

A Design Methodology for Incorporating Privacy Preservation into AI Systems

Jiehuang Zhang^{1,2} and Han Yu²,

¹School of Computer Science and Engineering, Nanyang Technological University

²Alibaba-NTU Singapore Joint Research Institute

{jiehuang001, han.yu}@ntu.edu.sg

Abstract

Artificial Intelligence (AI) has brought about paradigm shifts in how human societies work and live, from automating business processes to self-driving vehicles. With such tremendous impact comes ethical concerns. The privacy issue has been thrust to the forefront after multiple incidents involving well-known industry players. Although privacy preservation has been identified as a critical component of ethical AI, there is currently an absence of methodological tools to enable AI software teams to systematically surface and address privacy issues in the design and conception phase. We propose the Privacy-based Design (PbD) methodology to integrate privacy values early in the life cycle of AI products and services to address this gap. It allows AI software design teams to identify and analyse complex privacy issues in a systematic manner by guiding the envisioning of various scenarios. With PbD, we aim to reduce the barrier to entry, time and experience needed for AI practitioners to critically make well-thought-out decisions on incorporating privacy-preserving designs into AI solutions. User studies involving 29 participants found that the PbD methodology is useful and easy to use.

1 Introduction

Artificial Intelligence(AI) is a core part of the fourth industrial revolution [Schwab, 2017] and the digital age. It has enabled many advances in a plethora of fields such as healthcare [Tjoa and Guan, 2020], algorithmic crowdsourcing [Yu *et al.*, 2017] and autonomous driving [Zhang *et al.*, 2021]. The success of these AI technologies was made possible by the availability of big data generated in recent years, as well as novel machine learning techniques to achieve remarkable levels of performance. As new techniques facilitate the increasing impact of AI on everyday life [Makridakis, 2017], these technological breakthroughs allow the automation of tasks that brings many benefits.

Due to the far reaching impact of AI, there is a need to consider how it can adversely affect our societies due to potential ethical issues. Because of the improved capabilities of

AI, many businesses and organizations are transferring more responsibility and autonomy to algorithmic systems. As a result, the possibility of mistakes or unintended side effects are more likely to happen [Amodei *et al.*, 2016]. For example, a significant event that brought privacy into the spotlight was Facebook’s data breach in 2018 [Financial times, 2020] when the personal data of fifty million American voters from Facebook was gathered and then allegedly used by the political consultancy Cambridge Analytica. The incident raised questions on how giant technology companies can do more to protect their users’ interests. To ensure that AI development benefits humanity as a whole, we must collectively monitor its advancement and guide its trajectory towards human-centred and ethical AI solution design [Croeser and Eckersley, 2019; Yu *et al.*, 2018].

Privacy preservation research in AI aims to address the question “What are the privacy challenges in Machine Learning (ML) and how can we solve them?” [Liu *et al.*, 2021]. Federated learning [Yang *et al.*, 2019b] is the primary approach adopted by this research field. Multiple survey papers have provided overviews of the state of this field from diverse perspectives [Liu *et al.*, 2021; Tan *et al.*, 2022; Lyu *et al.*, 2022; Zhang and Yu, 2022b; Shi *et al.*, 2023]. As more groundbreaking techniques to achieve privacy are discovered, we must also consider the need of integrating these principles early in the early stages of the AI software life cycle. However, the following challenges hinder design teams to incorporate considerations for privacy preservation into their AI solutions during the conceptualization phase:

- Diverse Privacy Notions and Privacy Preservation Techniques:** Privacy is a complex and multifaceted concept. It can have different definitions and requirements in different application scenarios. Furthermore, when privacy is prioritized, its requirements may impact other metrics such as performance and accuracy. Hence, AI solution design teams who are not well trained on this topic might be overwhelmed when trying to grasp the different notions and techniques during the AI product and service life cycle.
- Diverse Groups with Different Interests and Agendas in Various Domains:** To different stakeholders, different AI application scenarios may require different notions of privacy to be prioritized. This complexity poses

85 significant challenges to AI solution teams to allocate
86 their limited resources to fulfil such requirements.

87 To address these challenges, we propose the Privacy-based
88 Design (PbD) methodological framework. It is an extension
89 of our previously proposed design methodologies for incor-
90 porating fairness [Shu *et al.*, 2021] and explainability [Zhang
91 and Yu, 2022a] considerations into AI solutions. The objec-
92 tive of PbD aims to facilitate software teams to systematically
93 analyse privacy issues during the conceptualization and brain-
94 storming phase, by lowering the barrier to entry and eliciting
95 systematic and deep thinking during team discussions. As an
96 important part of constructing ethical AI products and ser-
97 vices, the goal of the methodology is to create scaffolding
98 for conversations among scientists and engineers, thereby en-
99 abling them to reach an optimal solution for privacy preser-
100 vation in their AI solution designs. This is achieved by fac-
101 ilitating the process of brainstorming and investigating pri-
102 vacy requirements and topics surrounding the application do-
103 main and stimulating deep thinking from the shoes of various
104 stakeholder communities. Through preliminary user studies
105 involving 29 participants, we demonstrate that the proposed
106 methodology is useful and easy to use.

107 2 Related Work

108 Privacy preservation is a vital part of making AI safe and
109 beneficial for all. There is increasing public awareness about
110 large companies compromising on data security and user pri-
111 vacy. There has been much backlash in response to these
112 scandals, and many countries are improving their laws to ad-
113 dress data privacy and security [Yang *et al.*, 2019a]. For ex-
114 ample, the European Union (EU) instituted the General Data
115 Protection Regulation (GDPR) in order to enhance the pro-
116 tection of public users’ personal privacy and security [Euro-
117 peanUnion, 2016].

118 Newly emerging techniques in AI and machine learning
119 (ML) continue to increase the intricacy of privacy preserva-
120 tion. The challenges and problems associated with making AI
121 privacy respecting have been a central focus of the research
122 community, given the time-sensitive nature of the problem
123 before more government and regulatory laws are introduced
124 to protect data. Such works can be further grouped into sub-
125 categories, depending on whether the techniques are meant
126 for dataset or model protection, as well as whether ML tech-
127 niques are used for offence or defence.

128 Most responsible AI methodologies and frameworks are
129 influenced by the Value Sensitive Design (VSD) approach
130 [Friedman *et al.*, 2017], which was developed in human-
131 computer interaction (HCI) information systems design
132 (ISD). VSD gives importance to the ethical values of both
133 direct and indirect stakeholders and uses various methods to
134 engage with diverse values based on the application. This
135 allows designers to gain insights and integrate with other
136 methodologies. The main workflow of VSD involves stim-
137 ulating the perspectives of stakeholders and analyzing how
138 their values are affected. Direct stakeholders directly use the
139 AI product and are impacted, while indirect stakeholders are
140 not users but are still affected.

141 VSD has delivered two exploratory card games, Judgement

Call [Ballard *et al.*, 2019] and Envisioning Cards, to facil- 142
itate ethical AI design. Envisioning Cards encourage criti- 143
cal thinking about stakeholders, time, values, and motivation 144
to consider systemic long-term problems. Judgement Call 145
is a turn-based card game that AI development teams can 146
use to identify moral problems in an AI product, using cards 147
that focus on virtuous value, stakeholders, and review ratings 148
to encourage experimental thinking. Based on the Judge- 149
ment Call game design, the Fairness in Design (FID) [Shu 150
et al., 2021] and Explainability in Design (EID) [Zhang and 151
Yu, 2022a] approaches have been proposed to provide more 152
focused guidance for AI design teams on envisioning chal- 153
lenges and opportunities with regard to fairness and explain- 154
ability, respectively. The proposed PbD approach extends 155
FID and EID to provide support for incorporating privacy- 156
preservation into AI solution designs. 157

158 3 Preliminaries

159 For this section, we have classified the techniques of pri- 159
vacy into the respective four categories as shown in Figure 160
1: 1) Attack and Threat Models, 2) Private Machine Learn- 161
ing Schemes, 3) Privacy Attacks, and 4) Machine Learning- 162
enhanced Privacy Protection. 163

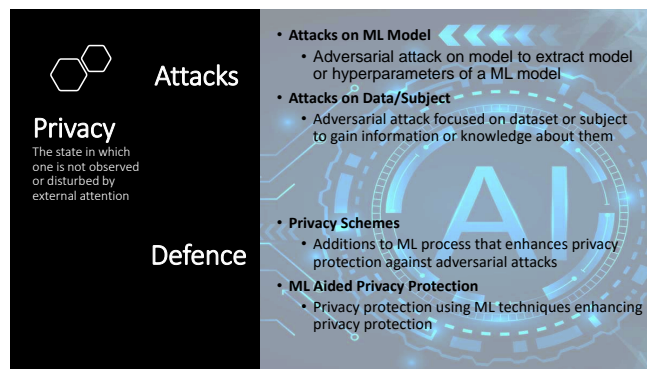


Figure 1: An overview of the main principles of privacy in AI

1. **Attacks on ML Model:** Adversarial attack on ML mod- 164
els to extract entire model or hyperparameters of an ML 165
model. Examples include: 166
 - (a) Model Extraction Attacks: Attacks aims to copy or 167
"extract" an AI model on a high level, resulting in 168
a function with parameters and coefficients that re- 169
sembles the original model [Tramèr *et al.*, 2016] 170
 - (b) Feature Estimation Attacks: Feature estimation at- 171
tacks seek to estimate specific features or statisti- 172
cal properties of the training dataset. This type of 173
attack is initiated through model inversion, power 174
side-channel attacks or shadow model [Fredrikson 175
et al., 2015] 176
 - (c) Membership Inference Attacks: Such an attack in- 177
volves determining whether a particular data point 178
belongs to the training dataset [Shokri *et al.*, 2017] 179
 - (d) Model Memorization Attacks: Such an attack seeks 180
to recover exact feature values on individual sam- 181

182 ples and involves stealing model parameters and
183 coefficient values [Song *et al.*, 2017]

184 2. **Privacy Schemes:** A privacy-preserving scheme is a
185 collection of techniques or algorithms that assist ML
186 models to improve their defence against adversarial pri-
187 vacy attacks.

188 (a) Encryption: Homomorphic Encryption applies a
189 computation to encrypt data, allowing sensitive
190 data to be used as a training dataset. However adds
191 an order of magnitude to computation complexity
192 [Bost *et al.*, 2014]

193 (b) Obfuscation: Obfuscation mechanisms aim to re-
194 duce the precision of privacy attacks using the in-
195 troduction of noise to the coefficients of the model
196 [McPherson *et al.*, 2016]

197 (c) Aggregation: Aggregation techniques involve mul-
198 tiple parties joining an ML scheme while at the
199 same time aiming to hide their own datasets, to be
200 applied during training or after training. Federated
201 Learning (FL) is grouped in this section

202 3. **Attack and Threat Models:** Types of attack an adver-
203 sary can employ to access a model’s parameters.

204 (a) Identification Attack: Specifies a user’s details or
205 identification on a shared dataset [Li *et al.*, 2016],
206 when anonymization is reversed, the attack is called
207 re-identification

208 (b) Inference Attack: This type of attack’s objective
209 is to explore data to obtain information on a target
210 [Nasr *et al.*, 2019]

211 (c) Linkage Attack: The counter party’s goal is to steal
212 the subject’s information by comparing or cross-
213 referencing different datasets from the source

214 4. **ML Privacy Protection:** Preemptive privacy measures
215 targeted at mitigating privacy risks

216 (a) Risk Assessment Protection: Evaluate and pre-
217 dict the risk for users during the process of access-
218 ing and sharing information. Algorithms are em-
219 ployed to predict data streams to find risks and sub-
220 sequently deploy countermeasures

221 (b) Personal Privacy Management: Policy evaluation,
222 user preference prediction, and management, of the
223 behaviour of the user

224 (c) Private Data Release: Disseminate datasets with a
225 privacy assurance

226 4 The Privacy by Design Methodology

227 The Privacy by Design (PbD) Methodology takes the form
228 of a tangible card game, which can be utilized by individu-
229 als with varying levels of expertise, ranging from novices to
230 professionals. The purpose of this methodology is to encour-
231 age brainstorming and the identification of potential privacy
232 issues within AI while remaining adaptable to different appli-
233 cations. The design team is given the freedom to dictate the
234 dimensions of their intended domain application, making the
235 methodology application agnostic.

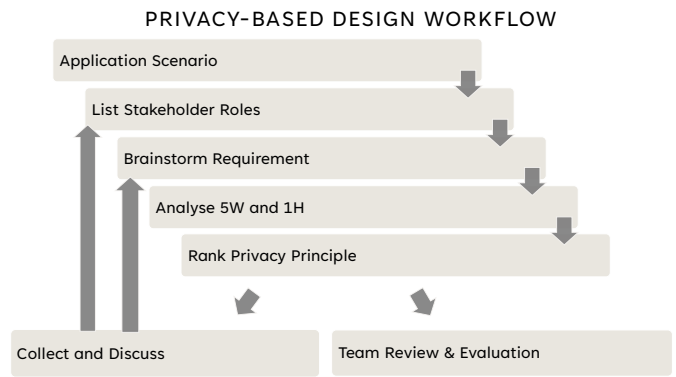


Figure 2: The Privacy-Based Design Workflow

236 The step-by-step framework for a software design team to
237 employ PbD to encourage conversations in the software team
238 surrounding privacy issues is illustrated in Figure 2. The sys-
239 temic guide to using the methodology is discussed below:

- 240 1. Initially, the AI design team is required to select an ap-
241 plication domain, which will provide the base environ-
242 ment for the purpose of the study. The environment, de-
243 pending on the teams, can either be a genuine or imag-
244 inary setting, with a strong preference for a domain
245 where privacy is a significant concern. It is ideal for
246 team members to possess knowledge about the domain,
247 enabling them to incorporate as many details as possi-
248 ble throughout the user study. This is particularly im-
249 portant as some domains may require specific consid-
250 erations and compromises that can influence the usage of
251 this methodological framework.
- 252 2. During Step Two, users are required to select a card
253 from the classification system that corresponds to their
254 context domain. The classification system used for this
255 purpose is based on Shneiderman’s work on usability
256 motivation in the field of Human-AI Interaction (HCI)
257 [Shneiderman and Hochheiser, 2001].
- 258 3. In Step Three of the PbD methodology, the group needs
259 to conduct an exploratory analysis and recognize the
260 stakeholders who play a vital role in the end-to-end AI
261 pipeline. Direct stakeholders are those who frequently
262 use the product or service, while indirect stakeholders
263 are not the end-users but are still influenced by the de-
264 ployment [Friedman *et al.*, 2017]. The team members
265 are required to take on the perspective of a stakeholder
266 and carry out an in-depth examination of the privacy de-
267 tails concerning that stakeholder. The PbD methodol-
268 ogy presents several guiding questions to streamline the
269 thought process. These critical thinking guides revolve
270 around the who, what, when, where, why, and how of
271 privacy-related topics. For instance:

272 After finishing all the steps in the methodology, the team
273 can have an insightful and deeper appreciation of privacy
274 concerns and their application domain. They can choose to
275 further investigate a specific topic by brainstorming and dis-
276 cussing it in detail as a team. This process of delving deeper

277 into a particular topic can help uncover any complex privacy
 278 issues that may have been missed otherwise.

279 The deliverables of the framework include the listed out-
 280 puts:

- 281 1. Understanding and Selecting the privacy principles that
 282 are relevant for the environment.
- 283 2. Rank priorities of specialized requirements for the anal-
 284 ysis of privacy measures.
- 285 3. Quantifying improvements in the privacy knowledge
 286 levels of methodology users.
- 287 4. Create a thinking guide to determine where to focus their
 288 attention and focus on during sprints

289 The methodology encourages a collaborative approach to
 290 AI design, with individuals stimulating the perspectives of
 291 direct and indirect stakeholders, and then conducting an in-
 292 depth exploration of privacy attributes. This approach helps
 293 ensure that all potential privacy concerns are addressed and
 294 that stakeholders' perspectives are taken into account in the
 295 design process. Using these outputs, the team can then use the
 296 insights and information gained to make informed decisions
 297 about improving their processes.

298 5 Empirical Evaluation

299 In this section, we discuss the process and analyse the results
 300 of our experiments with recruited participants to evaluate the
 301 proposed PbD Methodology and our hypotheses.

302 5.1 Study Design

303 We recruited 29 participants through the snowball sampling
 304 method. Our criteria for the qualification of the potential par-
 305 ticipants are such that they possessed experience being part
 306 of a team that worked on AI technologies. All of the partic-
 307 ipants were researchers, scientists or engineers that were able
 308 to understand privacy concepts in AI/ML, as well as consent-
 309 ing to be recorded and their insights published. We recruited
 310 participants from a diverse age range to investigate how the
 311 PbD Framework can affect end users of different levels of se-
 312 niority. However, the majority of participants belong to the
 313 20-30-year-old age group, as is the profile of the usual pro-
 314 posed users of the PbD framework.




Figure 3: Demographics of the Participants.

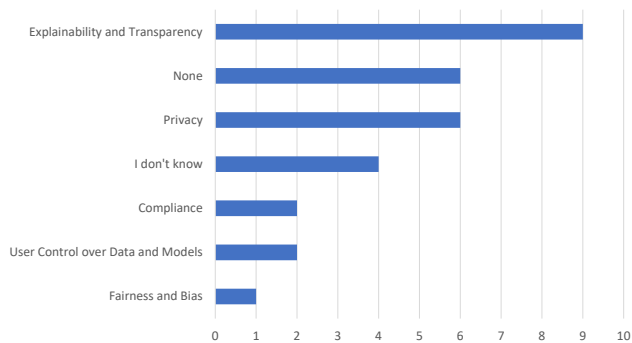


Figure 4: Participants' ethical AI prioritisation.

315 Before commencing the user study, we enquired about the
 316 order of priority of the type of responsible AI considerations
 317 in the AI product pipeline. According to Figure 5, nine users
 318 preferred explainability and transparency as their top consid-
 319 erations, while a significant portion of 6 participants indicates
 320 that none of the ethical values as part of the considerations for
 321 the software development pipeline. This observation was re-
 322 flected in the many feedback from the users that performance
 323 and efficiency were greatly valued over ethical AI principles.
 324 Most of the initiatives in ethical AI tend to be a reaction
 325 to government compliance processes or regulatory pressure.
 326 While privacy may not be at the top of the priority list, there
 327 is a need for more toolkits to assist software design teams to
 328 enhance privacy in their workflows.




Figure 5: Participants' Application Domains.

329 We also asked participants about their application domain
 330 which will form the context of the user study. According to
 331 Figure 6, most of them are working in the healthcare, gen-
 332 eral machine learning and media and entertainment sectors.
 333 The application domain will impact the decisions made while
 334 navigating the process of the PbD methodology. As a mea-
 335 sure to improve the consistency and validation of the ques-
 336 tionnaire results, we included a redundancy test by asking the
 337 same question twice. Based on the redundancy check, we
 338 identify and discard invalid responses. Furthermore, the par-
 339 ticipants were instructed to complete the post-study question-
 340 naire immediately after the user study, and most participants
 341 completed it on the same day as the user study.

- 342 We designed the questionnaire based on the 3 hypotheses:
- 343 1. PbD assists users to select the privacy concept that is
 - 344 appropriate for their applications.
 - 345 2. PbD improves participants' ability to identify privacy
 - 346 concerns in their AI applications.
 - 347 3. PbD helps users to stimulate the perspectives of different
 - 348 stakeholders.

349 Both the pre-study and post-study questionnaires consist of
 350 a main section where participants conduct a self-assessment
 351 of their understanding and ability to apply privacy concepts to
 352 their AI products and services. Each hypothesis is designed
 353 after exploring the literature on advances in the field of AI
 354 privacy and designed to rate the participant's individual abil-
 355 ity to brainstorm and surface privacy issues, design relevant
 356 and optimal strategies, as well as to stimulate the perspec-
 357 tives of stakeholders. They had to give themselves a score of
 358 their understanding of AI privacy issues on a Likert Scale of
 359 1 to 5, 1 being "strongly disagree" (SA) and 5 being "strongly
 360 agree" (SA). We conducted data analysis and hypothesis test-
 361 ing based on the results of the self-assessment questionnaire.

362 6 Results and Analysis

363 6.1 Hypothesis 1

364 Hypothesis 1: PbD assists users to select the privacy concept
 365 that is appropriate for their applications.




Figure 6: Participants' self-reported capability of making design decisions related to privacy before and after using PbD.

366 According to figure 6, it can be observed that the responses
 367 were mainly negative or neutral and can be said to follow
 368 a distribution roughly centred on "Disagree". This indicates
 369 that many of the participants are not so confident in their abil-
 370 ities to select the optimal privacy principle for their applica-
 371 tion domain. Since privacy is not a significant concern in
 372 many AI software teams, we expected that many participants
 373 are not well-versed in this area and require assistance in do-
 374 ing so. The findings also indicate that the distribution of the
 375 participants' self-assessed abilities to make relevant decisions
 376 pertaining to privacy is representative of a typical population
 377 of AI solution designers. After the participants proceeded
 378 through the PbD methodological tool, there was a significant
 379 increase in the number of participants who responded with

'Agree' and 'Strongly Agree', while the corresponding num-
 380 ber of responses with 'Disagree' decreased significantly. This
 381 observation showed that the participants perceived the PbD
 382 methodology to be effective in enabling them to think criti-
 383 cally about the privacy criteria that are relevant and optimal
 384 for their application scenarios.
 385

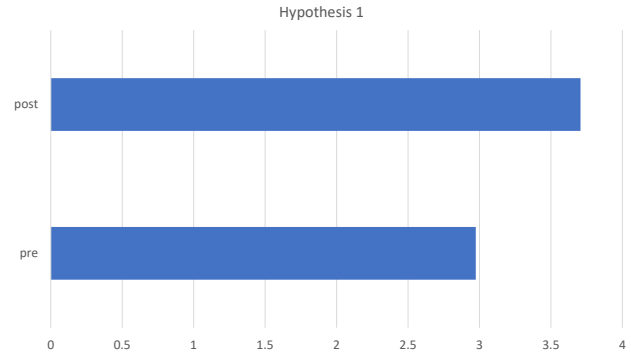


Figure 7: Participants' average scoring for the pre-and post-studies for hypothesis 1

386 According to figure 7, the average response score of the
 387 participants in the post-study was significantly higher than
 388 those in the pre-study, an increase of more than 0.6. After
 389 conducting statistical analysis and on the basis of the stu-
 390 dent's t-test of questionnaire results from H1, we concluded
 391 that the null hypothesis can be rejected at a 95 percent confi-
 392 dence interval with Cronbach alpha at 0.7462.

393 6.2 Hypothesis 2

394 Hypothesis 2: PbD improves participants' ability to identify
 395 privacy concerns in their AI applications.




Figure 8: Participants' self-reported capability of surfacing privacy concerns before and after using PbD.

396 According to figure 8, the results provide an overview of
 397 the participants' responses to surfacing or identifying privacy
 398 concerns in the AI pipeline ahead of time focusing on hy-
 399 pothesis 2. This is a useful skill that enables early detection
 400 of problems that can become exacerbated in the later stages
 401 of development. Similar to the previous observations, the
 402 pre-study responses were roughly centred on 'Agree', with




Figure 9: Participants' average scoring for the pre-and post-studies for hypothesis 2.




Figure 11: Participants' average scoring for the pre-and post-studies for hypothesis 3.

403 a slight increase in the number of responses of 'Agree' and
 404 'Strongly Agree' post-study. This observation might be re-
 405 flective of participants' relative confidence in being able to
 406 detect privacy issues during the development process.
 407 For hypothesis 2, we found that the average questionnaire
 408 response increased by more than 0.5 in the post-study com-
 409 pared with that in the pre-study, according to 10. After con-
 410 ducting a student's t-test, we were only able to reject the null
 411 hypothesis at a 90 percent confidence level with the Cronbach
 412 alpha at 0.7156.

413 6.3 Hypothesis 3

414 Hypothesis 3: PbD helps users to stimulate the perspectives
 415 of different stakeholders.




Figure 10: Participants' self-reported capability of stimulating stakeholders' perspective before and after using PbD.

416 Figure 10 illustrates the overview of participants' re-
 417 sponses on stimulating the perspectives of various stakehold-
 418 ers, both direct and indirect. According to figure 10, there was
 419 a significant increase in the number of responses of 'Strongly
 420 Agree', from 1 to 10 after the use study. We found that
 421 the methodology actively facilitates critical thinking and en-
 422 ables participants to filter through irrelevant information to
 423 find issues that stakeholders are concerned with. Indirect
 424 stakeholders are usually overlooked in most development sce-
 425 narios, and over the course of the user study, we constantly
 426 asked questions on addressing the needs of these commu-

nities that may not be directly impacted by the use of tech-
 nology. Hence, the methodology significantly improved the
 participants' self-assessed ability to think critically from the
 perspective of stakeholders.

According to figure 11, the average of questionnaire re-
 sponses increased by about 0.6 in the post-study compared
 with that in the pre-study. After conducting a student's t-test,
 we were able to reject the null hypothesis at a 95 percent con-
 fidence level with the Cronbach Alpha at 0.7683.

436 7 Discussions and Limitations

437 Over the course of multiple user studies, we found that the
 438 context of the application domain greatly impacts how deci-
 439 sions were made when measures for ethical AI are re-
 440 quired. Most participants feedback that the PbD methodology
 441 is effective for promoting discussions and facilitating criti-
 442 cal thinking regarding the various issues surrounding privacy
 443 in artificial intelligence. Especially when ethical issues sur-
 444 rounding privacy and other aspects can be difficult to uncover,
 445 deep insights driven by comprehensive exploratory thinking
 446 can avoid unnecessary issues in the future. However, ulti-
 447 mately measures designed to enhance ethical AI systems can
 448 also be perceived as trade-offs for performance. User 3 shared
 449 that the reality of privacy in the priority list of most large or-
 450 ganisations:

451 *"Realistically, privacy is low on business priorities as it is*
 452 *usually only appreciated in hindsight. However, I do recom-*
 453 *mend employing this methodology nonetheless as this study*
 454 *gave me the opportunity to pause and reflect on areas of*
 455 *weaknesses in my business with regard to privacy manage-*
 456 *ment. Being able to plan ahead is critical for companies to*
 457 *stay ahead of the game in the long run."*

458 Although the issue of trade-off is a major hindrance in en-
 459 hancing privacy in AI systems, the vision of the research com-
 460 munity is that eventually, teams can deliver a system that min-
 461 imises the compromise on performance and ethical values.

462 Despite the methodological toolkit being a useful starting
 463 point for software or AI design teams with no experience
 464 in addressing ethical AI issues, the complexities and frag-
 465 mented state of the field can be a significant challenge to over-
 466 come, especially when mired in technical details. Participant
 467 7 noted that the methodology can be used to generate industry

468 best practices and guidelines to plan the trajectory of building
469 ethical algorithmic systems:

470 *"This methodology serves as a good starting point to*
471 *incorporate privacy principles and considerations into ML*
472 *projects. Further on, it can be helpful to provide a simpli-*
473 *fied view of the complications when dealing with ethics and*
474 *inspire best practices in the process of enhancing private AI*
475 *systems."*

476 Furthermore, there are many stakeholders involved in
477 building and deploying large-scale AI systems, and each
478 group of stakeholders with diverging interests add to the diffi-
479 culty of building privacy in AI. For the purpose of this study,
480 we summarised the existing AI privacy principles into 4 main
481 groups of techniques to encourage the exploratory thinking
482 process. With additional time and resources, the team will
483 be able to build a more nuanced and balanced view of many
484 aspects of privacy and ethical AI. Participant 19 provided a
485 glimpse into the future direction of methodological tools for
486 building ethical AI:

487 *"This tool can serve as a privacy framework to be used in*
488 *the life-cycle of AI products, from design and conception to*
489 *deployment. To enable the framework to be conducted effec-*
490 *tively, the researchers can aim to standardise or estimate the*
491 *requirements, time and resources needed for each step of the*
492 *AI product pipeline. In this way, teams can work with this*
493 *information to better achieve their key goals, objectives and*
494 *deliverables."*

495 Additionally, we found that for several participants, their
496 internal model of privacy in AI was changed after being in-
497 troduced to the methodology. Due to the lack of importance
498 placed on privacy and the more general ethical AI principles,
499 AI team members may not fully understand the implications
500 and procedures of building these measures. Participant 11
501 noted that:

502 *"After the discussion and deep dive into the methodology,*
503 *the concept of privacy is completely different from what I ex-*
504 *pected. Upon learning what encompasses privacy in AI/ML, I*
505 *believe it is now even more important to implement preventive*
506 *measures against breaches of privacy."*

507 After each user study, the team discussed potential ways
508 to improve the process and each step of the methodology. In
509 some application scenarios, the participants shared more in-
510 depth principles and techniques that they employed. While in
511 other groups, less relevant concepts were discarded and new
512 topics were introduced to further facilitate discussion. Partic-
513 ipant 27 commented on the requirements of these ethical AI
514 tool-kits:

515 *"The framework provides a clear structure to navigate the*
516 *complexities and challenges in privacy-sensitive application*
517 *environments. It is comprehensive, clear and easy to fol-*
518 *low. However, it can be a significant challenge to build*
519 *application-agnostic methodological frameworks, given the*
520 *dynamic requirements of each field. One possible direction*
521 *moving forward is to group similar application domains to-*
522 *gether and define the common objectives, timelines and spe-*
523 *cific deliverables to provide a systematic way to address this*
524 *important field of ethical AI and privacy"*.

525 The user study consisted of 30 participants, many of which
526 are experts in their field of AI/ML. However, to evaluate

the PbD framework more effectively, a larger-scale online
study is required. We propose to use a crowdsourcing tool
such as Amazon Mechanical Turk to recruit participants from
the public for this large-scale study. Furthermore, there are
reports that self-assessed preferences and abilities usually
do not align completely with participants' actual behaviours
[Zell and Krizan, 2014]. Whether the findings from current
and past work can value add to objectives in our methodolog-
ical tool is still an open research question. When future works
in AI privacy is implemented, dividing the investigation into
multiple sub-categories seems to be the right course of action.

8 Conclusions and Future Work

The authors of this paper introduced a new methodology
called PbD to help design teams tackle complex ethical
dilemmas related to privacy in the development of AI prod-
ucts and pipelines. They identified gaps in current ethical AI
design methodologies and developed PbD using the VSD the-
ory and recent studies. PbD is designed to be user-friendly
and time-efficient to facilitate its adoption by design teams.
The authors plan to conduct user studies to assess the effec-
tiveness of PbD in various application domains, such as bank-
ing, finance, autonomous vehicles[Atakishiyev *et al.*, 2021],
and medical diagnosis [Tjoa and Guan, 2020]. They also plan
to include project management functions to make PbD more
accessible online. By promoting the use of PbD, the authors
hope to improve the analysis of the ethical implications of AI
products.

Acknowledgments

This research is supported, in part, by the National Research
Foundation Singapore and DSO National Laboratories under
the AI Singapore Programme (AISG Award No: AISG2-RP-
2020-019); Alibaba Group through Alibaba Innovative Re-
search (AIR) Program and Alibaba-NTU Singapore Joint Re-
search Institute (JRI) (Alibaba-NTU-AIR2019B1), Nanyang
Technological University, Singapore; and the Nanyang Assis-
tant Professorship (NAP).

References

- [Amodei *et al.*, 2016] Dario Amodei, Chris Olah, Jacob
Steinhardt, Paul Christiano, John Schulman, and Dan
Mané. Concrete problems in ai safety. *arXiv preprint*
arXiv:1606.06565, 2016.
- [Atakishiyev *et al.*, 2021] Shahin Atakishiyev, Mohammad
Salameh, Hengshuai Yao, and Randy Goebel. Explainable
artificial intelligence for autonomous driving: A compre-
hensive overview and field guide for future research direc-
tions. *arXiv preprint arXiv:2112.11561*, 2021.
- [Ballard *et al.*, 2019] Stephanie Ballard, Karen M Chappell,
and Kristen Kennedy. Judgment call the game: Using
value sensitive design and design fiction to surface ethical
concerns related to technology. In *Proceedings of the 2019*
on Designing Interactive Systems Conference (DIS'19),
pages 421–433, 2019.

- [Bost *et al.*, 2014] Raphael Bost, Raluca Ada Popa, Stephen Tu, and Shafi Goldwasser. Machine learning classification over encrypted data. *Cryptology ePrint Archive*, 2014.
- [Croeser and Eckersley, 2019] Sky Croeser and Peter Eckersley. Theories of parenting and their application to artificial intelligence. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 423–428, 2019.
- [EuropeanUnion, 2016] EuropeanUnion. Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation). 2016.
- [Financial times, 2020] Financial times. Facebook Data Breach, 2020.
- [Fredrikson *et al.*, 2015] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pages 1322–1333, 2015.
- [Friedman *et al.*, 2017] Batya Friedman, David G Hendry, and Alan Borning. A survey of value sensitive design methods. *Foundations and Trends in Human-Computer Interaction*, 11(2):63–125, 2017.
- [Li *et al.*, 2016] Huaxin Li, Haojin Zhu, Suguo Du, Xiaohui Liang, and Xuemin Shen. Privacy leakage of location sharing in mobile social networks: Attacks and defense. *IEEE Transactions on Dependable and Secure Computing*, 15(4):646–660, 2016.
- [Liu *et al.*, 2021] Bo Liu, Ming Ding, Sina Shaham, Wenny Rahayu, Farhad Farokhi, and Zihuai Lin. When machine learning meets privacy: A survey and outlook. *ACM Computing Surveys (CSUR)*, 54(2):1–36, 2021.
- [Lyu *et al.*, 2022] Lingjuan Lyu, Han Yu, Xingjun Ma, Chen Chen, Lichao Sun, Jun Zhao, Qiang Yang, and Philip S. Yu. Privacy and robustness in federated learning: Attacks and defenses. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [Makridakis, 2017] Spyros Makridakis. The forthcoming artificial intelligence (ai) revolution: Its impact on society and firms. *Futures*, 90:46–60, 2017.
- [McPherson *et al.*, 2016] Richard McPherson, Reza Shokri, and Vitaly Shmatikov. Defeating image obfuscation with deep learning. *arXiv preprint arXiv:1609.00408*, 2016.
- [Nasr *et al.*, 2019] Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *2019 IEEE symposium on security and privacy (SP)*, pages 739–753, 2019.
- [Schwab, 2017] Klaus Schwab. *The fourth industrial revolution*. Currency, 2017.
- [Shi *et al.*, 2023] Yuxin Shi, Han Yu, and Cyril Leung. Towards fairness-aware federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [Shneiderman and Hochheiser, 2001] Ben Shneiderman and Harry Hochheiser. Universal usability as a stimulus to advanced interface design. *Behaviour & Information Technology*, 20(5):367–376, 2001.
- [Shokri *et al.*, 2017] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017.
- [Shu *et al.*, 2021] Ying Shu, Jiehuang Zhang, and Han Yu. Fairness in design: A tool for guidance in ethical artificial intelligence design. In *International Conference on Human-Computer Interaction*, pages 500–510. Springer, 2021.
- [Song *et al.*, 2017] Congzheng Song, Thomas Ristenpart, and Vitaly Shmatikov. Machine learning models that remember too much. In *Proceedings of the 2017 ACM SIGSAC Conference on computer and communications security*, pages 587–601, 2017.
- [Tan *et al.*, 2022] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [Tjoa and Guan, 2020] Erico Tjoa and Cuntai Guan. A survey on explainable artificial intelligence (xai): Toward medical xai. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [Tramèr *et al.*, 2016] Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction apis. In *USENIX security symposium*, volume 16, pages 601–618, 2016.
- [Yang *et al.*, 2019a] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.*, 10(2):12:1–12:19, January 2019.
- [Yang *et al.*, 2019b] Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, Tianjian Chen, and Han Yu. Federated learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 13(3):1–207, 2019.
- [Yu *et al.*, 2017] Han Yu, Chunyan Miao, Yiqiang Chen, Simon Fauvel, Xiaoming Li, and Victor R Lesser. Algorithmic management for improving collective productivity in crowdsourcing. *Scientific reports*, 7(1):1–11, 2017.
- [Yu *et al.*, 2018] Han Yu, Zhiqi Shen, Chunyan Miao, Cyril Leung, Victor R Lesser, and Qiang Yang. Building ethics into artificial intelligence. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI’18)*, pages 5527–5533, 2018.
- [Zell and Krizan, 2014] Ethan Zell and Zlatan Krizan. Do people have insight into their abilities? a metasynthesis. *Perspectives on Psychological Science*, 9(2):111–125, 2014.

- 688 [Zhang and Yu, 2022a] Jiehuang Zhang and Han Yu. A
689 methodological framework for facilitating explainable ai
690 design. In *International Conference on Human-Computer*
691 *Interaction*, pages 437–446. Springer, 2022.
- 692 [Zhang and Yu, 2022b] Yanci Zhang and Han Yu. Towards
693 verifiable federated learning. In *Proceedings of the 31st In-*
694 *ternational Joint Conference on Artificial Intelligence (IJ-*
695 *CAI'22)*, pages 5686–5693, 2022.
- 696 [Zhang *et al.*, 2021] Jiehuang Zhang, Ying Shu, and Han Yu.
697 Human-machine interaction for autonomous vehicles: A
698 review. In *International Conference on Human-Computer*
699 *Interaction*, pages 190–201. Springer, 2021.