
ANALYZING THE AI NUDIFICATION APPLICATION ECOSYSTEM

A PREPRINT

Cassidy Gibson
University of Florida

Daniel Olszewski
University of Florida

Natalie Grace Brigham
University of Washington

Anna Crowder
University of Florida

Kevin R. B. Butler
University of Florida

Patrick Traynor
University of Florida

Elissa M. Redmiles
Georgetown University

Tadayoshi Kohno
University of Washington

November 18, 2024

ABSTRACT

Given a source image of a clothed person (an image subject), AI-based nudification applications can produce nude (undressed) images of that person. Moreover, not only do such applications exist, but there is ample evidence of the use of such applications in the real world and without the consent of an image subject. Still, despite the growing awareness of the existence of such applications and their potential to violate the rights of image subjects and cause downstream harms, there has been no systematic study of the nudification application ecosystem across multiple applications. We conduct such a study here, focusing on 20 popular and easy-to-find nudification websites. We study the positioning of these web applications (e.g., finding that most sites explicitly target the nudification of women, not all people), the features that they advertise (e.g., ranging from undressing-in-place to the rendering of image subjects in sexual positions, as well as differing user-privacy options), and their underlying monetization infrastructure (e.g., credit cards and cryptocurrencies). We believe this work will empower future, data-informed conversations — within the scientific, technical, and policy communities — on how to better protect individuals’ rights and minimize harm in the face of modern (and future) AI-based nudification applications.

Content warning: This paper includes descriptions of web applications that can be used to create synthetic non-consensual explicit AI-created imagery (SNEACI). This paper also includes an artistic rendering of a user interface for such an application.

1 Introduction

Computer vision and generative AI techniques can undress¹ someone depicted in a picture or video [61, 25, 20, 4]. Such technology is increasingly publicly accessible, leading to a proliferation of “nudification” applications available to end-users online.

Nudification applications enable end-users without technical or even Photoshop skills to artificially generate intimate imagery of someone *without* their consent. We refer to such a resulting image as a *synthetic non-consensual explicit AI-created imagery*, or *SNEACI*². Non-consensual imagery created through the use of AI has a starkly different paradigm compared to that of images created through Photoshop or technical skills. AI makes the creation of these images easier, faster, and more realistic because of how advanced generative AI has become. Similar to malware-as-a-service, synthetic non-consensual explicit AI-created imagery brings non-experts the ability to harm at scale. The creation of

¹We define undress functionality as that which takes a representation (image) depicting a clothed individual and produces an representation of that individual without clothes.

²Synthetic non-consensual explicit AI-Created imagery, or SNEACI, refers to images or videos that depict a nude or semi-nude subject, including those that contain intimate body parts and/or depict the subject engaged in a sexual act, without the subject’s consent.

SNEACI is a form of sexual abuse against the subject who, by definition, is non-consensually depicted in the resulting content [32]. In addition to the serious mental health impacts victim-survivors sustain from the violation of being depicted in SNEACI [18], SNEACI may be used by the creator to extort and/or otherwise harass the subject of the image [30]. As such, there is increasing concern amongst policymakers and new legislation about SNEACI and the applications that facilitate its creation [37, 22, 2]. Despite the severity of abuse possible with nudification applications, little is known about this software ecosystem. In order to secure potential victims of an abusive software ecosystem, we must first understand how that ecosystem operates in practice [3]. We take a first step toward filling this gap through measuring and characterizing the ecosystem by answering three key research questions.

Our first research question is:

- **RQ1.** How do nudification applications position themselves to clients via text and visual descriptions?

Acknowledging that nudification applications could be used with the consent of the image subject, we seek to understand whether existing applications foster and support those (and only those) use cases. Additionally, we seek to understand the general experience of users as they interact with these sites. Among the questions we ask: do users need to confirm that they are adults, do the sites support the nudification of all people, and how do the sites communicate about the role of consent from the image subject? Understanding the user experience on these sites, and how these sites position themselves and communicate with their users, is essential for having an informed conversation about their dominant use cases, the benefits and harms of these applications, and, to the extent that harms dominate benefits, how the broader research, industry, and policy communities might go about mitigating those harms.

Next, we seek to answer the following research question:

- **RQ2.** What features do nudification applications advertise?

In our study, we seek to understand and catalogue the advertised features of popular nudification applications — these are the features that they purport to offer, and hence the features that users believe that the applications will provide. We chose not to experimentally verify that each nudification site actually provides all (and has no hidden features) the features that they claim, nor did we upload images of people for nudification to these sites to test these features, as we did not believe it would be ethical to upload images of real people to potentially adversarial entities, including images of people on our research team, or even already-public stock photos or images of celebrities. Instead, for our purposes, we focus on what features users believe they will have access to if they choose to use the nudification site. Even if not all nudification platforms provide all the features that they claim to, if users seek those features today, we conjecture that future instantiations of these platforms will provide those features. Hence, knowing the features that these applications advertise, and what features these applications believe that their users want or will want, is valuable.

Lastly, we seek to answer the following research question:

- **RQ3.** How do nudification applications monetize?

Prior to the full initiation of our research study, we gained preliminary experience through the interaction with several nudification sites. Our preliminary interactions uncovered a diverse monetization ecosystem, including both conventional payment systems (like credit cards) as well as newer, less conventional payment systems (like cryptocurrencies). Under the hypothesis that at least some nudification sites are problematic, one approach toward curbing their existence might be to challenge their ability to monetize. Additionally, as part of studying the monetization ecosystem, we sought to assess these sites' monetization strategies, e.g., do certain features (like full nudification) require a paid subscription whereas other features (like changing clothing) do not?

To answer our research questions, we collected a sample of 20 nudification applications (websites) and systematically analyzed them using the application walkthrough method [24] — a methodology used in several fields within and beyond computer science (e.g., [31, 45, 39, 55, 21, 12]).

In analyzing the applications, we found a problematic ecosystem:

- 19 out of 20 applications explicitly specialize in the undressing of women; only half of the websites mention that they expect the user to have the image subject's consent and fewer ask for affirmation that consent has been obtained.
- Most of the applications allowed for additional features beyond “undressing” (e.g., making the image-subject nude with their breasts and vulva visible in the imagery). For example, half of the applications allowed users to put image-subjects into sexual acts.

- These nudification applications make up a commercial ecosystem and, hence, targeting their commercialization features might be one way to protect against SNEACI. Furthermore, we see purposeful repackaging of these nudification features with 5 out of the 20 applications offering API access to their highest paying customers.

Stepping back, the computer security research community, as it is often defined, focuses on computing in the presence of adversaries. In some cases, the research focuses on studying adversarial features and the adversarial ecosystem, and in other cases, the research focuses on studying defenses against said adversaries. The ecosystem of image-based sexual abuse is an example of an adversarial usage of technology, and this work sits within that context and builds on prior research in and adjacent to the computer security research community on understanding and studying the emerging realm of synthetic explicit non-consensual AI-created imagery [6, 56, 54].

2 Background and Related Work on SNEACI

In this section, we provide broader context and historical information regarding the rise of synthetic non-consensual explicit AI-created imagery (SNEACI) as well as related scholarly works. We discuss other related work inline in the body of this paper.

Media “deepfake”s³ (i.e., image, video, etc) have become more widespread in the last decade both in their frequency and uses [43, 35].

The most common response to stopping the dissemination of these deepfakes is through the creation of deepfake detectors. Prior work has developed machine learning models trained to detect if the video is fake or different methods to spot defects in the video to determine that they are generated media [52]. However, deepfake non-consensual explicit imagery (SNEACI) needs to be addressed as a form of sociotechnical violence, where detecting the output of that violence is not sufficient to remedy it.

SNEACI is a social issue of sexual harassment and abuse being perpetrated through deepfake models, largely against women [27]. Deepfake generators have been used to create SNEACI since 2017 and often were centralized around women politicians [27]. These political deepfake SNEACI’s have even been seen as recently as 2024 against a woman politician in Miami⁴ and have also recently been seen used against celebrities like Taylor Swift and social media content creators like Pokimane [23, 7, 50].

In the past, individuals have encountered image-based sexual abuse (IBSA) where real explicit images were spread across the internet or to specific parties to extort, exact revenge, or control a victim [56]. Since then, laws have been put in place in 47 of the 50 states in the United States of America which have criminalized IBSA in varying degrees [49], though such activities continue. SNEACI creates a form of IBSA where the images are “deepfakes”, or synthetically generated imagery.

New research within the computer security and adjacent communities have begun to evaluate awareness, prevalence, and societal norms around the creation, sharing and viewing of this content [6, 54]. Additionally, research has explored how those targeted by SNEACI might seek help [56]. Prior work has also studied perpetrator discourse [51], examining how deepfake creator communities discuss the technology (both in positive conversation and malicious use-cases) and how past attempts at mitigating abusive use cases have only caused the community to distrust moderation more [51]. In sum, past work on deepfakes primarily focus on human perspectives [5] — about the technologies they use to create deepfakes, victim-survivor experiences, or societal perceptions — rather than studying the systems capable of producing SNEACI. We seek to address the latter in this work through measuring and characterizing the AI nudification ecosystem. Such examination is necessary to understand how these technologies may be problematic and to think about how to defend against such abusive applications (e.g., intervene against their monetization sources, technical infrastructure, etc.) as has been done in other areas such as Stalkerware [16].

3 Methodology

Nudification applications have yet to be characterized. These applications are not only technical — their features build on computer vision and AI techniques — but sociotechnical: to the extent that they are used without the consent of the person in an image, their use is embedded in the sociocultural context of sexual abuse. Thus, we adopt a method to examine this ecosystem through a sociotechnical lens: Light et al.’s application walkthrough method [24]. This method expands on the HCI method of cognitive walkthroughs [57] used to assess issues of usability – which was previously

³“Deepfake” is the widely used colloquial terminology for media generated through AI. However, this terminology can often have a negative connotation because of its historical use in politics, especially against women.

⁴This person has since dropped from the race that she was in.

applied in the security literature as early as 1999 in “Why Johnny Can’t Encrypt” [58] — to examine an app’s interface in order to assess:

“its technological mechanisms and embedded cultural references to understand how it guides users and shapes their experiences. The core of this method involves the step-by-step observation and documentation of an app’s screens, features and flows of activity — slowing down the mundane actions and interactions that form part of normal app use in order to make them salient and therefore available for critical analysis . . . this process is contextualised within a review of the app’s vision, operating model and governance. [57]”

In this section, we both review the Light et al. method [24] and detail our approach to selecting a representative sample of applications, collecting the data, and analyzing that data via a walkthrough analysis. We additionally discuss ethical considerations, our positionality, and limitations of our approach.

3.1 Dataset

We leveraged two categories of sources for identifying AI nudification applications, i.e., websites offering AI “nudification” services. While there may have been other types of applications to generate SNEACI, we specifically look at AI “undressing” applications. During this time, we did find one application which was also an AI Face Swapping application; however, the method in which it was identified classified it as an AI “undressing” application. As this AI face swapping application may likewise be encountered by people seeking AI nudification applications, we kept it in our dataset. We continued to collect AI nudification applications until we hit saturation, i.e., no new websites appeared with further iterations. First, we reached out to a number of NGOs known for their interest in SNEACI and were given different lists of nudification applications. Second, we conducted preliminary Google searches to find “Top [X] Nudifying Apps” articles. These articles outline several AI nudification applications and often discuss features. After identifying AI nudification applications from the NGOs’ lists, we hit saturation of AI nudification applications after visiting three “Top [X]” articles. Between these two categories of sources, we found a total of 24 websites. While we initiated our study and developed our methodology much earlier, at the time of our final walkthroughs and data collection for this paper (July 2024), four of these websites were no longer accessible. As such, our researchers examined 20 website applications hosting AI nudification tools. While these websites would be easy for others to find, we choose to not name the websites in this paper because we do not believe that it would be responsible for driving more traffic to them.

3.2 Application Walkthrough

Light et al. [24] lays out a walkthrough of application UI’s that address the app in three stages: *entry*, *everyday use*, and *exit*. We use this methodology in our research to complete a critical analysis of AI nudification applications. The walkthrough analysis provides an understanding of the interactions on the website and a user’s likely expectations when interacting with the website. As such, throughout this analysis and under the Light et al. [24] approach, researchers should take screenshots of user interfaces and make notes on the provided features and how they are presented allowing the researchers to create the complete context of the infrastructure of an application.

We conducted a walkthrough of the 20 website applications hosting AI nudification tools, as identified in Section 3.1. These applications present a commercial storefront to purchase the applications and image generation with varying features. We completed a walkthrough of all 20 website applications in two parts. As an addition to Light et al.’s methodology, we double coded the first half of the repository. In the first part, two researchers walked through the first 10 websites separately, and we computed a Cohen’s Kappa for inter-rater reliability of 0.88 looking at our codes for how the applications present themselves both visually and textually, when age confirmation arises, the app’s features, payment processors, and monetization methods. Second, to ensure uniformity, one of the researchers from part one performed the walkthrough analysis for the remaining ten websites. The main researcher took screenshots of many of the user experiences upon entry and traversing the site. They also took notes. These screenshots and notes were shared with all of the other researchers for context about the applications and were discussed and refined over the course of numerous meetings.

3.2.1 Entry

The initial stage of the Light et al. [24] methodology, *entry*, refers to what the user first encounters as they access an application and extends into account registration and terms of services. In particular, researchers document the sign-up and registration process through well-articulated observation notes and examine the first elements the user interacts with. This includes required user information for sign-up (e.g., confirmations and questions asked), the registration and

sign-in process (e.g., Google sign-in or other sign-in options), and design elements of the entry pages. The *entry* defines the user's experience and affects the user's perception and understanding of the product before the user engages the product. As Light et al [24] argue, researchers can leverage this analysis to understand the intended use case of the platform.

In our approach: At the *entry* into a website, each researcher cleared their cookies and loaded the website in an incognito browser for coding and viewing. We denote the features on the *entry*, such as the presence of nude images, age verification, sign-in/sign-up, slogans, etc.

3.2.2 Everyday Use

After the *entry* phase, researchers move into the *everyday use* stage. The *everyday use* stage explores the functionality of the application and can vary on the extent and time requirements (e.g., basic walkthrough requiring minimal time or a thorough analysis requiring extensive time). While full analysis in this stage can provide a thorough understanding of the application, even walking through the basics of an app's functionality can provide an understanding of what activities it enables, limits, and guides the user towards. Using this approach, the researchers explore not only the available features but also the workflow of using the product (e.g., the various interfaces and how they interact).

In our approach: To map the *everyday use* of the website, we first traverse the links within the menu bar at the top of the webpage (e.g., dropdown menu). We collect the terminology used to describe the functionality of the website (e.g., taglines or slogans) that appears around the website. Although some of these AI Nudification Applications lock some or all uses of their platform behind a subscription paywall, they report the functionalities associated with each subscription tier. We collected the reported functionality from these descriptions and the associated cost. The advertised functionality indicates user expectations from the nudification platform. As such, it denotes the use case, priorities, and values of the creators of the websites. However, we find that the terminology used to discuss the functionality varied between applications. While we noted how these applications advertised themselves, we also coded functionality based on the application's functionality instead of only the term used by the application for the tool. For example, a website may describe itself as an "AI undressing application" but offer features that allow AI face swapping into explicit videos.

As discussed in Section 1, we specifically look at the *advertised* features mentioned on the AI nudification applications. A user purchasing these products will likely expect the features to match what the applications advertise. Furthermore, we presume that the advertised features are what the company strives to be able to provide.

3.2.3 Exit

Finally, the third stage of the walkthrough is the *exit*. This stage identifies how the user can exit the application. This includes but is not limited to logging out of the application as well as deleting their account. Particularly, walkthrough analysis classifies the user flows for deleting the account and the user's understanding of their data after account deletion.

In our approach: Finally, we finish the walkthrough at the *exit* phase by looking at if and how a user could delete the account they had created.

3.3 Ethical Considerations

Conversations on the ethics of researching these websites and the actions we took to study them were central to our meetings. While our work does not interact with human participants, there were several ethical considerations when conducting this research.

First, by researching the features of AI nudification applications, we may develop and communicate knowledge about these AI nudification applications that could be used by bad actors. Additionally, if we were to name or link to these sites, driving traffic to them could increase their ranking within search engine results. As such, we do not name the websites in our study when presenting our results. And, overall, we believe that the knowledge we develop and communicate will also be valuable for those seeking to mitigate the harms of SNEACI.

Whenever we experimented with accounts, e.g., sign-in with Google, we ensured that those accounts were not linkable to the researcher. While collecting data from a university campus, the researchers ensured that other people in the environment were not exposed to the imagery on the websites in our study.

Furthermore, and as we would encourage all research efforts that might expose researchers to potentially problematic and triggering content, we hired a therapist with open office hours every other week that our researchers could visit; the therapist was also available on demand, if needed.

The contents of this paper may be triggering to some, for a wide variety of reasons. In addition to including a content warning at the start of this paper, we are also hiring an expert for a sensitivity read of this paper, and will revise as appropriate. For example, in addition to thoughtful considerations of how to communicate about image-based sexual abuse, we believe that it is imperative to speak thoughtfully and respectfully about the differences between gender and sex even though the nudification applications themselves do not appear to make such a distinction. Whereas our team feels confident in considering and writing to these types of considerations, we believe it is valuable to hire an external expert for an independent sensitivity evaluation. In this paper, we chose to show artists renditions of the UI's presented on the AI nudification websites. We believe that the visual depiction more clearly articulates the features and viscerally expresses the harm which was think is important. We chose to do stylized sketches in order to balance reader comfort, expression of research results, and risks of displaying actual bodies from the website.

3.4 Author Positionality

As researchers, while we strive to be objective, we realize that our personal backgrounds and perspectives could influence our interpretation of the data that we collected. Thus, we describe relevant elements of our identities and positionality here. Half of our authors identify as woman, which is the demographic that we found was the primary target of these applications.

All members of the team believe that the non-consensual uses of these applications are *not* acceptable, which is consistent with the dominant view of acceptability [6, 54]. Thus, in interpreting our results, we focus on the harms that these systems can create and the rights that are violated when used non-consensually.

3.5 Limitations

We would like to acknowledge a number of different limitations; however, we do not believe that they take away from our analysis.

We acknowledge that our repository may not include every AI nudification application; however, we are only trying to characterize the ecosystem. This case study can be a representation of the population to build upon.

We acknowledge that the location from where we access these applications may affect what appears on their UI. In both creating our database and analyzing the ecosystem, we accessed all of these websites from the United States of America. Other advertised features may be present if accessed from other regions.

We acknowledge that our characterization of the ecosystem is based off of the advertised features. These features could differ from the actual features present when generating the images through their UI. However, any user accessing the website will believe these are the features as well. Furthermore, we believe as the AI features advance, the applications will get closer and closer to the advertised features.

3.6 Terminology

Throughout the paper we use three terms frequently: “women”, “application”, and “feature”. We use the term “women” to describe those depicted on the websites we analyze who appear to the researchers to present as women based on their dress or anatomy (visible breasts, vulva). We acknowledge that there are significant pitfalls in attempting to perceive gender based on visual depiction: gender is a complex concept and typically people self-identify their gender [46]. We use the term “application” to refer to the websites that serve as virtual UI store fronts to sell AI tools to consumers to automate the process where an image subject in the submitted image is returned in a “nude” form. Finally, we use the term “feature” to refer to specific features offered by these “applications” such as those listed in Table 1.

4 Entry: User Experience

This section reports on the first stage of the walkthrough, the entry phase. Here we study what users encounter as they enter the web applications. Users rarely have to create an account before encountering nude bodies though the option to login/create an account was always present. In this section, we will walk through what users experience upon entering the websites, what was required (or not) of them to be allowed entrance, how and when they had to make an account, and the terms of services of the application.

Aligned with qualitative methods, our analysis aimed to surface general themes about the AI nudification applications, rather than quantify their prevalence. Accordingly, we report the appearance of themes using the following terminology: a few (less than 25%), some (25-45%), about half (45-55%), most (55-75%), and almost all (75-100%).

While most websites ask that users are 18 or older, not all do. Recall from Section 3 that our data was collected from a region in the U.S. that does not have an age verification law for explicit content. As our researchers opened each website, in incognito mode and with cleared cookies, seven of the 20 (33%) websites immediately asked for confirmation that the user is 18 or older. As our team traversed the different pages, 14 out of the 20 *do* ultimately ask the user to confirm their age: seven when the user first loads the website, two when the user goes to login or create an account, and five when users start the process of generating “undress” images or purchasing the feature to do so. For the latter two categories (when age verification eventually happens but not at first-load), we find that five show nude content prior to confirming the user is over the age of 18; in more detail, three show blurred imagery where nipples and vulvas are blurred (the rest of the genitals are still visible) and two show complete nudity of the breasts and vulva.

Of the remaining six websites that *do not* verify the user is 18 or older throughout the regions of the sites that we navigated, we observe a spectrum of content visible to users. At the extreme, on its landing page, one of these six websites shows AI-modified images of celebrities engaged in sexual acts along with false news articles about their actions. In contrast, another of these six websites provides descriptive text of the form “fake nudes of a person using AI” but no images. While we have not used these six services to create images, as discussed in Section 3, our study of all six sites suggests that users will be able to generate “undressed” images without any age confirmation.

Additionally, in our walkthrough, we found that only seven of the 20 applications required the image subject to be over the age of 18, and all seven of those only mentioned that requirement in their terms of service.

Finding #1. For some sites, users can *generate* undress images without confirming that they are 18 years old or older. Further, only some sites mention that the image subject should be 18 years old or older, and on all those applications, we only encountered such mention in the terms of service.

19 of the 20 websites in our study explicitly specialize in the undressing of “women” or the nude female form, though the websites present it in different ways. Most of the websites show their focus on women through imagery. 17 of the 20 websites in our study show images of women in varying levels of nudity on the landing page. 13 of these 17 sites show women *without* clothing but not in sexual positions, and another three show womens’ faces and bodies but include mens’ body parts as the women and men engage in sexual acts; mens faces are never shown. Of the three websites that do not show photos on the homepage, one has a “gallery” with women’s nudified images on it with the caption “nudification samples” and another specifies in text that it is “specialized in the female form”.

Additionally, the website text also communicates a focus on women. In their text descriptions, 14 sites exclusively mention women and/or girls, and no other genders: two of 20 sites refer to the subjects exclusively as “women” (e.g., “AI algorithm for generating nudes from photos of clothed women”), five refer to the subjects as exclusively as “girls”, and seven refer to the subjects as either “women” or “girls”. For example, one site says “girl of your dreams”. The use of the word “girl” to refer to adult women can be interpreted in a number of ways. For example, it can be viewed as demeaning and sexualizing [41]. However, “girl” is also used in empowering ways by a number of different subcultures of hyper-femme expressions [27]. Further, two sites use the word “your” as in “your girl” or “your woman” or “your girlfriend”, potentially suggesting the user’s possessiveness of the woman subject.

In summary, 19 of the 20 nudification application web pages convey a specialization in the nudification of women, though they vary where and how that specialization is communicated. This focus aligns with findings from prior work on non-AI facilitated IBSA, which found that over 90% of content on “revenge porn websites” was of people who based on visual presentation appeared to the researchers to be women [53].

Finding #2. 19 of the 20 undress websites in our study explicitly specialize in the undressing of women, though they vary in the explicitness of the images shown on the landing pages.

Half of the undress websites in our study say that users must obtain “consent” of the image subject before undressing them. Our team additionally studied the consent-related text that users might encounter as they navigate the nudification applications as well as consent-related text in the sites’ terms of service, which users may not encounter but which are still present on these sites.




Figure 1: These are artistic renderings based on several real UI’s from the nudification applications we study. In the rendering, less details of the human form are shown. This figure shows the interface for uploading a photo and various different options that a user can pick from to customize the resulting image. On the left, a play button is added over the “sex” option to indicate that some applications can produce both image and video outputs. On the right, we indicate that some of the features are only available after payment. Artwork credit: Hanna Barakat.

Seven of the 20 sites include text during age verification or before image generation that states that the user must have consent of the image subject.

All seven of these sites — and an additional three sites for a total of 10 — include text in their Terms of Service that state that a user needs consent from the image subject to upload their image to the AI generator.

Consequently, 10 out of the 20 applications do not state a requirement in their Terms of Service that the user must have consent of the image subject in order to have their image modified, and we did not encounter consent-related text during the “normal” (non-Terms of Service) user flow in the 13 of the 20 applications.

Finding #3. Half of the websites mention in some form that they expect the user to have the image subject’s consent before using their tools; however, only some websites ask for the user to confirm this assumption.

Registering for an account was simple and mainly gated generation or product purchase. Our researchers were able to observe and traverse much of the web pages before being required to register for an account. While users immediately can login or register upon entrance to the web pages in most cases, three applications would allot users login credentials after purchase. For those three applications, users were not immediately able to login or register upon entrance to these sites. And in two of these websites, they could not register an account until after purchase (which the researchers did not do). In one application, users were never given logins. Instead, they can purchase a “key” which gives them an set number of generations.

On 10 out of the 20 of the websites, our researchers did not have to create or login until they were attempting to purchase features or generate images. In four applications, no login is ever required for the researcher to begin to generate imagery.

Many of the websites allowed researchers to register or login using authentication through other applications. In addition to traditional email-based account registration and login, we found that users could also login and register accounts via sign-in through Discord, Google, Twitter, and Apple. 4 websites allowed Discord-based account sign-in. Google was supported by 12 websites. Apple and Twitter were used to support logins to 3 and 1 websites, respectively.

5 Everyday Use: Features

In the second phase of our Light, et al. [24]-based walkthrough, we analyze the advertised features of the 20 nudification applications in our study.

Table 1: AI nudification applications have a wide variety of AI nudification features. Base features such as the AI Undressing Tool, Human-Assisted AI Undressing, and AI Face Swapping are inherently what classify these applications as AI nudification applications; however, some applications offer features that put image-subjects into sexual acts or allow the user to change the image subjects clothing. This table explains each terminology used to denote these features.

Likeness-Based Image Features		
Feature	Capabilities	Count
AI Undressing Tool	User inputs a photo of an image subject to receive back a “nude” of the image subject. The AI model is essentially trying to predict what is under the clothing.	18
Deepnudes	An instance of an AI Undressing Tool. Users input an image of an image subject. The AI model will modify the image to put the image subject in to different sexual scenarios as specified by the user such as sexual scenes (see rightmost image in Figure 1).	9
Clothing Changing	An instance of an AI Undressing Tool. Users input an image of an image subject. The AI model will modify the image to put the image subject in to different clothing such as lingerie or bikinis.	5
Human-Assisted AI Undressing	An instance of an AI Undressing Tool. Users give a photo of an image subject to “Experts” who use AI tools to make the image subject “nude”.	1
AI Face Swapping	User inputs a photo of an image subject with the option to also input another photo/video of a person and they receive back an image/video where the face of the image subject has been put onto the person of the second image/video. They use default destination photos/videos provided by the service if they do not upload one themselves.	1

The nudification ecosystem has a wide variety of features. However on each application there is at least some type of AI feature which allows a user to upload a photo and receive back a “nude” version of the image subject. While the exact feature set may vary, this underlying feature is what defines our repository — the ability to make someone “nude” through an AI feature. Many of these applications allow for users to pay for additional features that allow them to fulfill a number of tasks. This may range from changing what clothing the image subject is in (e.g., putting them into a bikini or in lingerie) to AI-based features that put image targets in AI-altered videos with sexual acts.

5.1 Likeness-Based Image Features

We focus our analysis first on the features that these nudification applications offer with respect to the manipulation of the likeness of an image subject (summarized in Table 1).

There are two broad categories of features that allow a user to see the image subject in the nude: undressing and face swapping. We find that 18 of the 20 applications offers an AI Undressing Tool in which a machine-learning model is purportedly trained to ‘predict how the breasts and vulva of an image subject look under clothes and alter the image to represent their as-predicted nude bodies (such ML functionality may be termed in-painting or body mesh estimation in the computer vision community [61, 25, 20, 4]).

Of the two applications that did not offer AI undressing, one application offers Human-Assisted AI Undressing features. Users pay per generation and attributes they want added to the image, but they hire an “expert” to do the generation using features that the expert has access to offline. The remaining application offers AI Face Swapping features. This feature essentially presents the image subject nude, but instead of using the machine model to predict how the image subject’s body may look under clothing, the AI model is trained to morph the face of someone in an image or video into that of the user’s uploaded image subject. In AI Face Swapping, at minimum the user must supply the machine learning model with the image subject’s face that will be put onto another person’s body; however, they also have the option to supply the materials that the face will be “swapped” onto.

It is unsurprising that 18 out of the 20 applications offer AI Undressing tools because our application selection process (see Section 3.1) specifically selected for this feature. However, the two applications that do not advertise AI Undressing but AI Face Swapping and Human-Assisted Undressing were included in the lists of AI Undressing tools we discovered or were provided. We retained them in our dataset because they are among the applications a user might encounter if they were to seek out AI nudification applications.

Website	Lowest Tier	Highest Tier	<div style="display: flex; justify-content: space-between; font-size: small; margin-bottom: 5px;"> AI UndressingFace SwapHuman-AssistedDeepnudeClothing ChangesBody ModificationAnimificationFur-ficationUnique FiltersAI Image GeneratorAI Hentai GeneratorLLM SeakingGeneration SpeedQueue PriorityImage QualityEarly AccessFile Size IncreaseWatermark RemovalSave HistoryDelete HistoryAPI AccessCustom WatermarksParallel Generation </div>																				
			Nudifyng	Manipulation	Modification	Additional Image Capabilities	Utility	Repackaging															
S01	Free (F)	T3	T3	T3 F T1-T3					T1-T3 T2-T3 T1-T3	T3													
S02	Free (F)	T3	T3	T3 F T1-T3					T1-T3 T2-T3 T1-T3	T3													
S03	Free (F)	T1	F-T1	T1 T1-T3				T1 T1 T1	T1											T1			
S04	Free (F)	T3	F-T3						T1-T3 T1-T3	T1-T3													
S05	Free (F)	T3	F-T3	T1-T3																			
S06	T1	T3	T1-T3																				
S07	Free (F)	T2	F-T2	T1-T2	T1-T2 T1-T3 T1, T2	F-T2			T1-T2	T2										T2			
S08	Free (F)	T3	F-T3						T1-T3 T1-T3	T1-T3													
S09	T1	T5	T1-T5						T1-T5 T1-T5	T1-T5													T5
S10	Free (F)	T3	T3	T1-T3 F-T3					T2-T3	T3													
S11	Free (F)	T3	T3	F-T3					T2-T3	T3													
S12	T1	T2	T1-T2	T2 T1				T1-T2	T1-T2											T2			Private
S13	T1	T4	T1-T4	T4 T4					T1-T4 T1-T4	T4													
S14	Free (F)	T2	F-T2	T1-T2																			
S15	Free (F)	T1	T1						F											T1			T1
S16	Free (F)	T1	T1						T1											T1			T1
S17	T1	T3	T1-T3						T1-T3 T1-T3											T1-T3 T1-T3			T1-T3
S18	T1			T1																			
S19	Free (F)	T3	F-T3						T1-T3 T1-T3											T1-T2			T3 T3 T1-T3
S20	Free (F)	T1	F-T1							T1													

Figure 2: The ecosystem offers a variety of features at different subscription tiers.

There are also a number of features which allow a user to put a subject in sexual scenes. The nudification applications in our dataset do not limit their features to “undressing” an image subject in-place. Rather, we encounter a diversity of features when studying the sites in our dataset.

One feature, which we refer to as Deepnudes, renders an image subject in a sexual act. See, for example, Figure 1, which is an artistic rendering of the UI of one application in our study. In this UI, a user can choose from “sex” (e.g., offering “missionary,” “cowgirl,” or “reverse cowgirl”), “cumshot” (ejaculation on the image subject’s face), “shibari” (a form of bondage), and “BDSM” (a term for a variety of sexual practices that include bondage, discipline, dominance, submission, and sadism). 9 of the 20 applications in our study offer similar “Deepnude” features.

Adjacent to the above category of features is the Clothing Changes feature, which is present in 5 of 20 applications we study. This feature allows a user to apply different clothing to the image subject ranging from bikinis to lingerie. The application that offers AI Face Swapping in images offered this feature for videos as well.

Finding #4. Half of applications allow for images to be modified such that the image subject is put into sexual scenes that are not in the original photo.

Some applications allow for modification of the image subject. Modifications can range from weight, body traits like breast or butt enhancements, or age (18+) changing to even altering the image subject into anime⁵ or furi⁶ versions of themselves. Body modifications for weight and age are seen in two and three nudification applications, respectively. Two applications offer unique filters and one application offers both animification and furi-fication models.

Overall, seven applications offer modes of modifying the images and how the image subject presents in the image.

5.2 Additional Features

We now summarize the additional features advertised on these sites, which speak to the broader set of services offered by these applications and, hence, services that might appeal to some of their users.

Some applications offer generation features. General image generation is offered by several of the applications we study. These can be used without any seed image uploaded by the user. Four applications offer general AI intimate image generation, one additional application offers AI Generation of Hentai⁷ specifically.

⁵A type of art/animation style that is attributed to Japanese animation.
⁶A subculture where people dress in fur suits depicting anthropomorphic animal characters.
⁷Hentai is an art style of pornographic animated video often attributed to or seen as a subset of Japanese anime.

Utility Features		
Feature	Capabilities	Count
Generation Speed Increase	Users get increased generation speed for each subscription tier they purchase.	8
Queue Priority	User purchase higher subscription tiers to increase their queue priority.	10
Generation Quality Increase	The subscription tier which the user purchases decides how high quality the images they generate can be.	7
Early Access	Users are given access to new features before lower paying users.	4
File Size Increase	Specific to AI Face Swapping, user’s can input files with larger file sizes for each subscription tier they pay for.	1
Watermark Removal	At the highest subscription tier, users can remove the watermarks from their images.	7
Repackaging	Users gain access to an API that can communicate with the AI Model and sell access to it on their own applications.	5
Saving History	Users can save images that they have created through the applications features.	4
Delete History	Users are able to delete their history of images or videos they have created through the applications features.	2

Table 2: AI nudification applications offer a number of features to help improve the user’s experience that we call “Utility” features. This table defines each term and the frequency with which we observed such features in our dataset.

One application offers synthetic sexual messaging. One application offers a trained LLM that will message with users and roleplay different sexual scenes with them. While this application’s LLM-based sexual messaging features do *not* appear to be connected with its AI nudification services, the fact that a single application offers both services suggests a future in which these services are combined, as hypothesized in prior research [9].

5.3 Utility

AI nudification applications offer a number of features to help improve the user’s experience that we call “Utility” features.

Enhancements to the generation and photo quality are available. Respectively, eight, 10, and seven applications offer increased generation speed, priority queues, and quality improvements through paid offers. We find that four applications also offer early access to future features with paid subscriptions. The application that offers AI Face Swapping also allows an increased file size for the option where paid users input a video for the image subject’s face to be swapped into.

By default, many of these applications watermark their images, and users have to pay for the product’s highest tier for them to be taken off. Through our application walkthrough, we find that seven of the undressing applications watermark the images they output and sell the removal of those watermarks at the highest tier of their application’s subscription.

Some companies allow their AI Model to be repackaged and sold on different applications as their own AI model by giving access to their API. *Repackaging* is when a company with a product, in this case the machine learning model for nudifying image subjects or any of the other features listed in this section, sells the product to other companies with the agreement that the other company can use their own label on the product but have little to no control over the product itself other than the label on it. In business literature, such repackaging is also referred to as “white labeling” [14], though we do not use that term here because of its similarity to “white listing”. We find that five of the 20 applications in our study offer access to their machine-learning model through an API. Additionally, one of these applications offers custom watermarks. Two applications explicitly offer parallel generation, where they can have multiple images being generated at the same time, which helps if someone wants to sell the product on their application.

Finding #5. Some applications offer access to their models through an API, allowing other applications to offer the same features without training their own AI nudification tools.

Many applications sell users the ability to save or hide their history. Applications are not in agreement on whether a premium feature is the ability to *save* or *delete* a user’s history.

We find that two applications make *deleting* a user’s history a feature. In contrast, the ability to save past generations is a feature on four applications.

6 Everyday Use: Monetization

As part of everyday use we additionally examined the cost of the applications’ features, how the applications monetized, the relationships to monetize, and the payment methods that enabled their monetization.

All nudification applications are commercial: they sell their features and offer very limited free functionality. Figure 2 details the features at different paid tiers of the applications we analyze. Fourteen applications offered free features, although the functionality of those features without payment was limited. Eight applications offered “free” nudification features, however all of these applications output the “free” nudified image in a blurred form or with a large enough watermark so as to incentivize payment to remove the blurring or watermarking. Four additional applications offered clothing changes for free, and two applications offered image generation for free. This was the extent of “free” features provided by the applications.

Across all 19 applications that offered AI undressing at scale, i.e., not the human-assisted AI undressing application, and if we consider the lowest-cost paid tier, the mean price per image was \$0.31 (min: \$0.06, max: \$1.00). At the mean price, the cost to generate 1000 such images would thus be \$310, but at it’s lowest pricing, that could be only \$6.00.

Finding #6. All AI nudification applications are commercial platforms, which could be a vector to incentivize guardrails against abusive use.

Platforms offer a symbiotic relationship with their users to give them a stake in the success of the platform. They do this in three relationships: Platform to User, Platform to Platform, and User to User(s). While traditionally we see monetization as a company selling either their product to other companies to sell (i.e., platform to platform) or company to consumer (i.e., platform to user), we see the rise a new third monetization relationship that essentially makes certain users a stakeholder in the company’s success (i.e., user to users) [17, 8]. In fact, eight of the AI nudification platforms in our study offer an affiliate program where select users make a percentage of the revenue generated by new users who create an account on the platform through their affiliate link.

In the nudification ecosystem we have seen affiliates being offered a range of 25% to 50% of the revenue. One application offered 25%, two offered 30%, two offered 40% and one offered 50%. Furthermore, two applications hid the split that they offer to their affiliates.

Because of the selective nature of these programs, our researchers could not identify all of the forms which an affiliate’s pay could be paid out; however, we were able to identify that three of the eight platforms which publicly discuss it offer payment in cryptocurrency.

Finding #7. Affiliate programs were offered by some of the applications. These applications offered between 25% and 50% of their revenue to their affiliates.

Five of the 20 applications offered a referral program to all of its users. All users could *refer* other users to the website. When the new user made an account, the account which referred them gains 2-3 credits that they could use for whichever tier they had already purchased.

Finding #8. Referral programs were offered to all users by some of the AI nudification applications. Referrals gave users 2-3 generations at their subscription level.

Traditional monetization relationships between the platform and users or to other platforms can be broken into three subgroups. Our research found there are 3 different monetization schemes used by the applications: *Feature Scaling* subscriptions, where the higher tier subscription a user pays for the more features they are given access to, normally with undressing and removal of watermarks at the highest tier and the quality, speed of generation, queue priority and number of generations allotted scaling with each tier; *Pro-Credit* subscriptions, where the higher tier

subscription a user pays for gives a higher base number of generations allotted as well as a higher discount from which they can purchase further generations but all of the features are *generally* the same; and *All-for-One* subscriptions where one price purchases all features of the product.

Eight of the 20 applications we study use the *Feature Scaling* subscription model. On average, these applications cost \$0.34 per generation of image over all tiers. As a general trend, purchasing a higher tier of subscription does decrease the price per generation; however, it is not a one-to-one because higher tiers allow the user access to a larger feature set. For example, as seen in Figure 2, Application S01 allows for Clothing Changes as a "free" tier, but Undressing is not accessible under Tier 3 and all images will be Watermarked till Tier 3 where the user gains access to Watermark Removal.

Seven out of the 20 web applications used the *Pro-Credit* subscription model. While AI nudification features are constant between the different tiers, this subscription also purchases scaling discounts for future credit purchases which need to be taken into account when looking at the cost-per-generation. On average, the cost-per-generation over all tiers for these applications is \$0.30 which is slightly less than the Feature Scaling subscription model.

The *All-for-One* subscription model was used by four web platforms. These applications only ever had one price. On average, they cost \$16 per month. Users could generate each image for an average \$0.35.

The one application that offered Human-Assisted AI nudification made users pay at a rate of \$20 per undressing image, and more for an image with greater sexual content. This is an outlier for our dataset and represents something closer to a crowdworker paradigm than a point-and-click application paradigm.

Five of the applications offer APIs for service resale. Five of the applications offered repackaging options where users could gain access to their model through an API. The cost of API access ranged from \$20 to \$299 with a mean of \$92. At the highest tier of the subscription, one platform offered repackaging options for \$299 per month. This platform offered those who purchase that tier to "resell our service using API". This tier also gave access to a manager and marketing plan.

Cryptocurrency is the most popular avenue for payment in this ecosystem. The commercial aspect of these platforms means that all applications must have some way to get their payments. Many of these applications rely on third party payment processors or cryptocurrency transactions.

Almost all of the nudification platforms heavily rely upon or incentivize paying in cryptocurrency. 17 out of the 20 applications offer users the option to pay through cryptocurrency. Of those, four of the platforms use incentives to encourage users to pay with cryptocurrency. They do this by offering an additional 10-20% discount on the subscription purchase. Furthermore, five of the 17 platforms offer cryptocurrency as their only payment method.

Of the three web platforms that did not offer cryptocurrency payment capabilities, based on their UIs, these platforms relied on PayPal and credit card payments to generate their revenue.

Eight of the 20 platforms offered the option of PayPal payment while seven accepted credit cards. Two nudification platforms offered to be paid through Patreon and three offered Apple Wallet. We also found that Steamskins, Venmo, PhonePe, Alipay, CashApp, Google Wallet, and Amazon Pay were each accepted by one platform. We include a figure representing a visual breakdown of these results in Appendix 6.

7 Exit: Account Deletion

At the third stage of an application walkthrough, we look at the ease of account deletion. For this step, we made 12 accounts through Google's single sign on when it was offered. On the other five websites, we created them through email. The final three websites did not allow us to make accounts.

Only one application let our researcher's delete their account through the website. Our researcher made accounts on 17 of the 20 (three websites did not allow for account registration) websites and looked to see how to delete their account. Only one of these websites allowed the researcher to delete their account.

8 Discussion and Conclusions

Using the walkthrough methodology of Light et al. [24], we study and develop an informed understanding of the nudification application ecosystem on the web. We find that point-and-click AI nudification features are ubiquitous, that these applications position themselves to specialize in the generation of nude or sexually explicit images of women.

Below, we reflect upon our findings, first by considering our findings in the context of the impact to potential victims and the role of consent (or the lack thereof), and then by exploring approaches to mitigate SNEACI.

8.1 False Representations of One’s Body

Representations of self are profoundly personal regardless of their form: static images, textual representations of self in early online communities through emerging representations like VR avatars.

Prior research has found that digital self-representation (e.g., avatars) can have significant effects on one’s self-concept, and on social behavior [60]. Since the MOO communities of the 90s, people have described experiencing harms to their digital representations much like harms to their physical bodies [26, 10]. Prior work on non-synthetic IBSA confirms this finding, noting that the “experience of image-based sexual abuse [is] embodied — experienced in and through their bodies, altering their sense of bodily integrity” [18]. Such harms are not merely “emotional”, but pervasive throughout victim-survivors’ lives. Prior work notes for example that “companies may refuse to interview or hire women and minorities because their search results include nude images or deep-fake sex videos” [10]. Even consensually-altered pictures of one’s self can have negative effects on people’s self-esteem and lead to self-objectification [33]. As Citron writes in “Sexual Privacy” [10], “even though deep-fake sex videos do not depict featured individuals’ actual genitals, breasts, buttocks, and anuses, they hijack people’s sexual and intimate identities.”

Historically, non-consensual photos and their sexual alterations were mainly through cutting out faces from magazines or public photos and tapping them to intimate imagery [15]. This type of alteration still meant that image subjects had to have their images publicly printed, and the alterations were relatively noticeable. After that, we saw a rise in the number of photoshopped images [19]. However, their resemblance to the original was mainly gated by the user’s skill level in Photoshop. Now, with the rise of commercial AI Nudification platforms, users without technical skills may still create altered images of a person, and they can do so with a few clicks of a button.

Our analysis finds that existing AI nudification tools appear positioned to specifically create abuse material of those who have not consented to their images being used. Below we offer suggestions for reducing the availability of these abusive tools.

We note that it is possible that there are legitimate sexual expression uses for such tools. Enabling potential non-abusive use cases while not allowing for abuse will hinge on verification of consent of the individual who is depicted in the images uploaded into such tools. This was not the case in the applications we studied. As seen in Section 4, only 10 of the 20 applications mention consent in their ToS. Applications aiming to offer consensual nudification functionality could draw on existing practices in the sex work industry to verify consent of the individual depicted in the input imagery, and that the individual depicted is over 18, by using a combination of: 2257 forms [11], which document two forms of ID to verify that performers are over 18, manual review of every piece of uploaded content as done by Pornhub [44], and cross-checking of identity verification via biometrics with those depicted in the uploaded content on a recurring basis, as done by OnlyFans [40]. We note that centralized approaches have inherent privacy tradeoffs, as they require the company verifying consent to maintain a collection of sensitive identity information [48]. We found no instances of such verifications in our study of 20 nudification applications.

8.2 Reducing the Availability of Abuse Tools

Prior work in computer security research suggests the value of considering the full pipeline of actions and actors in an abusive ecosystem [1], and then advocating that those stakeholders supporting elements of the abusive ecosystem remove their support (e.g., demonetization). Such advocacy requires (a) identifying stakeholders in the pipeline and (b) respecting freedom of expression, including sexual expression, by distinguishing abusive website behavior from non-abusive behavior such that actions are restricted to abusive websites.

Stakeholder Identification. Our research identifies the existing stakeholders in the nudification pipeline as: the creators of the models used to power the nudification websites, the nudification websites themselves, the payment processors for the nudification websites, single-sign on providers for the websites, and platforms that support discovery of the websites through search or app stores. We identify and catalogue these stakeholders via manual analysis. Future work may seek to build an automated stakeholder identification pipeline. This could be a difficult process; however, our research by characterizing the ecosystem begins the work of identifying stakeholders such as payment processors that enable the commercialization of AI notification applications.

Such information could be used by regulatory bodies or voluntarily by platforms that provide single-sign on or application discovery to reduce their contribution to the abuse ecosystem.

There may be technical challenges to identifying some stakeholders. Many of the nudification websites we analyze refer to having been built on top of other open-source models. Indeed, the functionality used by these platforms has long been a subject of study in the computer vision community (e.g., inpainting, body mesh estimation). Consequently, there is a long history of open-sourcing models that offer the functionality necessary to launch an AI nudification website, without requiring the users of those models to have significant expertise. The proliferation of open-source models on platforms such as Hugging Face in recent years have added to the availability and ease of offering such functionality at commercial scale.

While the computer science community has long centered the importance of open-source software and data, recently the ML community has shifted toward considering whether access to high risk research tools should be restricted: “Models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, e.g., by requiring that users adhere to a code of conduct to access the model” [34]. Such steps have precedence in other domains such as medicine: NIH for example requires an application to access genetic data they hold in their database [38] and computer science community projects like the National Internet Observatory similarly require an application and multiple board reviews before data / tool access are provided [36].

Future work should consider what steps for responsible use should be taken by open-source model providers, including but limited to tracking downstream uses of their models for abuse and governing access to model features, as well as technical interventions for identifying “responsible” model providers whose features underlie downstream abusive websites or copycat models.

Additionally, our research identified that five of the applications we analyzed offered access to their functionality via an API such that a user could repackage the application as their own on another website. This option to repackage the AI nudification application onto another store front brings up a number of concerns. For one, while these commercial sites already decrease the cost of entry for users to create abusive content, repackaging these applications also decreases the cost of entry for new applications building on existing functionality. Whether the initial application means for their AI nudification features to be used in non-consensual use cases or not, they do not have control over how the new application decides to market these AI nudification features. Furthermore, repackaging could act as a shield for malicious actors allowing for a chain of shell companies to hide behind.

Signals to Identify Abusive Behavior. There are a number of signals that could be used to identify abusive nudification websites, with the goal of distinguishing them from (a) general GenAI websites and (b) non-abusive GenAI focused on sexual expression: e.g., the text used on the website, the SEO terms and/or advertisements used to promote it, and the prevalence of use of the site for abusive behavior.

Our work finds that the lexicon used by these websites to promote their services is still evolving, and varied, with sites describing the same features in several different ways. However, a common lexicon begins to emerge, particularly around the promoted use cases of the services with e.g., the use of phrases like “AI Undressing” or “AI Deepnude.” A common lexicon of terms that arise with high frequency in the text or SEO keywords of abusive websites and distinguishably lower frequencies on non-abusive websites could allow for automated identification of those sites to, e.g., payment processors, single sign-on providers, model owners whose models are being used by those platforms.

Future work may seek to map the pipeline of clearly abuse-suggestive advertisements that link back to particular websites. We hypothesize based on emerging evidence from investigative journalism [28, 29] that the advertisements used by abusive sites may be clearly distinguishable from those used by non-abusive sites.

Finally, our research finds that some AI nudification platforms already watermark the content they produce. Regulators with a centralized processing infrastructure for reports of abusive imagery [13] could use these watermarks for attribution: to keep track of which websites are generating the most reported abuse material. There are two current limitations to this approach. First, our analysis found that several of the AI nudification websites we examined offered the option for users to pay to remove the watermark applied by the platform. Regulation to require AI-generators to watermark their content is being considered in multiple jurisdictions [47] and would be necessary to give regulatory enforcement power against platforms that do not comply. Second, regardless of any regulations, watermarks would need to be resilient to removal by users or others [59, 42].

Acknowledgements

This work was supported by the US National Science Foundation under grants CNS-2205171 and CNS-2206950. We would also like to acknowledge the work done by Hanna Barakat to help us with our figures, and the sensitivity reads completed by Miranda Wei and Vaughn Hamilton.

References

- [1] Coalition Against Stalkerware. <https://stopstalkerware.org/>, 2023. Accessed: 2023-02-05.
- [2] Executive Order No. 14110 - Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, 2023.
- [3] R. Anderson. Why cryptosystems fail. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 215–227, 1993.
- [4] A. O. Bălan and M. J. Black. The naked truth: Estimating body shape under clothing. In *Computer Vision–ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part II 10*, pages 15–29. Springer, 2008.
- [5] A. R. Brieger. Empowerment or exploitation: A qualitative analysis of online feminist communities’ discussions of deepfake pornography, 2024.
- [6] N. Brigham, M. Wei, T. Kohno, and E. Redmiles. “Violation of my body:” Perceptions of AI-generated non-consensual (intimate) imagery. In *33rd USENIX Security Symposium (USENIX Security 24)*, 2024.
- [7] B. Britton. They appeared in deepfake porn videos without their consent. Few laws protect them. *NBC News*, Feb. 2023.
- [8] M. R. Bruce, A. F. Adekoya, S. Boateng, and P. Appiahene. Prevalent user-centered monetization techniques in social media. *International Review of Management and Marketing*, 13(1):19–28, Jan. 2023.
- [9] I. Cheong, A. Caliskan, and T. Kohno. Safeguarding human values: Rethinking US law for generative AI’s societal impacts. *AI and Ethics*, May 2024.
- [10] D. K. Citron. Sexual privacy. *Yale LJ*, 128:1870, 2018.
- [11] C. Division. 18 u.s.c §§ 2257- 2257a certifications.
- [12] S. Duguay. Dressing up tinderella: Interrogating authenticity claims on the mobile dating app tinder. *Information, communication & society*, 20(3):351–367, 2017.
- [13] eSafety Commissioner. What you can report to esafety.
- [14] G. A. Gabison. White label: The technological illusion of competition. *The Antitrust Bulletin*, 67(4):642–662, 2022.
- [15] J. Geary. Judges deciding if cut-and-paste photos equal child porn, Oct. 2010.
- [16] C. Gibson, V. Frost, K. Platt, W. Garcia, L. Vargas, S. Rampazzi, V. Bindschaedler, P. Traynor, and K. R. B. Butler. Analyzing the monetization ecosystem of stalkerware. In *Privacy Enhancing Technologies Symposium*, July 2022.
- [17] I. Golmgrein. A comprehensive overview of monetization strategies in creative industries. *International Journal of Latest Engineering and Management Research (IJLEMR)*, 8, April 2023.
- [18] N. Henry, C. McGlynn, A. Flynn, K. Johnson, A. Powell, and A. J. Scott. *Image-based sexual abuse: A study on the causes and consequences of non-consensual nude or sexual imagery*. Routledge, 2020.
- [19] J. Hoggard. Face swap in photoshop leads to child pornography arrest, Mar. 2022.
- [20] Z. Huang, R. Khan, et al. A review of 3d human body pose estimation and mesh recovery. *Digital Signal Processing*, 128:103628, 2022.
- [21] M. H. Jarrahi, W. Sutherland, S. B. Nelson, and S. Sawyer. Platformic management, boundary resources for gig work, and worker autonomy. *Computer supported cooperative work (CSCW)*, 29:153–189, 2020.
- [22] A. P. Jennifer Klein. A call to action to combat image-based sexual abuse. Technical report, White House, 2024.
- [23] C. Kraft. Trolls used her face to make fake porn. There was nothing she could do. *New York Times*, July 2024.
- [24] B. Light, J. Burgess, and S. Duguay. The walkthrough method: An approach to the study of apps. *New media & society*, 20(3):881–900, 2018.
- [25] H. Luo, M. Wang, P. K.-Y. Wong, and J. C. Cheng. Full body pose estimation of construction equipment using computer vision and deep learning techniques. *Automation in construction*, 110:103016, 2020.
- [26] R. MacKinnon. Virtual rape. *Journal of Computer-Mediated Communication*, 2(4):JCMC247, 1997.
- [27] S. Maddocks. ‘A deepfake porn plot intended to silence me’: Exploring continuities between pornographic and ‘political’ deep fakes. *Porn Studies*, 7(4):415–423, 2020.
- [28] E. Maiberg. Apple’s huge “dual use” face swap app problem is not going away, Aug. 2024.

- [29] E. Maiberg. Instagram advertises nonconsensual ai nude apps, Apr. 2024.
- [30] N. Marchal, R. Xu, R. Elasmr, I. Gabriel, B. Goldberg, and W. Isaac. Generative AI misuse: A taxonomy of tactics and insights from real-world data. *arXiv preprint arXiv:2406.13843*, 2024.
- [31] T. Mattinen, J. Macey, and J. Hamari. A ruse by any other name: Comparing loot boxes and collectible card games using magic arena. *Proceedings of the ACM on Human-Computer Interaction*, 7(CHI PLAY):721–747, 2023.
- [32] C. McGlynn, E. Rackley, and R. Houghton. Beyond ‘revenge porn’: The continuum of image-based sexual abuse. *Feminist legal studies*, 25:25–46, 2017.
- [33] S. McLean, S. Paxton, E. Wertheim, and J. Masters. Photoshopping the selfie: Self photo editing and photo investment are associated with body dissatisfaction in adolescent girls. *The International Journal of eating disorders*, 48, Aug. 2015.
- [34] neurips. Ethics guidelines.
- [35] Y. Nirkin, I. Masi, A. Tran Tuan, T. Hassner, and G. Medioni. On face segmentation, face swapping, and face perception. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, page 98–105. IEEE Press, 2018.
- [36] N. I. Observatory. FAQ | The National Internet Observatory.
- [37] U. M. of Justice and L. Farris. Government cracks down on ‘deepfakes’ creation. <https://www.gov.uk/government/news/government-cracks-down-on-deepfakes-creation>, 2024.
- [38] I. of Medicine (US) Committee on Assessing Genetic Risks. Social, legal, and ethical implications of genetic testing, Jan. 1994.
- [39] D. O’Neill, M. V. Birk, and R. L. Mandryk. Unpacking norms, narratives, and nourishment: A feminist hci critique on food tracking technologies. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–20, 2024.
- [40] OnlyFans. Age & identity verification.
- [41] H. J. Parkinson. Stop calling women ‘girls’. it’s either patronising or sexually suggestive. *The Guardian*, July 2015.
- [42] F. Petitcolas, R. Anderson, and M. Kuhn. Information hiding: A survey. *Proceedings of the IEEE*, 87(7):1062–1078, 1999.
- [43] F. Pierri and S. Ceri. False news on social media: A data-driven survey. *SIGMOD Rec.*, 48(2):18–27, Dec. 2019.
- [44] PornHub. Model agreement – PornHub help.
- [45] L. Reime, V. Tsaknaki, and M. L. Cohn. Walking through normativities of reproductive bodies: A method for critical analysis of tracking applications. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2023.
- [46] M. K. Scheuerman, K. Spiel, O. L. Haimson, F. Hamidi, and S. M. Branham. HCI guidelines for gender equity and inclusivity. 2020.
- [47] C. Senate. Cantwell, Blackburn, Heinrich Introduce Legislation to Increase Transparency, Combat AI Deepfakes & Put Journalists, Artists & Songwriters Back in Control of Their Content, Jul 2024.
- [48] Z. Stardust, D. Blunt, G. Garcia, L. Lee, K. D’Adamo, and R. Kuo. High risk hustling: Payment processors, sexual proxies, and discrimination by design. *CUNY L. Rev.*, 26:57, 2023.
- [49] R. Su. Here’s a map showing which US states have passed laws against revenge porn — and those where it’s still legal. *Business Insider*, Oct. 2019.
- [50] K. Tenbarage. Nude deepfake images of Taylor Swift went viral on X, evading moderation and sparking outrage. *NBC News*, Jan. 2024.
- [51] B. Timmerman, P. Mehta, P. Deb, K. Gallagher, B. Dolan-Gavitt, S. Garg, and R. Greenstadt. Studying the online deepfake community.
- [52] X. Tong, L. Wang, X. Pan, and J. Wang. An Overview of Deepfake: The Sword of Damocles in AI. pages 265–273, July 2020.
- [53] C. A. Uhl, K. J. Rhyner, C. A. Terrance, and N. R. Lugo. An examination of nonconsensual pornography websites. *Feminism & Psychology*, 28(1):50–68, 2018.
- [54] R. Umbach, N. Henry, G. F. Beard, and C. M. Berryessa. Non-consensual synthetic intimate imagery: Prevalence, attitudes, and knowledge in 10 countries. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI ’24, New York, NY, USA, 2024. Association for Computing Machinery.

- [55] G. Wang, J. Zhao, M. Van Kleek, and N. Shadbolt. Protection or punishment? relating the design space of parental control apps and perceptions about them to support parenting for online safety. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2):1–26, 2021.
- [56] M. Wei, S. Consolvo, P. G. Kelley, T. Kohno, T. Matthews, S. Meiklejohn, F. Roesner, R. Shelby, K. Thomas, and R. Umbach. Understanding Help-Seeking and Help-Giving on social media for Image-Based sexual abuse. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 4391–4408, Philadelphia, PA, Aug. 2024. USENIX Association.
- [57] C. Wharton, J. Rieman, C. Lewis, and P. Polson. The cognitive walkthrough method: A practitioner’s guide. In *Usability inspection methods*, pages 105–140. 1994.
- [58] A. Whitten and J. D. Tygar. Why Johnny can’t encrypt: A usability evaluation of pgp 5.0. In *USENIX security symposium*, volume 348, pages 169–184, 1999.
- [59] Y. Wu, Z. Hu, J. Guo, H. Zhang, and H. Huang. A resilient and accessible distribution-preserving watermark for large language models. 2024.
- [60] N. Yee and J. Bailenson. The proteus effect: The effect of transformed self-representation on behavior. *Human Communication Research*, 33:271 – 290, July 2007.
- [61] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5505–5514, 2018.

A Appendix A: Age Confirmation

These applications create sexual content that may not be appropriate for younger ages. During Section 4 we discuss when a user is asked by the application to confirm their age. To visually represent this, we show a flow diagram of when all 20 applications ask a user to confirm their age. This "confirmation" is only a button that they have to click to continue through the website.

B Appendix B: Imagery Present

In Section 4, we discuss the types of imagery the user may come across during their entry walkthrough: women’s bodies both clothed, nude, in sexual acts. However, the further flesh out the spread of imagery, we created a bar graph to show what a user may view on the homepage.

C Appendix C: Pricing breakdown

To further discuss and lend intuition into how these applications are being funded, users can visually see how the payment processor breakdown looks, and the prices per tier that users pay as well as the average price-per-image on each application.




Figure 3: 14 out of the 20 websites will ask the user if they are 18 years old or older; however, when this occurs varies. 7 websites ask upon loading up the website; 2 ask when the user registers/logs in; 1 asks when users attempt to view the gallery; 3 ask before purchase; and 1 asks at generation. These periods reveal a variety of nudity to the user before their age is asked, and in the case of 6 applications, they are able to generate images without being asked their age.




Figure 4: All 20 websites specialize in AI Undressing Tools for the female form; however, the levels of explicit imagery visible on the homepage of each application vary between websites. This can range from blurred naked “undressing” photos to celebrities being posed naked in sexual positions.

U = Unlimited

Website	Scaling	Tier 1		Tier 2		Tier 3		Tier 4		Tier 5		Avg Cost per Image
		Cost	Credits	Cost	Credits	Cost	Credits	Cost	Credits	Cost	Credits	
S01	Feature	\$11	15	\$24	90	\$49	600					\$0.06
S02	Feature	\$24	15	\$38	90	\$60	600					\$0.70
S03	All For One	\$10	50									\$0.20
S04	Pro Credit	\$15	50	\$30	120	\$100	500					\$0.25
S05	Pro Credit	\$10	50	\$20	500	\$40	U					\$0.12
S06	Pro Credit	\$10	10	\$20	30	\$50	100					\$0.72
S07	Feature	\$15	900	\$18	1,200							\$0.02
S08	Pro Credit	\$2	3	\$6	15	\$15	45					\$0.47
S09	Pro Credit	\$10	30	\$29	120	\$49	160	\$99	600	\$299	3,000	\$0.23
S10	Feature	\$11	15	\$24	90	\$49	600					\$0.36
S11	Feature	\$8	15	\$20	90	\$35	900					\$0.26
S12	Feature	\$10	120	\$68	990							\$0.08
S13	Feature	\$30	100	\$37	110	\$99	420	\$85	270			\$0.30
S14	Pro Credit	\$25	100	\$55	300							\$0.22
S15	All For One	\$20	20									\$1.00
S16	All For One	\$13	100									\$0.13
S17	Pro Credit	\$1	5	\$10	100	\$40	500					\$0.13
S18	-	\$20	1									\$20.00
S19	Feature	\$30	500									\$0.06
S20	All For One	\$30	500									\$0.06

Figure 5: This ecosystem offers a variety of price points. On average, AI nudification generations cost between \$1.00 and \$0.06 depending on the application that the user bought from and the subscription tier that they purchased.




Figure 6: AI Nudification Platforms offer a number of payment processors to complete transactions with them and allow users to use their AI Nudification Tools. The most popular is cyptocurrency, followed by paypal and creditcards. However, some websites allow for apple pay, paetron, alipay, venmo, stream skins, phonepe, amazon pay, google pay, or cash app to complete the transaction.