude images

NSFW: an abbreviation for 'not safe for work' or 'not suitable for work'. A warning used to flag links, videos, images or text messages containing inappropriate, graphic or offensive material that you would not wish to view in a public, professional or school setting.

Open web: the internet that is open and public to everyone, also known as clearnet.

Pony model: an open-source model designed for generating AI-generated images

Prompt: a prompt is a text, question, or instruction given to an AI or computer system to guide its response or action.

Stable Diffusion: AI software released by Stability AI

When machines learn too much

Findings from dark web discussions regarding the production of AI-generated child sexual abuse material and the technologies used for that purpose

Text: Mikko Ahtila, Maija Ikonen, Heidi Näppi ja Paula Soini
Illustration: Red and Blue
Layout: Days Agency
Save the Children 2026

Save the Children

Save the Children is the world's leading independent children's organization, creating both immediate and lasting changes in children's lives. Together, we secure every child's right to life, protection, education, and participation in Finland and around the world.

Because children's future starts now.

SAVETHECHILDREN.FI

FUNDED BY
Ministry of
Social Affairs and Health

Co-funded by
the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.