NordDesign 2018
August 14 – 17, 2018
Linköping, Sweden

Machine learning as a design material: a curated collection of exemplars for visual interaction

Danwei Tran Luciani[1], Martin Lindvall[2], Jonas Löwgren[3]

[1, 2, 3]Linköping university
[2]Sectra AB
[1]danwei.tran.luciani@liu.se
[2]martin.lindvall@sectra.com
[3]jonas.lowgren@liu.se

Abstract

Although machine learning is not a new phenomenon, it has truly entered the spotlight in recent years. With growing expectations, we see a shift in focus from performance tuning to awareness of meaningful interaction and purpose. Interaction design and UX research is currently in a position to provide important and necessary knowledge contributions to the development of machine learning systems.

Machine learning can be viewed as a design material that is arguably more unpredictable, emergent, and "alive" than traditional ones. These characteristics suggest practice-based work along the lines of research-through-design as a promising approach for machine learning system development research. Design researchers using a research-through-design approach agree that a created artifact carries knowledge, but there is no consensus on how such knowledge is best articulated and transferred within academic discourse. Knowledge contributions need to be abstracted from the particular to a higher level. We suggest *curated collections*, a variation of annotated portfolios, as a way to abstract and communicate intermediate-level knowledge that is suitable and useful for the research-through-design community. A curated collection presents thoughtfully selected and inter-related exemplars, articulating their salient traits. The insights collected in a curated collection can be used to inform future design in related design situations.

This paper provides a curated collection addressing the fine-grained details of interaction with machine learning systems. The examples are drawn from highly visual interaction, predominantly in the domain of digital pathology. The collection of interaction examples is used to elicit a set of salient traits, including the preservation of visual context, rapid real-time refinement, leaving traces, and applying judicious automation. Finally, we show how this curated collection could inform the design of a future system in a different domain. The insights are applied to a case of interaction design to support air traffic controllers in their collaboration with future agentive systems.

Keywords: digital design, interaction design, machine learning, curated collection

1 Introduction

Machine learning has existed for a long time as an engineering practice. Recent waves of public attention have raised the expectations for machine learning to support people's professional as well as leisurely practices in many application domains. Currently, successful innovation in machine learning depends on designing for meaningful human use rather than solving technical challenges. For practical applications, the performance of the joint human-machine ensemble is arguably more important than, say, the accuracy of standalone automated predictions. In other words, the disciplines of UX and interaction design need to contribute to future research relating to the development of machine learning systems. Deploying a design research approach to machine learning systems development opens several areas of inquiry, including the choice of training datasets, curation of data, and optimization of targets – but the most obvious research orientation concerns the fine-grained interaction between human user and machine learning system, including feedback and validation mechanisms for what is known as machine teaching (Lindvall et al, 2018a). This is also the topic of our present work.

2 Related works

Design research into machine learning systems development is an emerging field, and the available literature is somewhat scarce. However, we find current notions of machine learning as a design material to be potentially fruitful. Furthermore, there is a need to survey what is known about methods and knowledge forms in general interaction design research with a view towards adapting it for machine learning systems development research.

2.1 Machine learning as a design material

The current hype around machine learning has probably fueled the misconception of what machine learning is and what it is capable of. It is not a magic ingredient that can solve any kind of problem; adding machine learning does not necessary mean that the result will be better or "smarter". Other scholars have recently suggested that framing machine learning as a design material is an appropriate way of working with it (Dove et al, 2017; Holmqvist, 2017; Yang, 2017). It is a design material that needs close consideration of the context of how and where it will be applied. Let us continue this line of thought of dealing with machine learning as a design material with its own unique characteristics.

Compared to traditional design materials, machine learning is more unpredictable, emergent, and "alive". The need for explorative research-through-design for inquiry through design practice is clear (Löwgren, 2016). In this kind of research practice in general, and when designing with machine learning in particular, there is no clear separation between framing the problem and finding a solution. It is something that is done in parallel with both the problem and solution becoming clearer during an iterative process. Consequently, a promising way to approach machine learning system development research is to apply research-through-design.

2.2 Research-through-design as an approach

In human-centered design, research activities are conducted early in the process to gain an understanding of intended users and the current situation. The result of the research informs the design of the product or service, usually with a focus on usability and utility. This approach may not be optimal when it comes to dealing with more complex design situations that require more than just incremental improvements of an existing system.

Designers are often faced with situations that are messy and the problem definition is so fuzzy that it requires them to start by creating a possible solution to help understand the problem. In research-through-design, research and design are two activities that are not just connected, but inseparable. Knowledge is gained through an iterative process of making, framing and reframing (Zimmerman & Forlizzi, 2008; Gaver, 2012). A key component in research-through-design is the creation of artifacts. They are often considered the embodiment of the knowledge produced. While there seems to be a consensus among design researchers that research-through-design produces transferable knowledge, a remaining challenge is how to document and communicate it. An artifact on its own is to some extent a black box. How do we make the box more transparent or unfold it to highlight important experiential qualities or traits?

3 Proposal for a designerly approach to machine learning system development research

An artifact needs to be abstracted to a higher level for the knowledge to be useful beyond the particular context. Abstracting all the way would produce what could be considered general theory, but there are other forms of intermediate-level knowledge populating the space between the particular artifact and the general theory (Löwgren, 2013).

3.1 Curated collection as a form of intermediate-knowledge representation

Research-through-design needs to result in knowledge that is of use to other research-through-design researchers (Zimmerman et al., 2010). In other words, what research-through-design contributes must be accessible for research-through-design research peers to appropriate and use in their own research. This includes knowledge contributions that are conducive to creative design in the context of research-through-design processes. It is common practice in design to use exemplars (exemplary examples) as appropriable and useful knowledge in the context of design knowledge dissemination. Abstracting to a level slightly above specific instances makes it possible to formulate insights on core concepts and essential qualities in a form that is amenable to mediated communication in scholarly media (i.e. text with a few images/illustrations). It also enables assessment of the scope of knowledge contributions (i.e. the breadth of applicability of a specific articulation or insight).

This forms the basis for a kind of knowledge contribution that combines a number of exemplars with abstractions and insights drawn from them. Such a collection of exemplars and articulations is closely related to what Gaver and Bowers (Gaver & Bowers, 2012) call annotated portfolios. The main difference is that Gaver and Bowers limit their selection of exemplars to work originating from the research-through-design researchers formulating the annotated portfolio. Our approach, drawing slightly more widely on perspectives of criticism complementing creative practice, is open to the inclusion of exemplars from other sources. We call this kind of collection a *curated collection*. A curated collection is not exactly a conventional designer's portfolio, since it is not intended to showcase an individual designer's proficiency. It draws upon exemplars that are curated and analyzed in a form of erudite scholarship similar to the work of criticism in more mature fields of creative practice (Bardzell, 2011). Exemplars can be made by the research-through-design researcher providing the knowledge contribution or they can be selected from other sources and analyzed equivalently. Each exemplar is presented in a semi-structured format with enough introduction of context to make the design situation clear, the salient experiential qualities of the design, and the core design idea abstracted to a level where it becomes transferable to new design situations. Exemplars can be positive or negative. Positive exemplars contribute to articulations of promising roads for future development. Negative exemplars help identify directions that may not be very interesting for further development, and thereby assist in determining the scope, or breadth of applicability, of a proposed abstraction.

3.2 A curated collection for fine-grained visual interaction with machine learning systems

It is generally more important to design the interaction for a better performance of the joint human-machine ensemble than to aim for higher diagnostic performance of the standalone AI. This is not only a general design stance – human empowerment rather than replacement – but also a practical concession in our domains of interest where full automation is unrealistic as well as undesirable in the short term. Our specific interest is in designing the fine-grained details of the interaction with machine learning systems. By gathering previous and ongoing projects that share this focus we have composed an example of a curated collection that highlights the following traits: (1) preservation of visual context, (2) rapid real-time refinement, (3) leaving traces, and (4) judicious automation.

3.2.1 Preservation of visual context

For certain domains we might want to design for a situation where the user can improve upon the machine’s output, both in the short term for immediate decision aid and in the long term for continuous improvement of the model’s predictions. In one project described in (Lindvall et al., 2018a), careful attention was paid to human validation of system output.

The task in this example concerns labeling of tissue in digital pathology. The prototype applies a grid pattern over a user-selected area and extracts a small image patch for each point in the grid. Each patch is then fed to a machine learning algorithm that classifies the patches into different categories, which are then shown in a sorted gallery (see figure 1). Each defined class in the trained model is shown as patches in the same gallery. The user can click on a patch to see it in the main view to get a sense of its context in the tissue, and change a label by either dragging the patch to the correct category or by clicking on the button or the corresponding shortcut key.

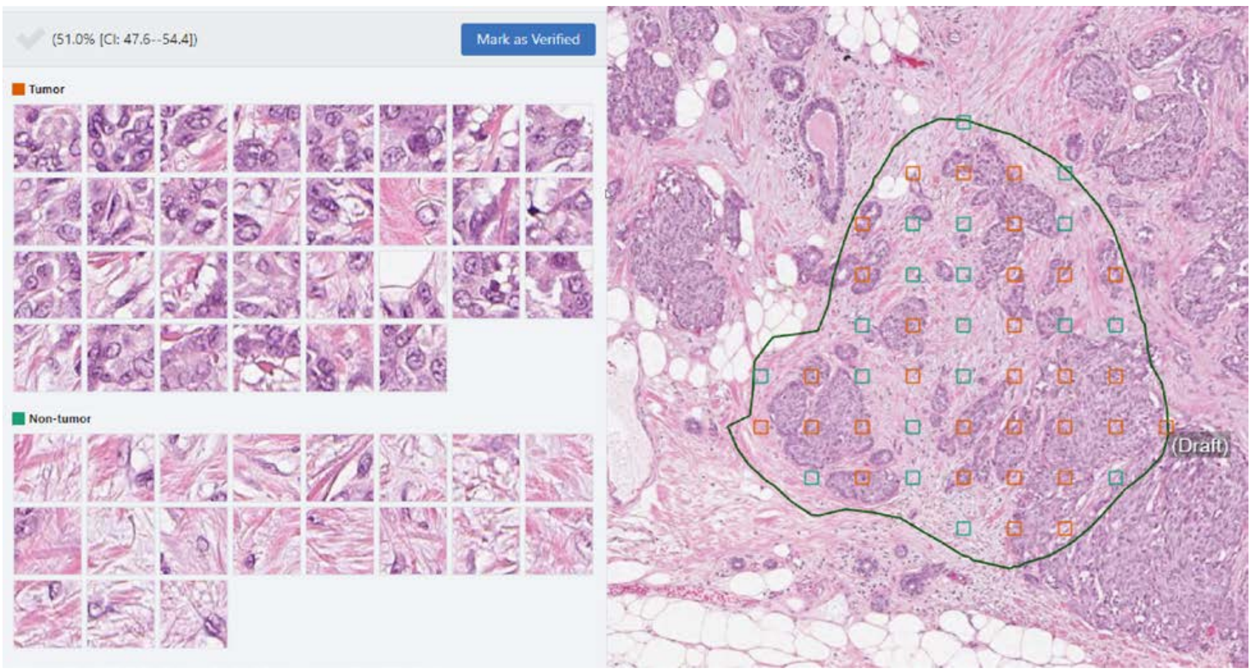

Figure 1. An interface preserving visual context. To the left are image patches sorted into different categories and to the right is an overview of its context in the tissue.

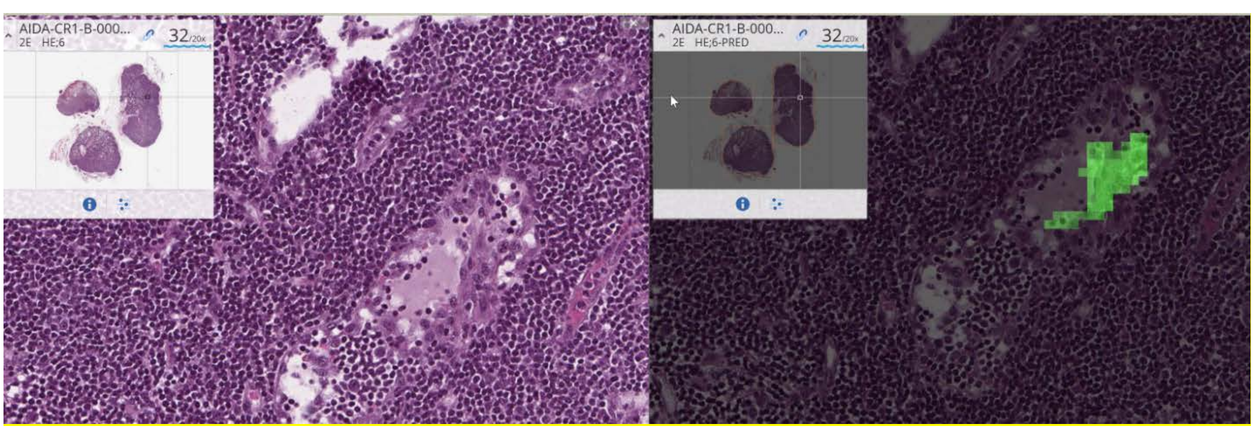

Figure 2. A view of a region preserving visual context (left) compared to the region presented through a heatmap (right).

This approach should be contrasted with the more common way of showing the output of such a classifier for a region through a heatmap (see figure 2). A heatmap occludes the phenomena of interest (the cancerous cells) and makes correction cumbersome, especially for large areas where the user would need to "paint" over all pixels to correct classifications. Categorization tasks in this domain are typically discrete, and each point of interest is deeply dependent on its visual context. Presenting proposed categorizations as separate tiles sorted into buckets emphasizes the discrete and correctable nature of the system's suggestions and preserves the visual context necessary to validate and manually change the labeling if needed.

The core design idea of this exemplar is to preserve the visual context needed to assess actions and suggestions from the automation.

3.2.2 Rapid real-time refinement

For many domains there are no viable alternatives to having humans provide the training data for the machine learning algorithms. With the increasing importance of machine learning follows the importance of having efficient annotation tools. One way to make annotation more efficient is to use machine learning for machine learning in what is called interactive segmentation tools. Research surrounding these types of tools focuses on the correctness of the produced areas and efficiency of the algorithm in terms of computational resources, but there are few studies of how efficient people become at annotating by using these tools.

In a recent work, a human-centered design approach was applied in the construction of such a tool for annotating healthy and non-healthy tissue. In the initial prototype, interaction was experienced as a trading of control between human and machine. After a noticeable delay, the results were received, and the pathologist would make one correction, wait and repeat. Typically, the pathologist would be both intrigued and annoyed by the automatic assignment of the areas that were not specifically drawn over, sometimes resulting in long back-and-forth correction cycles without noticeable progress.

The prototype was revised to allow for more rapid interchange (see figure 3). The tool was rebuilt so that the response to user input typically arrives in less than 40 milliseconds. Drawing on previous work in interactive visualization (Löwgren, 2007), we postulate that this kind of real-time interaction, with the ability to quickly observe and correct many predictions that the model produces, lets the pathologist gain an understanding of the underlying mechanism and its limitations. Depending on the intent of the annotation, the pathologist might want a conservative or a more aggressive use of automation. In order to let the pathologist control this without disruption, a viable approach might be to use stylus input with the pressure mapped to classification sensitivity.

The core design idea is that a near-real-time interaction loop enables a better understanding and training of the automation.

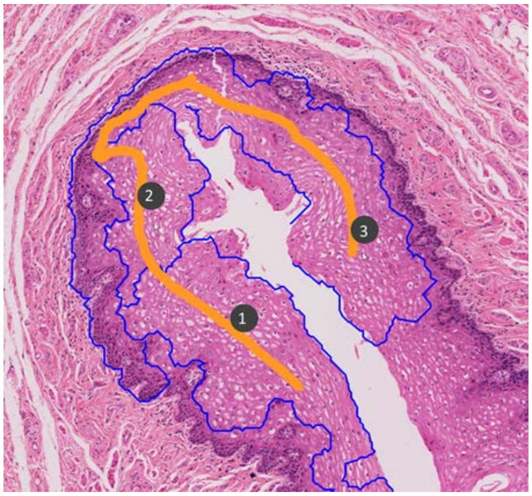

Figure 3. A prototype with a more rapid interaction for pathologists to annotate tissue.

3.2.3 Leaving traces

In an experimental project on sketching common objects collaboratively together with a neural network (Ha & Eck, 2017) the color of the strokes and the persistency differ depending on who made them. The strokes drawn by the system disappear and are replaced by other alternatives quite rapidly while the strokes drawn by the user become permanent. It makes the user feel in control of the sketch with the system only offering suggestions. However, the suggestions flash by a bit too quickly constantly leaving only the lines drawn by the user. There might be an argument for slowing down the flashes of suggestions in order for them to influence the user's drawing. One way would be to have the generated strokes leaving traces by remaining visible for a longer period of time as a way to provide a temporal context.

The core design idea is to leave traces to provide raw material for understanding and training the temporal behavior of the automation.

3.2.4 Judicious automation

In another project (Lindvall et al., 2018b) the aim was to create an aid for the pathologist to be put to practical use in the short term. Supporting primary diagnostics, while important, is also very challenging in terms of meeting the requirements for validation and robustness. By using high degrees of involvement and creating a cross-disciplinary team, an end-to-end aid was created that could reduce diagnosis times up to one week. By using regular human-centered design approaches and analyzing the sensemaking process in clinical decision making, the foraging loop (requesting supporting evidence) was identified as a candidate for automated predictions since a failure on behalf of the system would not have harmful consequences to patients. The result is an agentive system that automatically orders extra experiments with no user involvement. The design of the interactions for monitoring and providing corrections are currently a work in progress.

Not all decisions in a joint human-automation system are equally critical. Decisions that carry comparatively insignificant costs are candidates for automatic machine predictions even in situations where the performance of the automation is not strictly validated. This amounts to reducing the overall demands on the human user in terms of interaction and attention resource expenditure.

The core design idea is to consider how the scarce resource of human attention and effort should be spent, and specifically whether there are aspects that should be fully automated.

3.3 Using the curated collection to inform future design

For a curated collection to be a relevant transfer format, the knowledge it represents needs to be applicable to further design research. The scope of applicability, however, may depend on the nature of the knowledge contributions. We argue that our current curated collection, focusing on fine-grained interaction design, is applicable to a range of design situations where users are engaged in visual interaction with machine learning systems. To illustrate this point, we provide an example where we apply the insights from our curated collection to a different research-through-design case: a future system to support air traffic controllers in collaborating with agentive systems for planning flight routes. It might be worth mentioning that the designers working on this case have not been involved in making the prototypes selected as examples in the curated collection presented above.

We are currently creating possibilities to support air traffic controllers in exploring future air traffic scenarios by manipulating suggested flight routes using a stylus on a responsive system. The envisioned system "talks back" and guides the air traffic controller by way of visual inertia, making it more difficult to direct an aircraft towards a less desirable direction. Our goal is to provide a closer collaboration between the air traffic controller and automated systems by using the system as a channel to communicate. The system makes its future plans visible and editable (see figure 4). The air traffic controller can modify the suggestions while getting a sense of how desirable the change is depending on factors important to air traffic control and management, such as weather condition and surrounding aircraft.

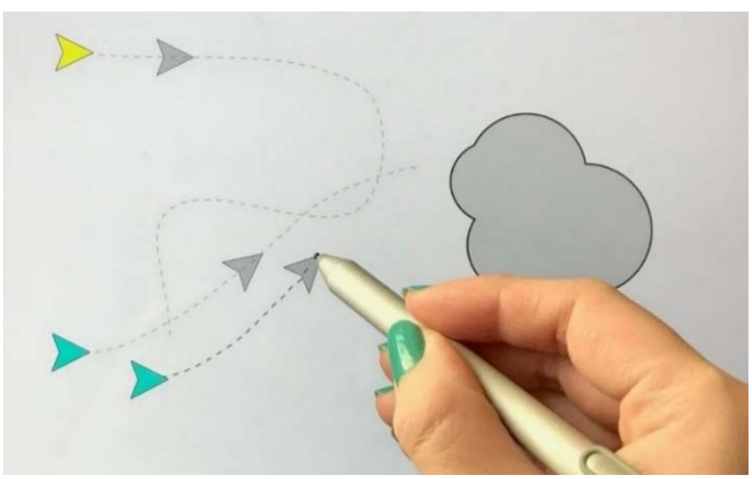

Figure 4. Planning future flows of aircraft through fine-grained interaction using a stylus. Through visual inertia the system can guide and "talk back" to the air traffic controller who is editing a suggested path.

The current prototype is based on notions of static pre-programmed automation and has no elements of learning. Based on the insights from the curated collection above, a future system for air traffic control looking to incorporate machine learning aspects could be developed along the following lines.

Design for preservation of visual context: in the current prototype visual inertia guides the air traffic controller to avoid directing aircraft toward less desirable paths. The visual inertia affects if, and how much, the aircraft lags behind the stylus and thus provides “fake” inertia. The inertia is a cue that does not add clutter to the interface. If the air traffic controller ignores the inertia, only then does the system mark out the possible risk for collision with a translucent circle.

Design for rapid real-time refinement: the main concept for the current prototype is already to support rapid real-time interaction. By using a stylus to interact with the system, immediate feedback in the form of visual inertia is continuously fed back to guide the path being drawn toward a desirable direction according to the system. In a refined version of the system the air traffic controller could use the pressure or speed when drawing with the stylus to indicate how much guidance is desirable at any given moment.

Design for leaving traces: in the current prototype, only one future path is visible for each aircraft and the path changes as soon as it is modified, leaving no trace of what the planned path was like before. The look of the dashed line also remains the same at all times. A possible refinement is to visualize multiple future paths as predicted suggestions and a distinction could be made to the path if it has been manually modified. Perhaps historical planned paths could even remain visible as shadows and fade away with time.

Design for judicious automation: many factors affect the decision-making process of planning aircraft routes and the robustness of a future system needs to be proven before it can be deployed in operative use. An obvious strategy for further development is to analyze the current air traffic control practice to identify less critical factors that could be prime candidates for machine learning.

4 Conclusions

Machine learning is arguably more dynamic than conventional design materials within interaction design. We aim to show how a research-through-design approach can yield transferable knowledge for future design of machine learning systems. Specifically, our focus of interest is in the fine-grained details of interaction, where we suggest (in the form of exemplars and annotations) that preservation of visual context, rapid real-time refinement, leaving traces and judicious automation can be significant for the interaction experience and the joint human-machine performance of an interactive machine learning system.

On the level of methodology, our curated collection approach combines the known intermediate-level knowledge forms of annotated portfolios and criticism. We argue that this combined approach has the potential to accommodate transferable knowledge in the context of research-through-design, as demonstrated in the application of the core ideas on fine-grained visual interaction to a case from a different domain.

Acknowledgements

This work was partially supported by the Swedish Transport Administration and the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Citations and References

Bardzell, J. (2011). Interaction criticism: An introduction to the practice. *Interacting with computers*, *23*(6), 604-621.

Dove, G., Halskov, K., Forlizzi, J., & Zimmerman, J. (2017, May). UX design innovation: Challenges for working with machine learning as a design material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 278-288). ACM.

Gaver, W. (2012, May). What should we expect from research through design?. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 937-946). ACM.

Gaver, B., & Bowers, J. (2012). Annotated portfolios. *interactions, 19*(4), 40-49.

Ha, D., & Eck, D. (2017). A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477*.

Holmquist, L. E. (2017). Intelligence on tap: Artificial intelligence as a new design material. *Interactions*, *24*(4), 28–33.

Lindvall, M., Molin, J., & Löwgren, J. (2018a). The importance of UX for machine teaching. Proc. *AAAI Spring Symposium Series*, pp. 407–410.

Lindvall, M., Skoglund, K., Rose, J., & Lundström, C. (2018b). Working methods for end-to-end development of automated assistants in pathology. Submitted in-review for *Proceedings of 5th Nordic Symposium on Digital Pathology*. Helsinki.

Löwgren, J. (2007). Pliability as an experiential quality: Exploring the aesthetics of interaction design. *Artifact, 1*(2), 85-95.

Löwgren, J. (2013). Annotated portfolios and other forms of intermediate-level knowledge. *interactions*, *20*(1), 30-34.

Löwgren, J. (2016). On the significance of making in interaction design research. *interactions, 23*(3), 26-33.

Yang, Q. 2017. The role of design in creating machine-learning-enhanced user experience. *The AAAI 2017 Spring Symposium on Designing the User Experience of Machine Learning Systems Technical Report SS-17-04* 406–411.

Zimmerman, J., & Forlizzi, J. (2008). The role of design artifacts in design theory construction. *Artifact, 2*(1), 41-45.

Zimmerman, J., Stolterman, E., & Forlizzi, J. (2010, August). An analysis and critique of Research through Design: towards a formalization of a research approach. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems* (pp. 310-319). ACM.