Design Patterns for AI Curated Content

Khan, Awais Hameed*[a]; Wijenayake, Senuri[b]; Al Lawati, Sara Fahad Dawood[b]; Spina, Damiano[b]; Hettiachchi, Danula[b]

[a] The University of Queensland, Brisbane, Australia
[b] RMIT University, Melbourne, Australia

Advancements in artificial intelligence (AI) technologies have led to a rapid increase in the use of AI-enabled tools to curate content across various application contexts (e.g., news, social media, search engines, product reviews). This content can include multiple data sources and formats. Some commercially available AI-enabled content curation integrations seem useful, while others are underdeveloped, or not considered in terms of utility, contextual relevance, or user experience. We present an annotated portfolio of nine discrete interface design patterns, exploring how an AI-in-the-loop approach can be used to present contextually relevant, AI curated content — across varying degrees of AI involvement. To illustrate these patterns, we use a case study of online news content, to reflexively examine how different content types and use cases are suited to the different interface design patterns. We view this work as a provocation for advancing the discourse on AI-enabled content curation applications.

Keywords: Design Patterns; Artificial Intelligence; User Interface Design

1 Introduction

This pictorial introduces a catalogue of nine interface design patterns for AI-enabled content curation, catering to different levels of AI involvement, user engagement, and the complexities of content properties and context. As AI-in-the-loop automated content curation becomes more prevalent, these proposed patterns provide a blueprint for technology researchers, designers, and practitioners to explore various potential applications of AI-enabled content curation. To our knowledge, this is the first attempt to explicitly characterise the wide range of design alternatives for AI-enabled content curation that current interface design patterns have not covered. These interfaces act as useful provocations (see: 'provotypes' (Boer & Donovan, 2012)), and boundary objects (Star & Griesemer, 1989), enabling different stakeholder groups to develop a shared understanding of the vocabulary of the patterns to explore content curation possibilities. These patterns are closer to Alexander's original notion of architectural patterns (Alexander, 1977), modular building blocks that can be combined into different configurations. Since Alexander's original conception of a pattern language, design patterns have become popular in software engineering (Gamma et al., 1994). These design patterns are

interface-level concepts, distinct from software engineering patterns. We have opted to present these patterns in the form of an annotated portfolio (Gaver & Bowers, 2012), as an assemblage of the patterns, i.e., the actual interface design pattern concepts, coupled with high-fidelity renderings of the patterns in use, complemented by reflexive, textual annotations, drawing attention to particular salient features of these patterns (i.e. description, type of AI involved, requirement, solution, and context of use).

Effective content curation extends beyond simply filtering and ranking; it requires careful consideration of both the properties of the content, and the context in which it is presented. Content formats vary, with some consisting solely of text or video, while others incorporate interactive elements such as user inputs e.g., comments. Additionally, contextual factors such as urgency, play a crucial role in how content should be presented. For instance, with *breaking news*, users are primarily focused on staying up-to-date with the latest developments, hence it makes sense to highlight the most recent, important updates. In contrast, an *in-depth explainer on the causes of bushfires* to educate the public, might be better suited for a more considered pace, allowing users to investigate related information at their own leisure. Without appropriate presentation strategies, AI-enabled curation risks obscuring relevance, diminishing usability, or even distorting meaning. Thoughtful interface design is therefore essential to ensuring that curated content is structured in ways that align with both its characteristics and context.

2 Methodology & Positionality

The design patterns presented, were iteratively developed over multiple generative design workshops involving the authors, who are HCI, Design, and Information Retrieval researchers. The team drew upon their domain expertise about content design, AI-curation, and ongoing collaborations with news and media practitioners. We used exploratory, participatory, and generative design methods to identify user content needs, brainstorm AI-curation concepts, and explore human-AI cooperation futures within the news and media context. An initial workshop involved generating low-fidelity sketches (hand-drawn concepts and digital mock ups using *Adobe Illustrator*, *Figma*, and *Procreate*) of AI curated content interface concepts. Of these initial concepts, twelve were transformed by the first author into higher-fidelity mock-ups, preliminary designs embodying the broad range of ideas generated during the workshop. This was followed by a second workshop, where the team systematically reviewed the initial designs on *Figma*, removing duplicates and combining designs that shared similar characteristics, but lacked sufficient distinction to create more comprehensive patterns. The concepts were again refined, and a second iteration was evaluated in a collaborative workshop using *Miro*. Pattern concepts were examined for considerations relating to (1) the type of *content:* does the pattern cater to multiple formats of curated content? i.e., sequential or concurrent; static or dynamic; single or multimodal formats; (2) the impact on *user experience* i.e., what is the role of user and their engagement with the interface i.e., do users proactively or passively engage with the content? What agency do users have to append and modify the curated content?; and lastly (3) the level of *AI involvement* i.e., what is the role and behaviour of the AI in curating content? leading to the emergence of the different AI-involvement types identified (aggregation, exploration, recommendation and generation). Through critical evaluation of each design's unique contribution, this process resulted in finalising nine distinct interface design patterns which were then refined and consolidated into the final AI curated content design pattern portfolio (See: Figure 1).

3 Portfolio

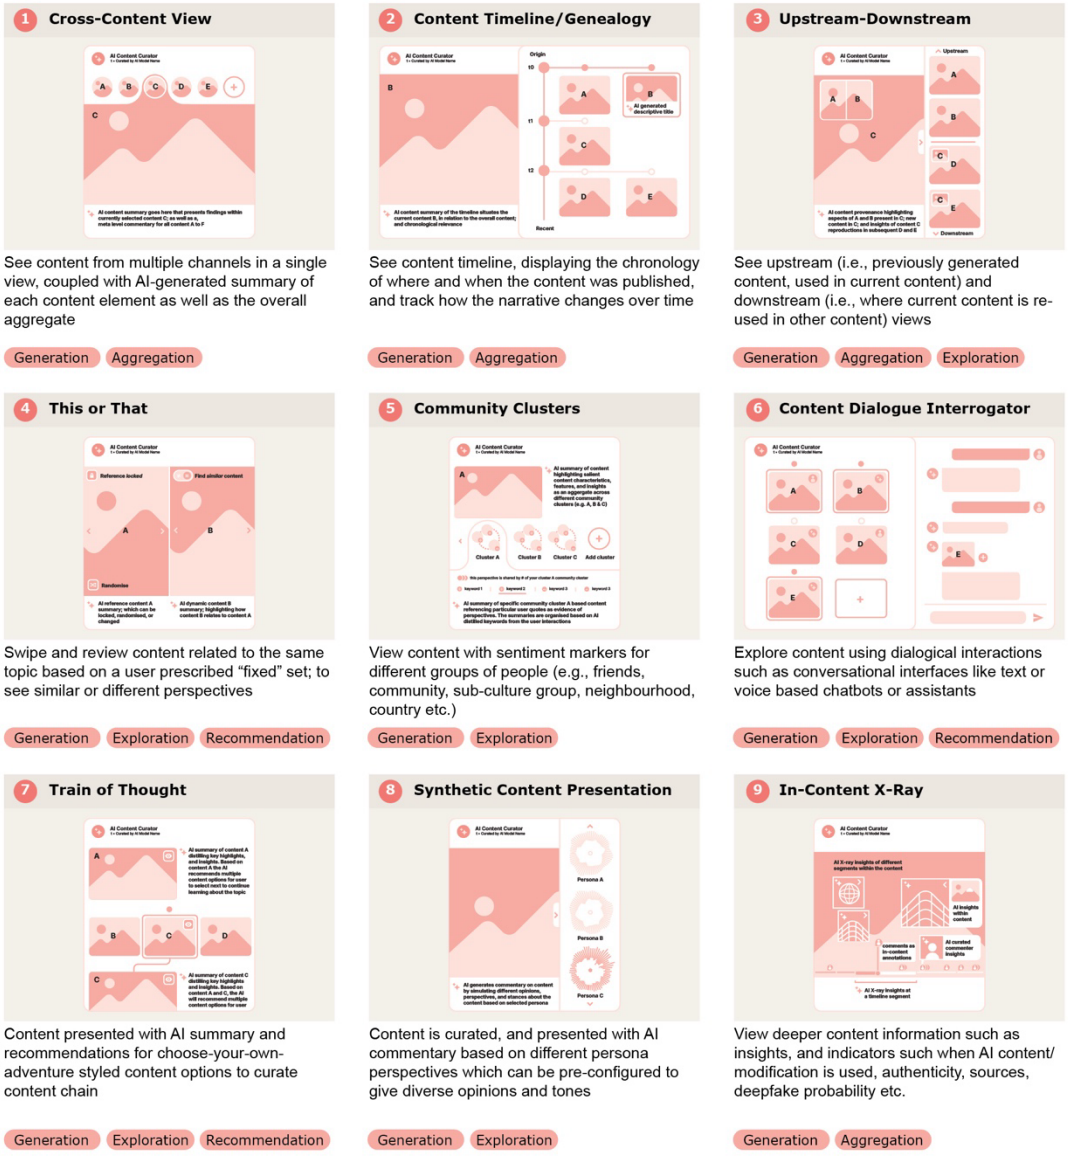

Figure 1. Portfolio of Interface Design Patterns

This design pattern portfolio presents different ways to explore how AI-curated content might be presented to users. The content can comprise of text, images, or video media. Each concept is presented as an interface design pattern card, drawing on a familiar visual vocabulary to current social media user interface content. To illustrate how these patterns can be used, we present high-fidelity interfaces depicting a hypothetical case of news content about wildfires in Australia. All illustrated content is purely fictional, and the use of mainstream news outlet imagery is for illustrative purposes only. All content was created by the authors, using licensed images from *Adobe Stock*, generating images from *Adobe Firefly 3*, (some modified using *Adobe Photoshop*). Placeholder text content within the interface is created by adapting content generated through *ChatGPT-4o*. The interface design patterns were created by sketching ideas, later mocked up using *Adobe Illustrator*. The high-fidelity interfaces were developed using *Figma* and collated on *Adobe InDesign*.

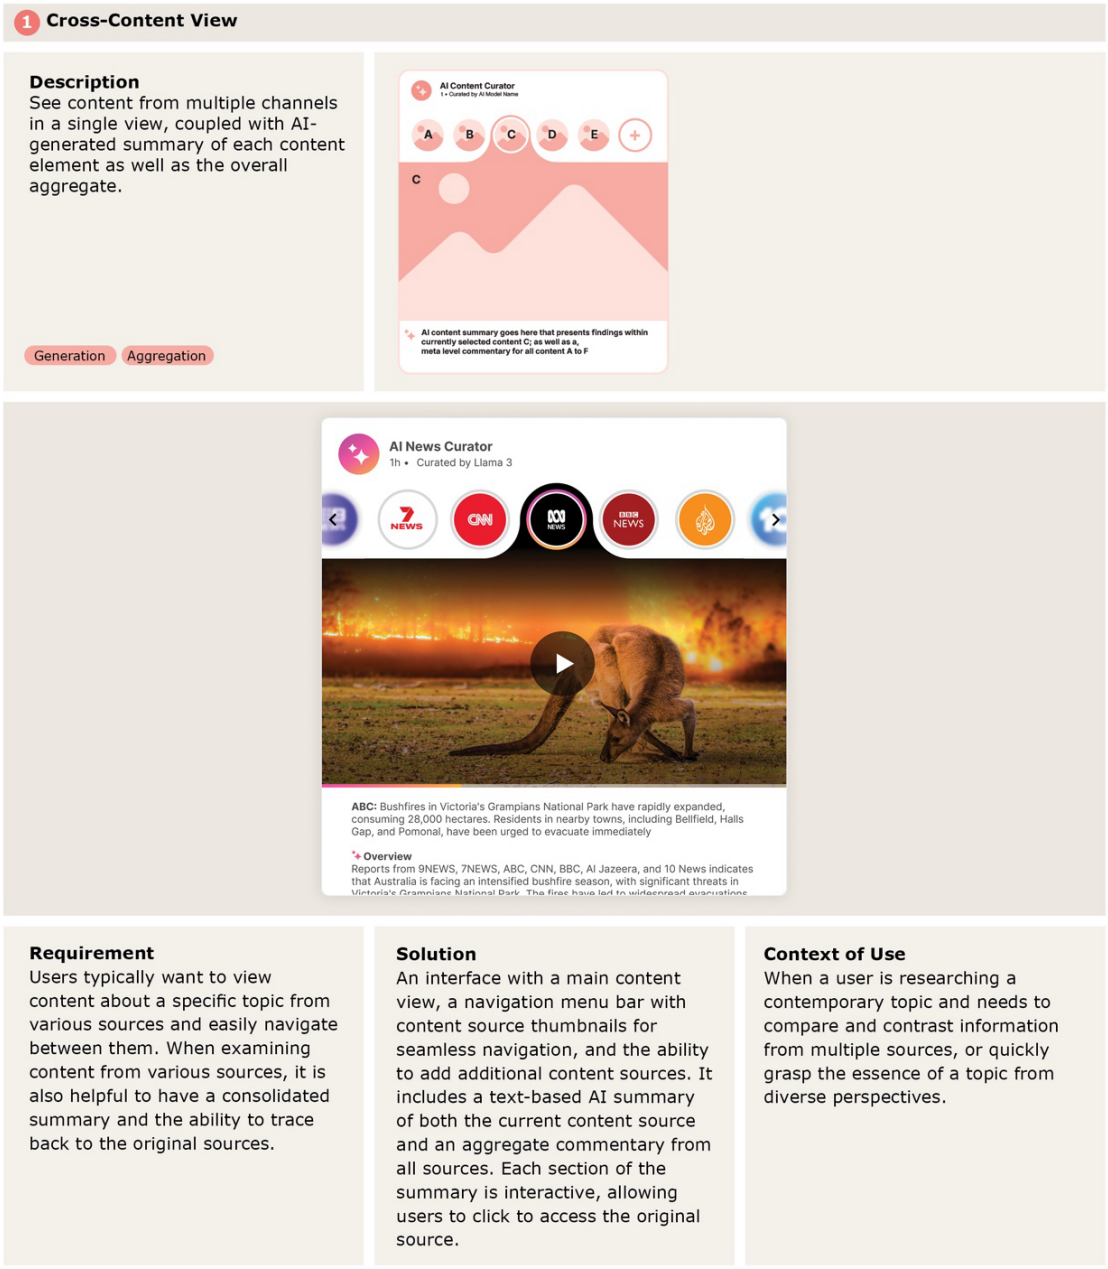

Figure 2. Design Pattern: Cross-Content View

4 Commentary & Discussion

Level of AI Involvement in Content Curation. The level of AI involvement in content curation directly impacts how users interact with the content and their agency for content consumption. At a basic level, AI can aggregate content, structuring it based on predefined categories or timelines (*Aggregation*). More advanced approaches can assist users in exploring content by surfacing relevant connections or highlighting key themes (*Exploration*). AI can also take a more active role through personalised recommendations, adapting to user preferences and behaviours (*Recommendation*). At the highest level of involvement, AI can generate (*Generation*) new content, such as summaries, explanations, or even reformulated narratives tailored to different audiences. Generation applies to

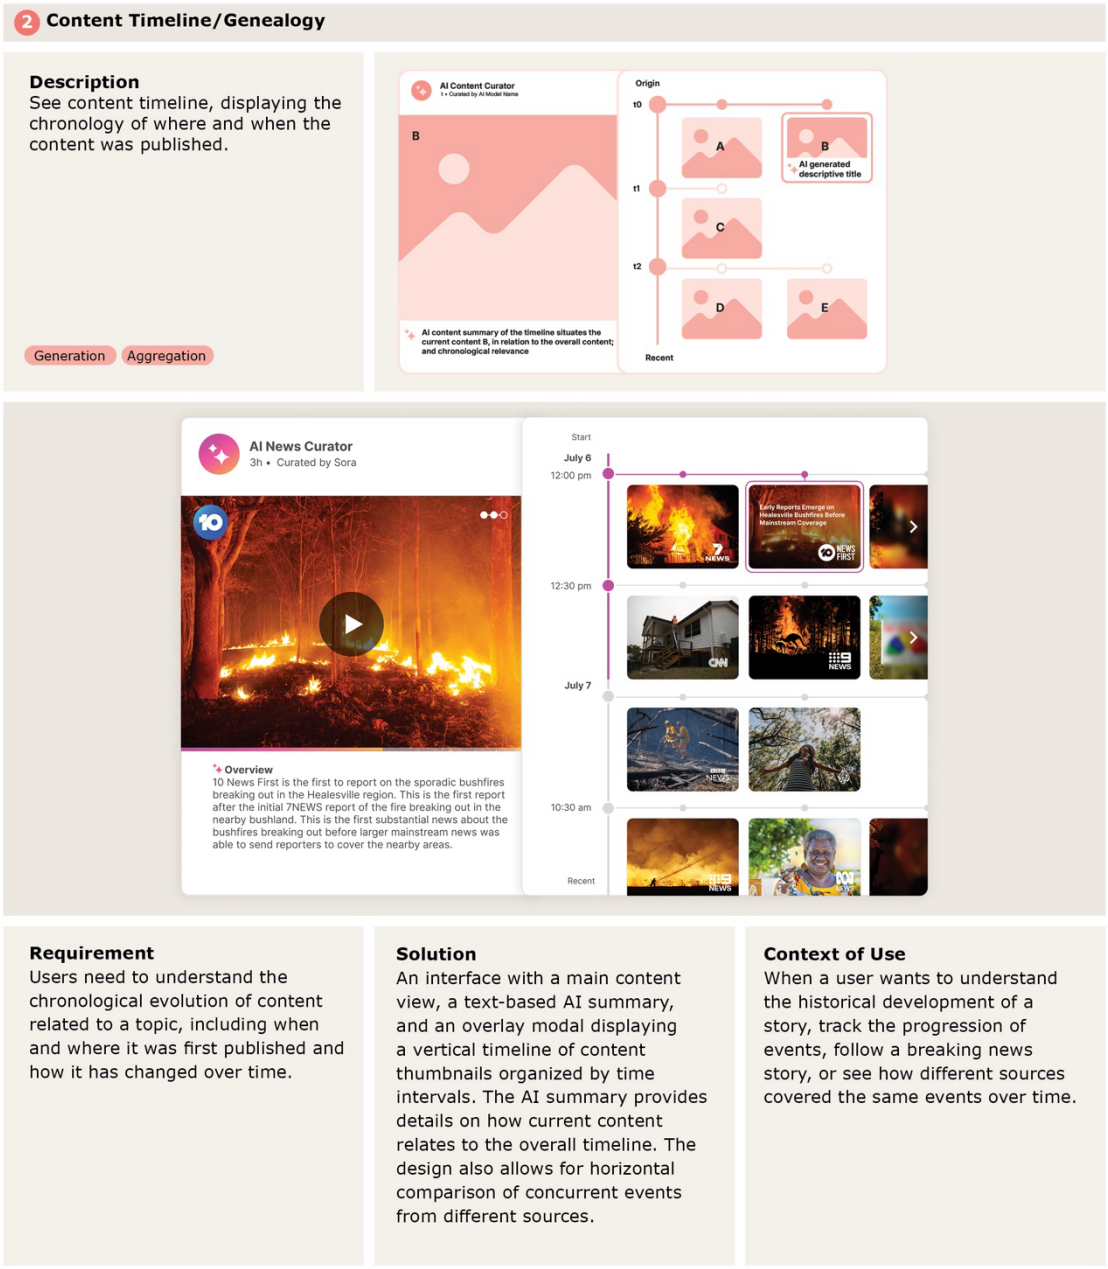

Figure 3. Design Pattern: Content Timeline/Genealogy

all nine patterns. Each of these strategies addresses distinct user needs, from passive consumption to active engagement, and requires careful design considerations to balance automation with user control and trust. Our interface design pattern catalogue explores different ways of applying AI-enabled content curation – such as aggregating, recommending, or exploring content – covering a range of interaction types.

Impact on User Engagement. Some concepts require more significant user input and hence require proactive user engagement. Others are more tailored to passive consumption. Formatting the concepts as in this portfolio opens them up as boundary objects (Star & Griesemer, 1989), i.e., externalised design concepts that can be used to have a dialogue with possible users. This enables us to explore concepts such as user agency, control, and perspectives on how AI-content integrations are

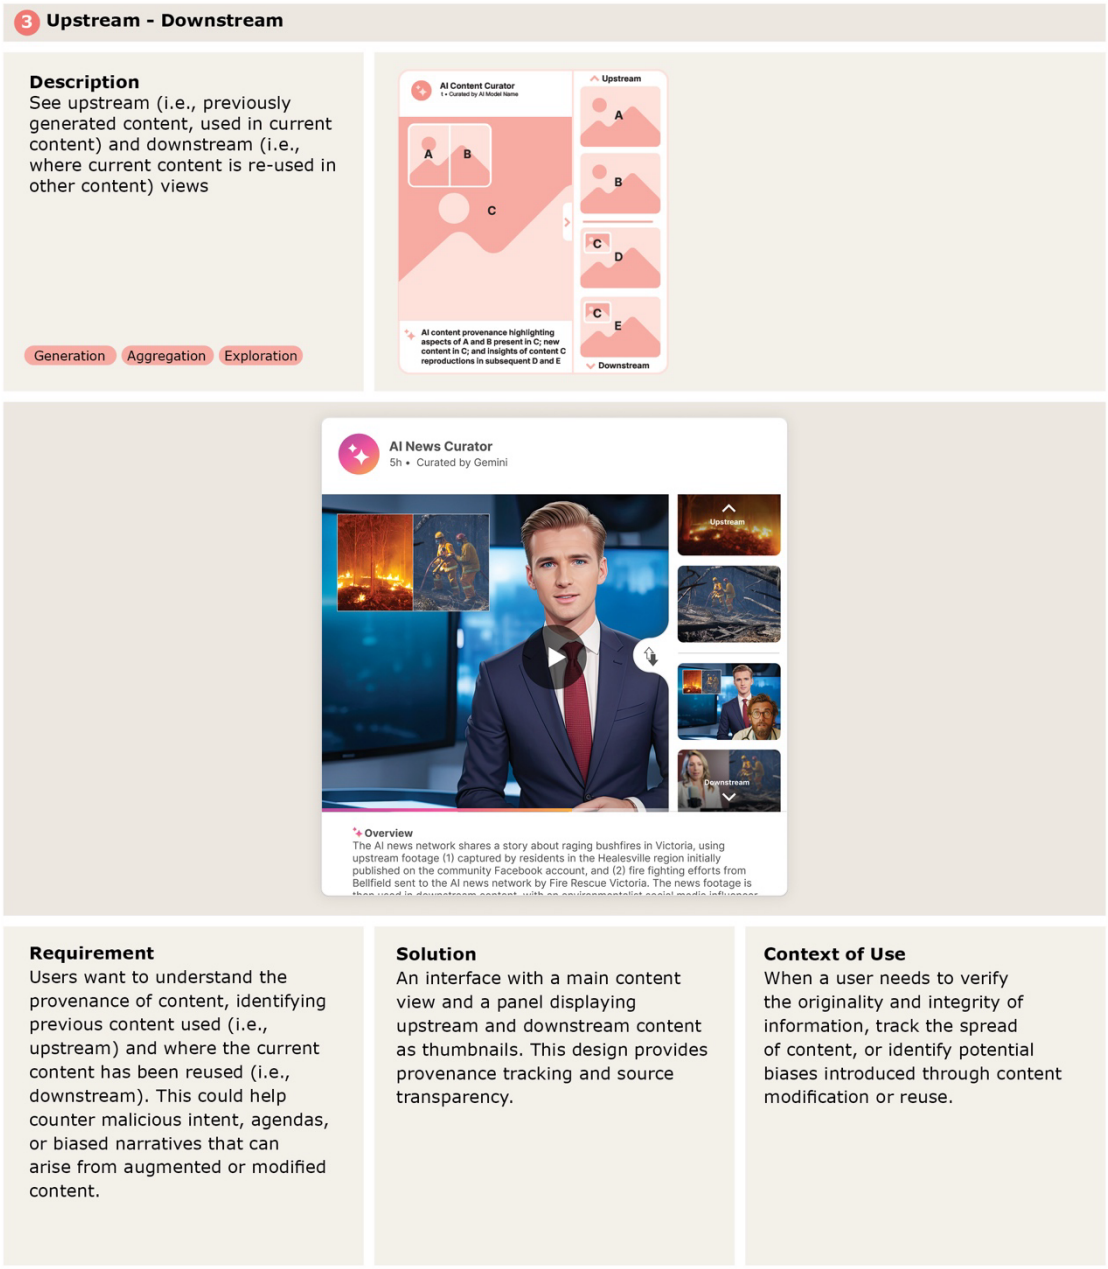

Figure 4. Design Pattern: Upstream-Downstream

being rolled out — and in what contexts. The use of examples – like the illustrated news on wildfires – allows us to move a step further from the patterns as building blocks, to illustrate ideas that can serve as provocations (Boer & Donovan, 2012), where news only serves as one of a multitude of application context possibilities that can be explored. This approach provides the ability to explore these concepts and where users might provide feedback (e.g., opportunities to gauge usefulness within the interface itself), that might otherwise be challenging to conceptualise.

Designing for Content Properties and Context. The different patterns are also attuned to the different characteristics of the content itself— i.e., patterns like *Content Timeline/Genealogy* (See: Figure 3) and *Upstream-Downstream* (See: Figure 4) content deal with temporality, whereas for other patterns, i.e., *Community Clusters* (See: Figure 6), *Synthetic Content Presentation* (See: Figure 9), and *In-Context*

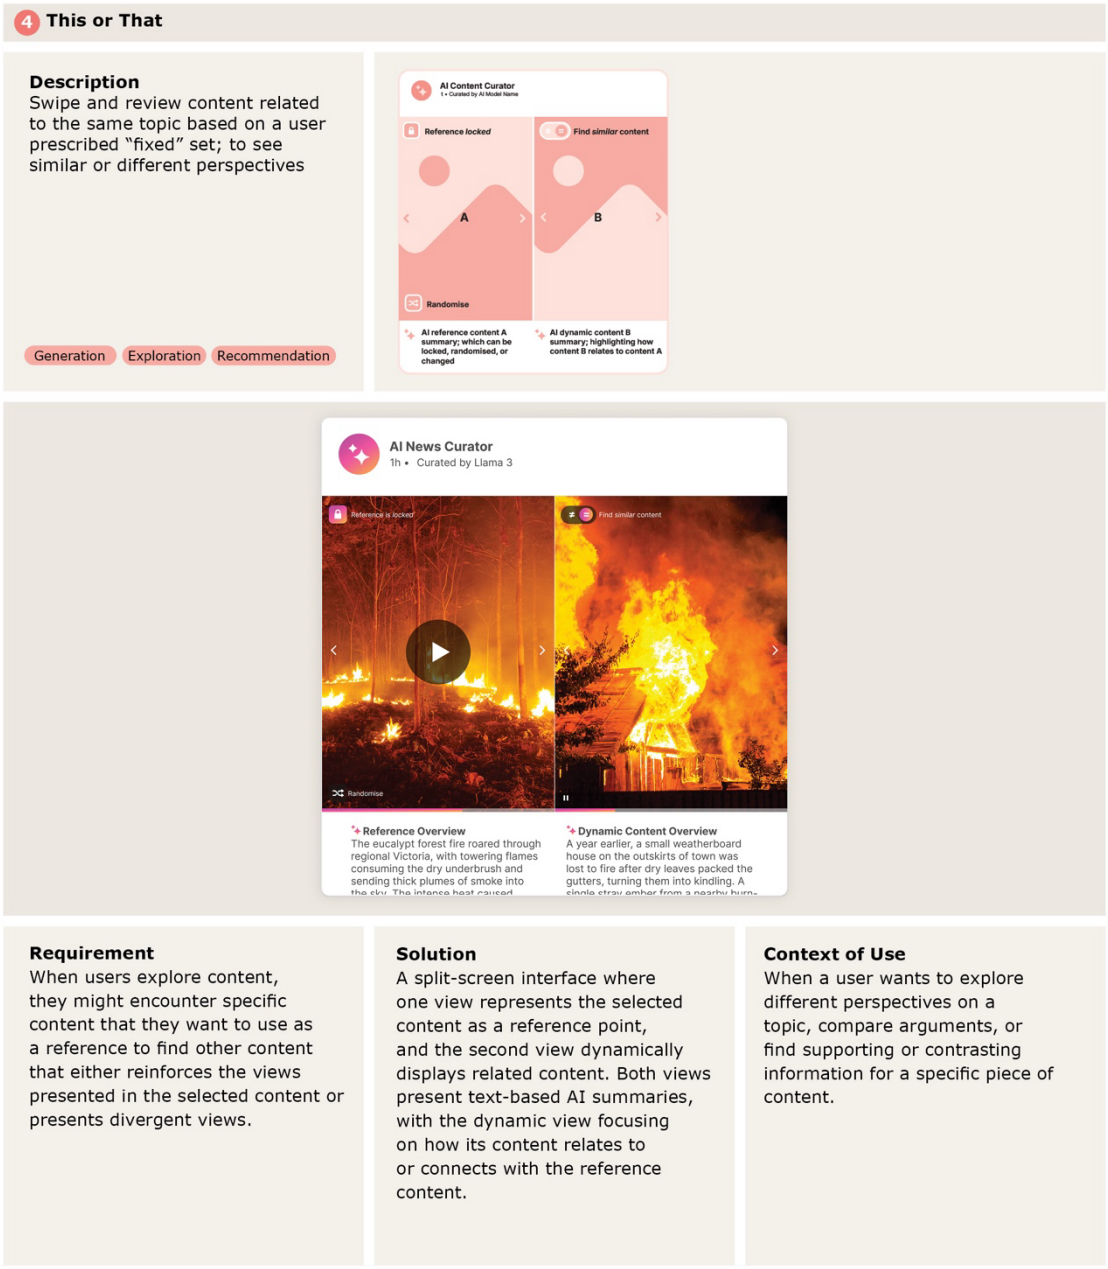

Figure 5. Design Pattern: This or That

X-Ray (See: Figure 10), temporality is not as apparently consequential. This work focuses predominantly on the interaction layer that connects the content input (i.e., elements that compose the data) with the application use cases (i.e., how the content is consumed by users) and the motivations that over others. This then means that the decision to select one pattern, (or assembly of patterns) over others can privilege certain perspectives, goals, and values that can impact user experience and agency. For instance, certain design patterns in the catalogue support specific characteristics (and use cases) more prominently than others. The *Cross-Content View* (See: Figure 2) offers transparency by displaying a summary of all sources, where each piece of summarised text is clickable, allowing users to navigate directly to the original source. *Content Timeline/Genealogy* offers

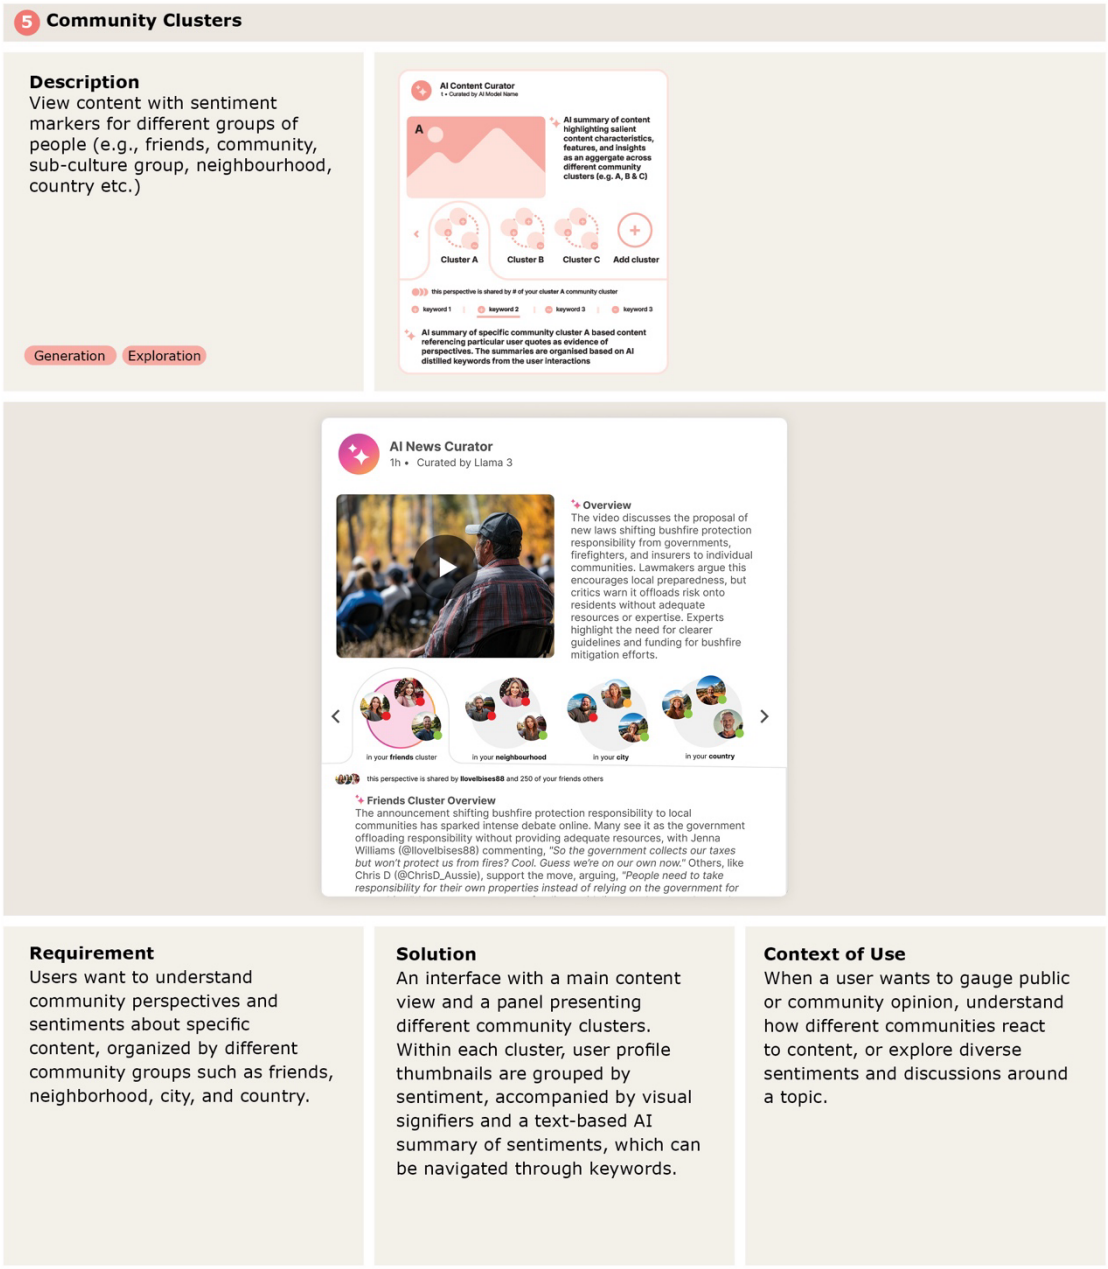

Figure 6. Design Pattern: Community Clusters

temporal information of the events and could allow the user to contextualise information and understand how the story is developed, while potentially eliminating bias as the user is looking at different sources at the same time. *Upstream-Downstream* aims to prevent cherry picking and the selection of slices of information that, without context, can represent specific arguments and support particular agendas by providing users the ability to see where the source content comes from, and how the content being viewed is further propagated and spread. The *This or That* pattern (See: Figure 5) encourages user agency by enabling the user to have control over their confirmation bias i.e., the user gets to decide whether they only want to consume content that reaffirms their views, and echoes sentiments they wish to hear, or allows them to compare completely different perspectives and narratives.

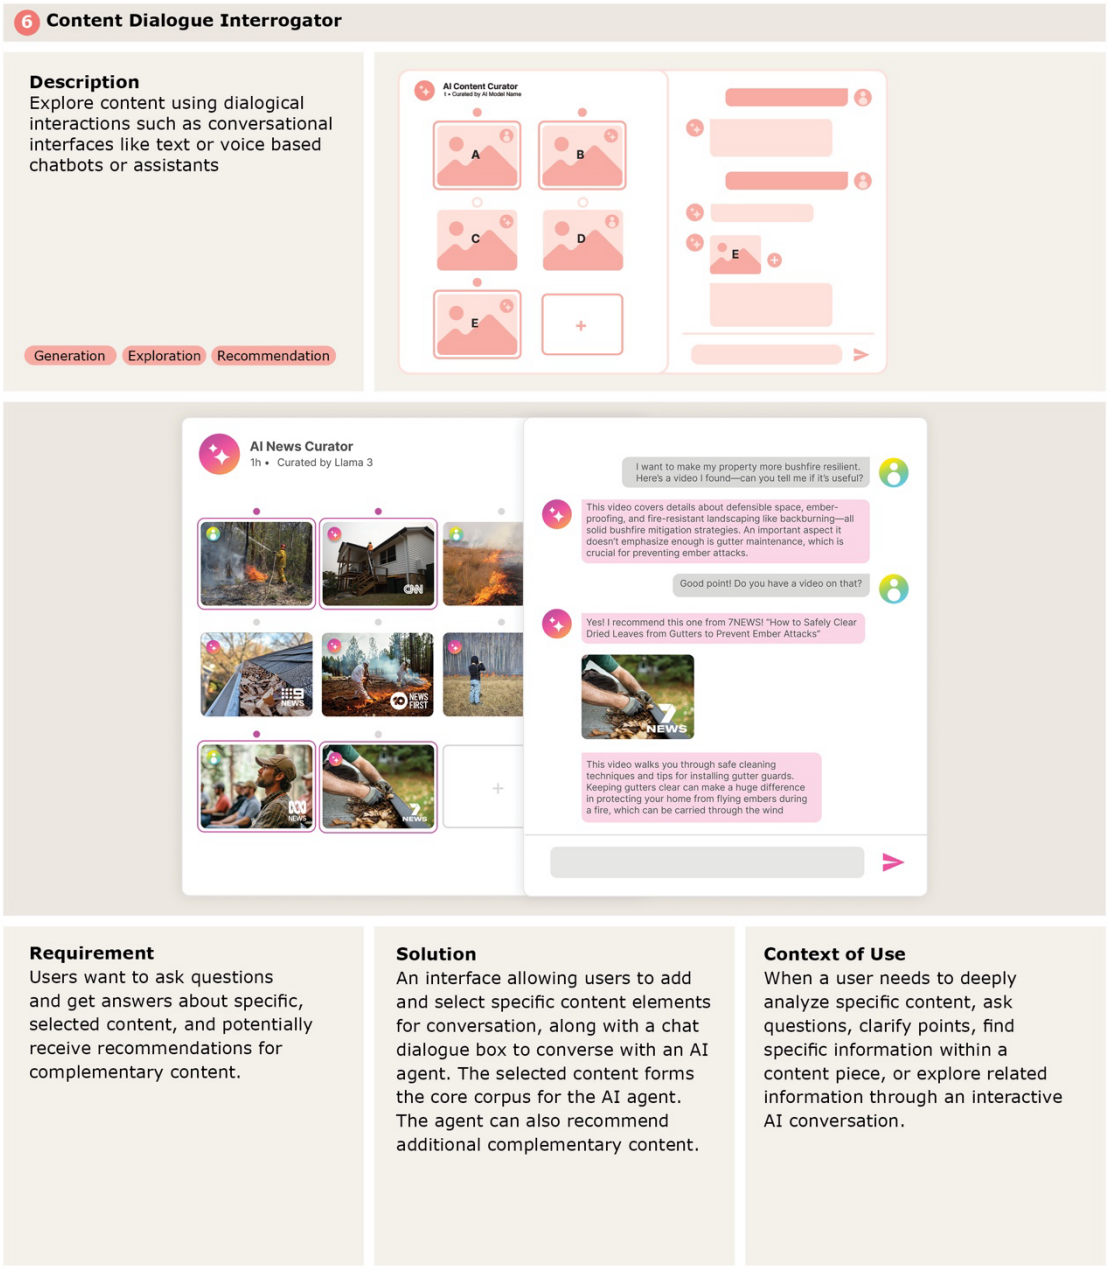

Figure 7. Design Pattern: Content Dialogue Interrogator

Transparency in AI-Generated Content. A common thread that runs through most of the patterns is the prioritisation of transparency through the presentation of both the original source content, coupled with AI-generated summaries. The interfaces enable users to not only engage in content that has been aggregated, but also explore the source materials, make their own inferences and come to their own conclusions. Whilst there are aspects of AI summaries, and recommendations, the users can possibly use these content consumption tools to fact check content veracity and accuracy. This aims to help increase trust and allow easier verification of the AI's outputs, which is essential given the often described *'black box'* nature of AI's output generation, where responses can sometimes include un-factual results, as well as *'hallucinations'* (Yeon et al., 2024).

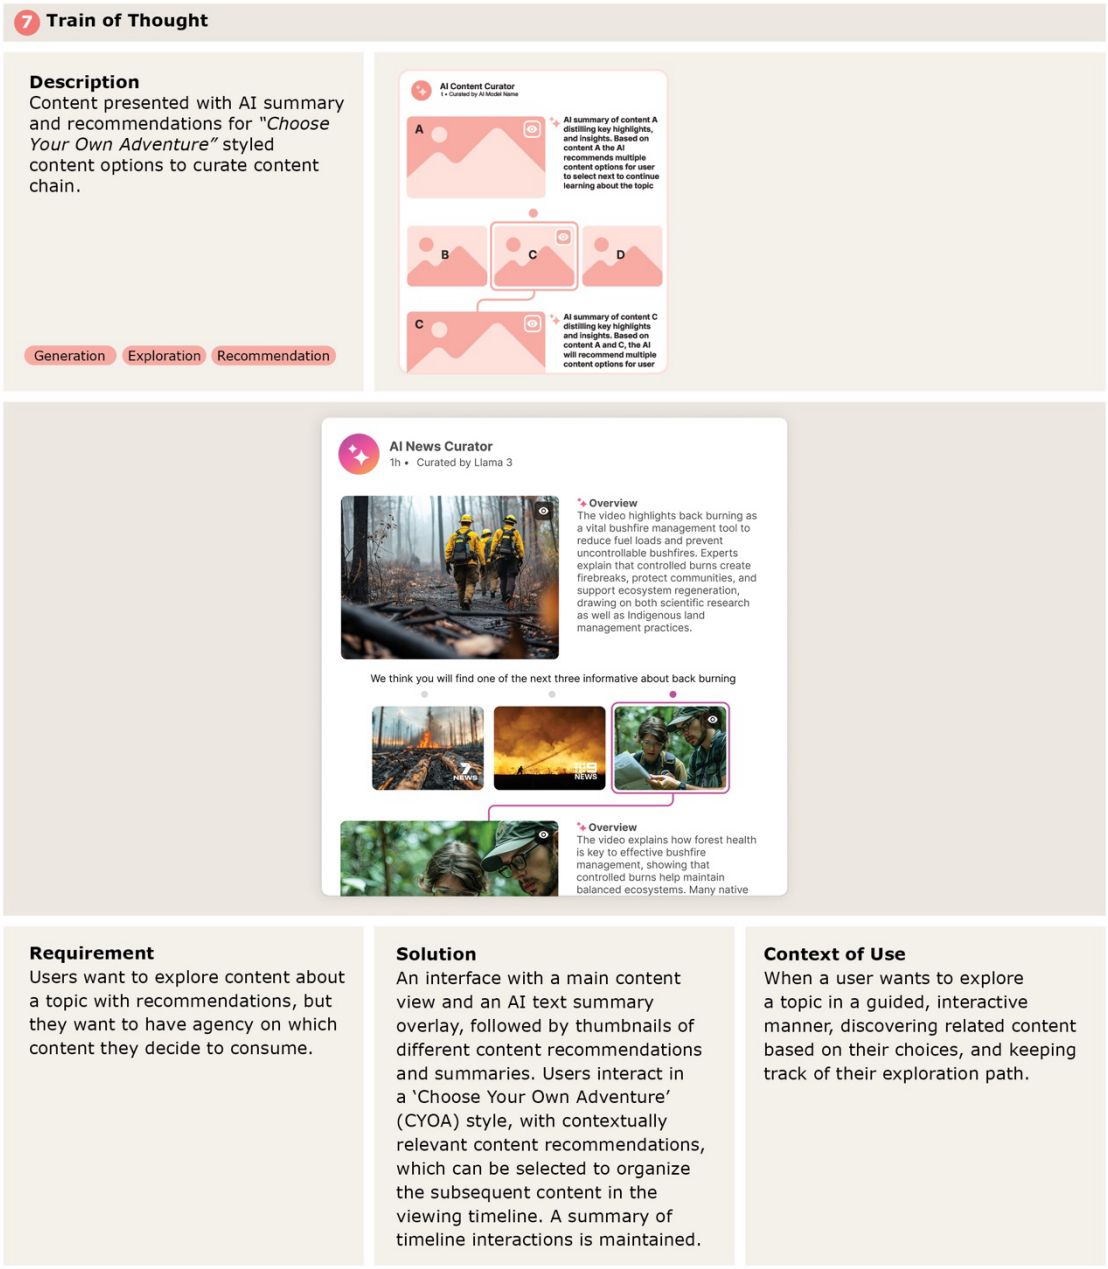

Figure 8. Design Pattern: Train of Thought

5 Conclusion & Future Work

The central contribution of this pictorial is the introduction of a vocabulary and visual language of design patterns for thinking about AI curated content. We see these patterns as a necessary first step of pre-work, especially in exploring how AI curated content can be designed in collaboration with and validated by end users. This is simply because it is not an easy ask, even with the recent proliferation of consumer-facing AI tools and products, to expect end-users to understand and be comfortable with engaging in dialogue about co-designing AI curated content systems, without having a basic understanding about the underlying mechanics of AI. Hence, we see these design patterns play an integral role as boundary spanning objects that can help bridging these conceptual gaps as participatory design tools, restoring user agency for interface design (Rezk et al., 2024). The modular

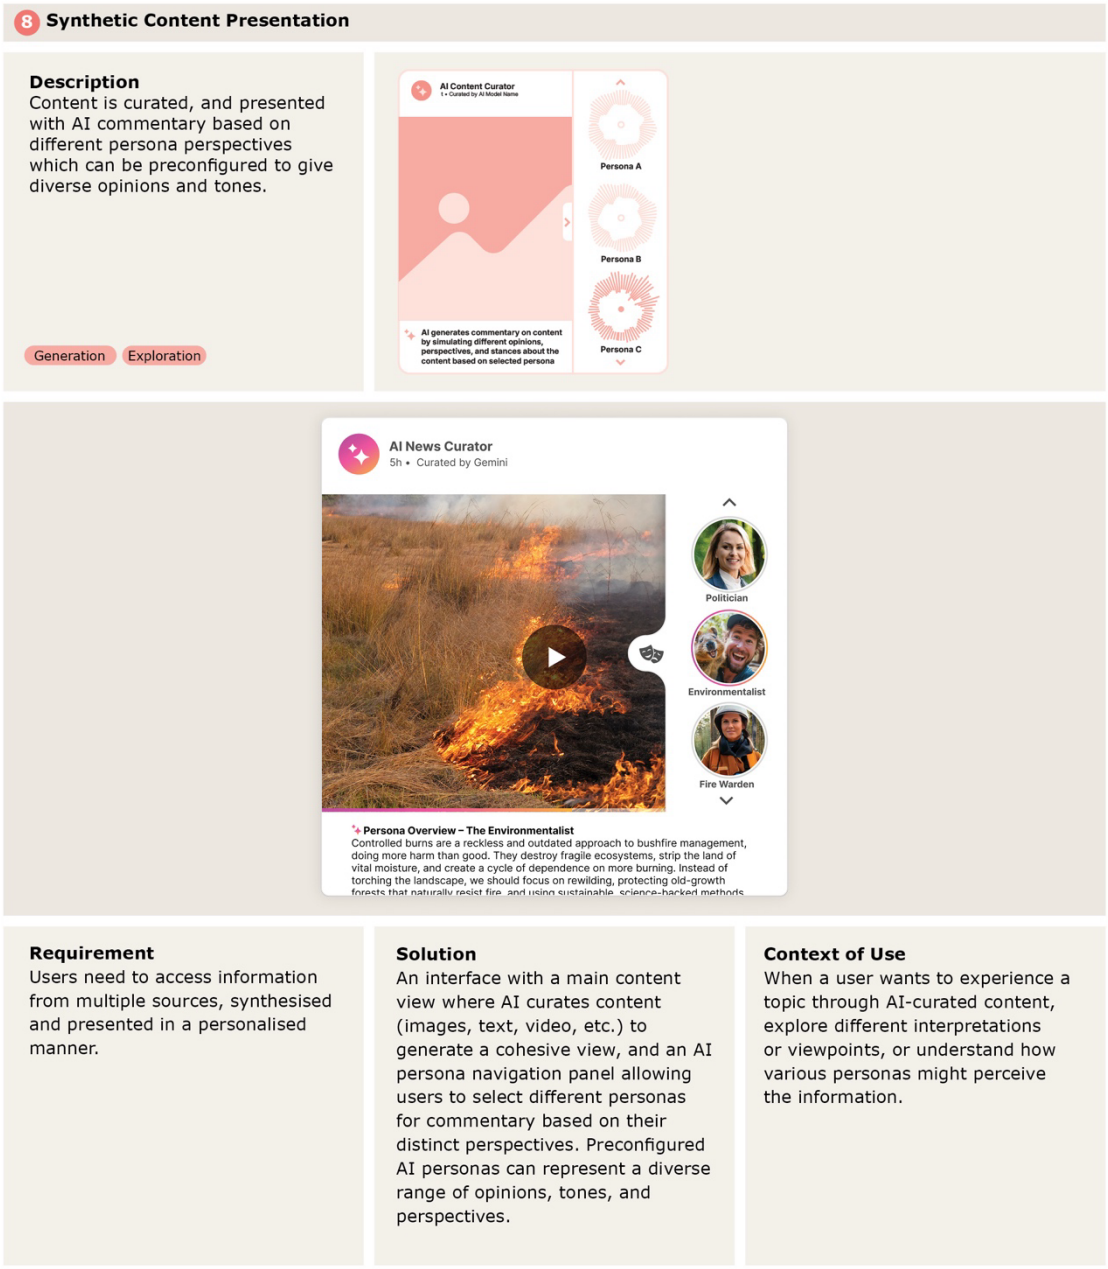

Figure 9. Design Pattern: Synthetic Content Presentation

structure of the design patterns makes them highly adaptable — enabling us to operationalise them by reformatting the patterns into participatory materials such as design cards (See: Khan et al., 2025), a very popular design material (e.g., see: (Hsieh et al., 2023; Roy & Warren, 2019). A lot of recent design cards have been developed to surface broader questions about AI (e.g., (Croisdale et al., 2023; Ghajargar & Bardzell, 2022; Khan et al., 2023; Microsoft, 2019)). Transforming this design pattern portfolio into a deployable participatory toolkit is in a similar vein to previous work where conceptual ideas and frameworks (Hornecker & Buur, 2006; Khan et al., 2024; Li et al., 2020) have been adapted into participatory materials as design cards (Hornecker, 2010; Khan et al., 2023; Li et al., 2021) to engage participants in collaborative activities and make complex concepts easier to understand.

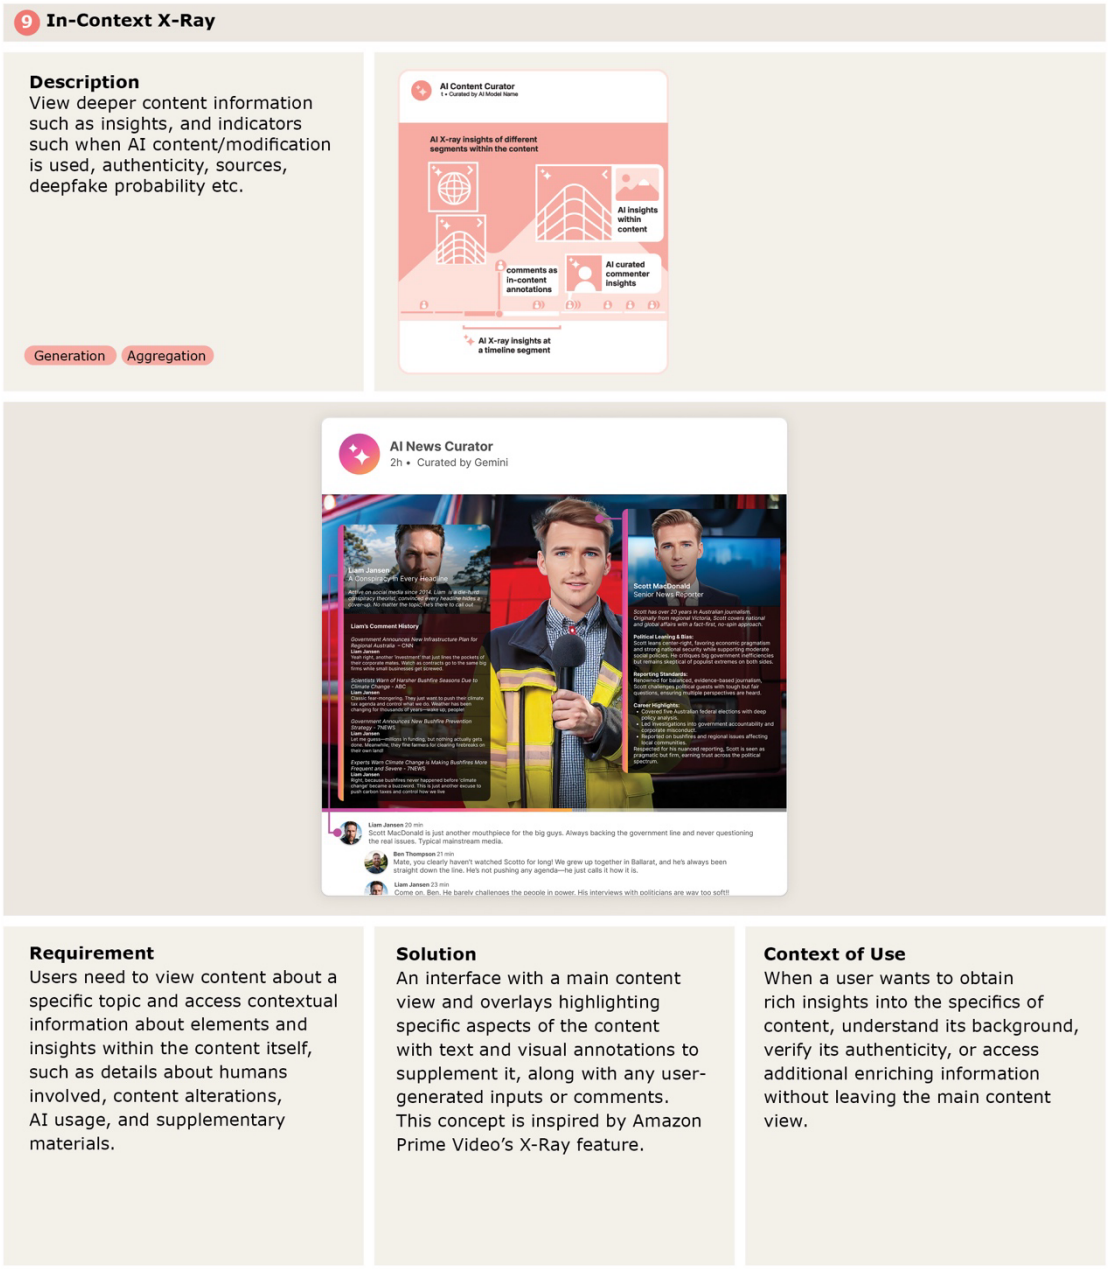

Figure 10. Design Pattern: In-Context X-Ray

We also reflect on the underlying dilemmas of fidelity that emerged through the process of creating these patterns, as we saw that different levels of pattern concept complexity can significantly impact how easily one understands and communicates a pattern. For instance, while commonly agreed upon in most interface design discourse, low fidelity artefacts are better suited for user feedback as they invite more critique. However, here we actually saw the converse. We observed that the higher fidelity renderings of some of the patterns were actually instrumental in bridging the conceptual gaps, and presenting the patterns as meaningful, easy to understand, provocations. This was especially evident in a pattern like *Upstream-Downstream*. This draws parallels with the importance that design fiction scholarship puts on the notion of 'plausibility' of design fiction artefacts, to be convincing provocations (Bleecker et al., 2022). We see this as a very interesting strand of future research to explore.

There is a further need to explore ethical considerations more closely than what has been presented in the portfolio, emphasising accountability, privacy, and user consent (Pranav Dixit, 2024) when designing future iterations. The intention behind these design patterns is to make news consumption more efficient and engaging. For users, the patterns enhance how information is explored and understood. For curators, they offer flexible tools to diversify the structure and presentation of content in meaningful ways. Some patterns focus on breadth by bringing together multiple perspectives, such as *Cross-Content View* or showing how information evolves through *Content Timeline* and *Upstream-downstream* views. Others support depth, like *This or That* and *Train of Thought* (See: Figure 8), which guide users to explore a topic in detail.

However, these designs also raise important ethical considerations. While AI systems play a major role in curating content, human news curators still make key decisions about which patterns to use, how to combine them, and what parameters to apply – i.e., these designs are meant to be used within and contribute towards a human-in-the-loop or human-AI cooperation approach (Breckner et al., 2025; Schneider et al., 2025; Spina et al., 2023). For example, in the *Community Clusters* pattern, curators decide which clusters to include, and these choices can affect which voices or viewpoints are highlighted. This interaction between algorithmic curation and human judgment shapes both what content is shown and how it is interpreted. Moreover, when AI involvement is high, such as in synthetic content generation, it can subtly influence meaning in ways that audiences may not recognise. Therefore, it is crucial to maintain transparency about AI's role in the curation process and to attribute all content clearly to its original authors and sources. Such transparency and attribution support accountability and help users understand how both human and algorithmic decisions shape what they see. Additionally, users should retain autonomy over what they choose to see, with AI primarily assisting the navigation process. Without such autonomy, biases embedded in algorithms could determine which content is highlighted, even in patterns like cross-content view, ultimately limiting the diversity of perspectives. Furthermore, some patterns, such as the *Content Timeline*, could be exploited by bad actors to amplify misinformation, for instance, by presenting sequences of false stories in a format that appears more credible. Similarly, patterns designed for deep exploration i.e., *This or That* or *Train of Thought*, could be misused to create echo chambers, much like the phenomenon of doomscrolling on platforms like TikTok.

We encourage future work to use these patterns for empirical studies, including systematic validation (user testing and evaluation, usability studies, expert reviews, application in real-world contexts) as well as co-design for further pattern iterations. We also see incremental value in explorations of single pattern variations (e.g., layout, content, domain), which can help designers further explore trade-offs between patterns variants, and reflect on their design decisions, their implications, and priorities. There is also scope for conducting more comparative analysis with other design pattern libraries to situate these patterns, examining congruencies and blind spots. Future work should also explore how to make the proposed design patterns more adaptive, ensuring they remain user-friendly across diverse usage contexts, including different devices, user groups, and environments (Hussain et al., 2018). There is a need to further experiment using these patterns across different user studies and contexts e.g., they can be used to interrogate concepts for user engagement (O'Brien et al., 2018); be the basis for developing an AI curated content based heuristic evaluation (Kim et al., 2024).

This work represents an early step towards using design patterns to help make AI curated content that balances efficiency with agency, prioritising transparency, user autonomy, utility and accountability.

Acknowledgements

This research was conducted by the ARC Centre of Excellence for Automated Decision-Making and Society (CE200100005), and funded fully by the Australian Government through the Australian Research Council.

References

Acharya, P.S., Sahu, T., & Dixit, P. (2024). Ethical Considerations in AI-Driven User Interfaces. Journal of Informatics Education and Research, 4(1). https://doi.org/10.52783/jier.v4i1.643

Alexander, C. (1977). *A pattern language: Towns, buildings, construction*. Oxford university press.

Bleecker, J., Foster, N., Girardin, F., Nova, N., Laboratory, N. F., Frey, C., & Pittman, P. (2022). *The Manual of Design Fiction*. Near Future Laboratory.

Boer, L., & Donovan, J. (2012). Provotypes for participatory innovation. *Proceedings of the Designing Interactive Systems Conference*, 388–397. https://doi.org/10.1145/2317956.2318014

Breckner, K., Neumayr, T., Mara, M., Streit, M., & Augstein, M. (2025). The Changing Nature of Human-AI Relations: A Scoping Review on Terminology and Evolvement in the Scientific Literature. *International Journal of Human–Computer Interaction*, *0*(0), 1–58. https://doi.org/10.1080/10447318.2025.2482742

Croisdale, G. T., Chung, J. J. Y., Huang, E., Birchmeier, G., Wang, X., & Guo, A. (2023). DeckFlow: A Card Game Interface for Exploring Generative Model Flows. *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 1–3. https://doi.org/10.1145/3586182.3615821

Gamma, E., Richard Helm, Ralph Johnson, & John Vlissides. (1994). *Design patterns: Elements of reusable object-oriented software*. Addison-Wesley.

Gaver, B., & Bowers, J. (2012). Annotated portfolios. *Interactions*, *19*(4), 40–49. https://doi.org/10.1145/2212877.2212889

Ghajargar, M., & Bardzell, J. (2022). Making AI Understandable by Making it Tangible: Exploring the Design Space with Ten Concept Cards. *Proceedings of the 34th Australian Conference on Human-Computer Interaction*, 74–80. https://doi.org/10.1145/3572921.3572942

Hornecker, E. (2010). Creative Idea Exploration Within the Structure of a Guiding Framework: The Card Brainstorming Game. *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction*, 101–108. https://doi.org/10.1145/1709886.1709905

Hornecker, E., & Buur, J. (2006). Getting a Grip on Tangible Interaction: A Framework on Physical Space and Social Interaction. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 437–446. https://doi.org/10.1145/1124772.1124838

Hsieh, G., Halperin, B. A., Schmitz, E., Chew, Y. N., & Tseng, Y.-C. (2023). What is in the Cards: Exploring Uses, Patterns, and Trends in Design Cards. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–18. https://doi.org/10.1145/3544548.3580712

Hussain, J., Ul Hassan, A., Muhammad Bilal, H. S., Ali, R., Afzal, M., Hussain, S., Bang, J., Banos, O., & Lee, S. (2018). Model-based adaptive user interface based on context and user experience evaluation. *Journal on Multimodal User Interfaces*, *12*(1), 1–16. https://doi.org/10.1007/s12193-018-0258-2

Khan, A. H., Kegalle, H., D'Silva, R., Watt, N., Whelan-Shamy, D., Ghahremanlou, L., & Magee, L. (2024). *Automating Thematic Analysis: How LLMs Analyse Controversial Topics*. https://doi.org/10.48550/arXiv.2405.06919

Khan, A. H., Whelan-Shamy, D., Watt, N., D'Silva, R., Kegalle, H., Ghahremanlou, L., & Magee, L. (2023). *AI Sub Zero Bias Cards*. ARC Centre of Excellence for Automated Decision-Making and Society. https://espace.library.uq.edu.au/view/UQ:f998cc0

Khan, A. H., Wijenayake, S., Lawati, S. F. D. A., Spina, D., & Hettiachchi, D. (2025, October 20). *Design Patterns for AI Curated Content Toolkit*. ARC Centre of Excellence for Automated Decision-Making and Society. https://espace.library.uq.edu.au/view/UQ:24dc3a6

Kim, T.-S., John Ignacio, M., Yu, S., Jin, H., & Kim, Y.-G. (2024). UI/UX for Generative AI: Taxonomy, Trend, and Challenge. *IEEE Access*, *12*, 179891–179911. https://doi.org/10.1109/ACCESS.2024.3502628

Li, H., Häkkilä, J., & Väänänen, K. (2020). Towards a Conceptual Design Framework for Emotional Communication Systems for Long-Distance Relationships. In A. Brooks & E. I. Brooks (Eds.), *Interactivity,*

Game Creation, Design, Learning, and Innovation (pp. 103–123). Springer International Publishing. https://doi.org/10.1007/978-3-030-53294-9_8

Li, H., Khan, A. H., Hurtig, K. M., Jarusriboonchai, P., & Häkkilä, J. (2021). Flexi Card Game: A Design Toolkit for Unconventional Communication Systems for Long-Distance Relationships. *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*, 1–16. https://doi.org/10.1145/3430524.3440650

Microsoft. (2019). *Guidelines for Human-AI Interaction*. https://www.microsoft.com/en-us/research/uploads/prod/2019/04/AI-Design-guidelines_041519.pdf

O'Brien, H. L., Cairns, P., & Hall, M. (2018). A practical approach to measuring user engagement with the refined user engagement scale (UES) and new UES short form. *International Journal of Human-Computer Studies*, *112*, 28–39. https://doi.org/10.1016/j.ijhcs.2018.01.004

Pranav Dixit, P. S. A., Dr. Tripti Sahu,. (2024). Ethical Considerations in AI-Driven User Interfaces. *Journal of Informatics Education and Research*, *4*(1). https://doi.org/10.52783/jier.v4i1.643

Rezk, A. M., Simkute, A., Luger, E., Vines, J., Elsden, C., Evans, M., & Jones, R. (2024). Agency Aspirations: Understanding Users' Preferences And Perceptions Of Their Role In Personalised News Curation. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3613904.3642634

Roy, R., & Warren, J. P. (2019). Card-based design tools: A review and analysis of 155 card decks for designers and designing. *Design Studies*, *63*, 125–154. https://doi.org/10.1016/j.destud.2019.04.002

Schneider, D., Chaves, R., Pimentel, A. P., de Almeida, M. A., De Souza, J. M., & Correia, A. (2025). AI-mediated Collaborative Crowdsourcing for Social News Curation: The Case of Acropolis. *Proceedings of the 2025 ACM International Conference on Interactive Media Experiences*, 395–401. https://doi.org/10.1145/3706370.3731644

Spina, D., Sanderson, M., Angus, D., Demartini, G., Mckay, D., Saling, L. L., & White, R. W. (2023). Human-AI Cooperation to Tackle Misinformation and Polarization. *Commun. ACM*, *66*(7), 40–45. https://doi.org/10.1145/3588431

Star, S. L., & Griesemer, J. R. (1989). Institutional ecology,translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, *19*(3), 387–420.

Yeon, J., Jung, Y., Baek, Y., Lee, D., Shin, J., & Chung, W. Y. (2024). User Preferences on a Generative AI User Interface Through a Choice Experiment. *International Journal of Human–Computer Interaction*, 1–12. https://doi.org/10.1080/10447318.2024.2400379