



HAL
open science

Immersive Situated Data Representations: Using the Physical World as a Visualization Canvas

Ambre Assor

► **To cite this version:**

Ambre Assor. Immersive Situated Data Representations: Using the Physical World as a Visualization Canvas. Human-Computer Interaction [cs.HC]. Université de bordeaux, 2023. English. NNT: . tel-04429392v1

HAL Id: tel-04429392

<https://hal.science/tel-04429392v1>

Submitted on 29 Jan 2024 (v1), last revised 31 Jan 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE PRÉSENTÉE
POUR OBTENIR LE GRADE DE
DOCTEUR
DE L'UNIVERSITÉ DE BORDEAUX

ECOLE DOCTORALE MATHÉMATIQUES ET INFORMATIQUE

SPÉCIALITÉ INFORMATIQUE

Par **Ambre ASSOR**

Représentation de Données Immersives et Situées: Utiliser le Monde
Physique comme Canevas de Visualisation

Sous la direction de : **Martin HACHET**
Co-directeur : **Pierre DRAGICEVIC**

Soutenance le 12 décembre 2023.

Membres du jury :

Mme. Céline Coutrix, Chargée de Recherche (HDR), CNRS, Rapporteur
M. Emmanuel Pietriga, Directeur de Recherche, Inria Saclay, Rapporteur
Mme. Fanny Chevalier, Associate Professor, University of Toronto, Examinatrice
M. Fabien Lotte, Directeur de Recherche, Inria Bordeaux, Examineur
M. Arnaud Prouzeau, Chargé de Recherche, Inria Bordeaux, Invité
M. Martin Hachet, Directeur de Recherche, Inria Bordeaux, Directeur de thèse
M. Pierre Dragicevic, Chargé de Recherche, Inria Bordeaux, Co-directeur de thèse

Représentation de Données Immersives et Situées : Utiliser le Monde Physique comme Canevas de Visualisation

Résumé :

Les affichages immersifs de type casques de réalité mixte permettent de visualiser du contenu numérique directement dans l'espace physique. Bien que cela rende possible l'affichage d'informations près d'objets, de personnes et d'espaces du monde réel auxquelles elles se rapportent, en faisant ainsi des visualisations dites situées, très peu de travaux en Réalité Mixte explorent les connexions entre monde physique et numérique pour la visualisation de données. Dans cette thèse, nous explorons ces connexions. Dans un premier temps, nous nous intéressons aux usages et effets de visualisations que nous nommons concrètes (c'est-à-dire dont la forme ressemble à des objets ou espaces physiques), affichées à l'échelle, en réalité augmentée et dans l'environnement direct des utilisateurs. En particulier, nous évaluons les réponses émotionnelles des sujets dans le cadre de données à caractère environnemental. Puis, nous explorons et évaluons des transitions animées entre visualisations concrètes et celles que l'on appelle abstraites (comme des diagrammes en barre ou des graphes), comme moyen de bénéficier des avantages de ces deux types de visualisation, en environnement immersif. Enfin, nous avons produit un espace de conception qui a pour but d'aider les concepteurs de visualisation à créer des visualisations situées lorsqu'un objet d'intérêt physique est non-visible temporairement ou de manière permanente. Les travaux présentés dans cette thèse contribuent à étendre les connaissances concernant la conception de visualisations immersives situées ainsi que d'identifier de nouvelles directions de recherches.

Mots-clés : Réalité Augmentée, Réalité Mixte, Visualisation, Interaction Humain-Machine

Immersive Situated Data Representation: Using the Physical World as a Visualization Canvas.

Abstract:

Immersive displays such as Mixed Reality headsets enable digital content to be viewed directly in the physical world. Although this makes it possible to display information close to real-world objects, people, and spaces to which it relates, thus making them so-called situated visualizations, very little work in Mixed Reality explores the connections between the physical and digital worlds for data visualization. In this thesis, we explore these connections. First, we investigate the usage and effects of visualizations that we call concrete (i.e., employing marks that resemble physical objects and/or that are positioned with respect to physical objects) displayed at scale, in augmented reality, and in the user's surroundings. In particular, we evaluate subjects' emotional responses to environmental data. Then, we explore and evaluate animated transitions between concrete and what we name abstract visualizations (such as bar charts or graphs) as a way of benefiting from the advantages of both types of visualizations in an immersive environment. Finally, we have produced a design space that aims to help visualization designers create situated visualizations when a physical object of interest is temporarily or permanently non-visible. The work presented in this thesis contributes to extending knowledge related to the design of immersive situated visualizations, as well as identifying new research directions in the field.

Keywords: Augmented Reality, Mixed Reality, Visualization, Human-Computer Interaction

Equipe POTIOC Inria/LABRI

UMR 5800 Université de Bordeaux, 33000 Bordeaux, France.

Contents

1	Introduction	16
1.1	Motivation	16
1.1.1	Visualize the unseen	16
1.1.2	Display within the world	18
1.2	Contributions of this Ph.D	20
1.2.1	Research questions	20
1.2.2	Contributions	21
2	Background	23
2.1	Mixed-Reality Visualizations	23
2.1.1	Information Visualization	23
2.1.2	Immersive Analytics	24
2.1.2.1	Definitions and concepts	24
2.1.2.2	Enabling Technologies	25
2.1.2.3	Challenges	26
2.2	[Immersive] Situated Data Representations	27
2.2.1	Definitions and Concepts	27
2.2.2	Enabling technologies	29
2.2.3	Existing Work in Immersive Situated Visualizations	30
2.2.3.1	Frameworks and Design Space Explorations	30
2.2.3.2	Situated Visualization Systems and Studies	31
2.2.4	Challenges	33
3	Designing Immersive Concrete Visualizations for Eco-Feedback	35
3.1	Introduction and contributions	37
3.2	Background	39
3.2.1	Concrete Visualizations	39
3.2.1.1	On-screen Concrete Visualizations	39

CONTENTS

3.2.1.2	Immersive Concrete Visualizations	41
3.2.2	Eco-Feedback Technologies	41
3.2.3	Eco-Feedback Studies	42
3.3	Illustrative scenarios	44
3.4	Prototypes and Initial Tests	46
3.4.1	Implementation	46
3.4.2	Early User Feedback	47
3.4.3	Lessons Learned	49
3.5	Designing ARWAVs	49
3.5.1	Type of representation	50
3.5.2	Realism of the representation	50
3.5.3	Animation and Interaction	51
3.5.4	Framing	52
3.5.5	AR Display Technologies	52
3.6	Experiment: are ARWAVs Emotionally More Engaging?	53
3.6.1	Experiment Design	53
3.6.1.1	Visualized Data	54
3.6.1.2	Physical setup	54
3.6.1.3	Presentation formats	55
3.6.1.4	Design	56
3.6.1.5	Procedure and Measurements	56
3.6.1.6	Hypotheses	57
3.6.1.7	Participants	57
3.6.2	Results	57
3.6.2.1	Planned Analysis	57
3.6.2.2	Additional Analyses	62
3.6.3	Discussion and Qualitative Feedback	63
3.6.4	Limitations of this Experiment	64
3.7	General Discussion	65
3.7.1	Exploring Presentation Modalities for ARWAVs in Varied Contexts	65
3.7.2	Exploring More Design Options for ARWAVs	65
3.7.3	Scale Issues with ARWAVs	66
3.7.4	Generalizations of ARWAVs	67
3.8	Conclusion	69
4	Binding Bridges Between Abstract and Concrete Immersive Visualizations With Animated Transitions	70

CONTENTS

4.1	Introduction and contributions	72
4.2	Background	73
4.2.1	Unit Charts	73
4.2.2	Hybrid Abstract/Concrete Systems	74
4.2.3	Animated transitions	75
4.3	Illustrative Scenarios	76
4.4	Designing animated transitions between concrete and abstract visualizations	78
4.4.1	A Design Space for Concrete and Abstract Visualizations	79
4.4.2	Initial Design Guidelines	81
4.5	Study	84
4.5.1	Choice of transitions for study	84
4.5.2	Choice of the task	85
4.5.3	Transfer function and duration	86
4.5.4	Hardware	88
4.5.5	Details of the task	88
4.5.6	Variables Measured	91
4.5.7	Pilot Study	91
4.5.8	Full Experiment	92
4.5.9	Results	93
4.5.9.1	Planned Analysis	93
4.5.9.2	Subjective Feedback and Exploratory Data Analysis	93
4.5.10	Discussion	96
4.5.10.1	On the design of animated transitions	96
4.5.10.2	On the implementation of animated transitions in AR	97
4.6	Conclusion	97
4.7	Future Directions	97
5	Handling real-world constraints for situated visualizations using MR	99
5.1	Introduction and contributions	100
5.2	Related Work	101
5.3	Terminology	102
5.3.1	Illustrative Scenarios	102
5.3.2	Basic Terminology	103
5.4	Handling Non-Visible Objects in Immersive Environments: A Survey	105
5.4.1	Methodology and Scope	105
5.4.2	Worlds-In-Miniature	106
5.4.3	Virtual Twins	107

CONTENTS

5.4.4	Physical Replica	108
5.4.5	Telepresence Systems	109
5.4.6	X-Ray Vision	110
5.4.7	Out-Of-View Pointing	111
5.4.8	Labeling	112
5.4.9	Virtual Mirrors	113
5.5	Design Space	114
5.5.1	Category: General	114
5.5.2	Category: Representation	117
5.5.3	Category: Visualization	120
5.6	Discussion	123
5.6.1	Use Cases and Design Trade-Offs	123
5.6.2	Link to Other Conceptual Frameworks	125
5.6.3	Descriptive and Generative Power	126
5.6.4	Techniques not Discussed	127
5.7	Conclusion	128
6	Conclusion and Perspectives	129
6.1	Summary	129
6.2	Future work and Perspectives	131
6.2.1	Exploratory Data Analysis with Concrete Visualizations	131
6.2.2	Immersive Data Storytelling with Concrete Visualizations	132
6.2.3	ARWAVs for a Sustainable Behaviour Change	133
6.2.3.1	Ecological cost of ARWAVs	133
6.2.3.2	ARWAVs for Behaviour Change	133
6.2.4	Engage Multiple and Diverse Display Mediums when Designing Immersive Situated Visualizations	134
6.2.5	More on the Study of Situatedness	136
6.2.6	Human Perception of the physical world with Immersive Situated Visualizations	136
6.2.7	Immersive Situated Art	137
6.2.8	Reflections on ethical stakeholders associated with Immersive Situated Visualizations	138
A	Appendices	159
A.1	Glossary	159
A.2	Supplementary material chapter 3	160
B	Experiment questionnaire	160
B	Videos	164

CONTENTS

B	Pre-Registration	164
A.3	Supplementary material chapter 4	164
C	Experiment Questionnaire	164
C	Videos	168
C	Pre-Registration	168
A.4	Supplementary material chapter 5	168
D	Complementary documents	168
D	Video	171

List of Figures

1.1	Brasilia’s area mentioned by Tom Hulme	17
1.2	Strava user-interface.	17
1.3	Three interfaces displaying route information.	18
1.4	A 3D representation of New York City’s daily carbon dioxide emissions.	19
2.1	Visualization model presented in Mackinlay and Card’s work [9]	24
2.2	Reality-Virtuality Continuum by Milgram and Kishino [224]	25
2.3	Commonly-used AR-enabling devices. (a) Ikea app https://apps.apple.com/fr/app/ikea/id1452164827 . (b) Hololens 2 ad https://cdn-dynmedia-1.microsoft.com/is/image/microsoftcorp/gldn-HIGHLIGHT-hololens-2-4000x2100:VP2-859x540	26
2.4	Visualization pipeline extended to the physical world [322].	28
2.5	A physicalization: the robotic plant [170].	29
2.6	Two AR-based situated visualizations systems.	30
2.7	Three AR systems providing situated visualizations.	32
2.8	HydrogenAR [318].	32
2.9	AR condition of Guarese et al.’s study [158].	33
3.1	ARWAVs in a cafeteria.	38
3.2	Examples of physical waste accumulations.	39
3.3	A concrete visualization of CO ₂	40
3.4	Illustrative examples of ARWAVs we prototyped.	44
3.5	Two of the three experiment conditions on ARWAVs.	55
3.6	ARWAVs experiment’s demographics.	58
3.7	Overall emotion intensity for each of the three techniques.	58
3.8	Within-subject differences in overall emotion intensity.	59
3.9	Overall intensity of positive and negative emotions.	59
3.10	Reported intensity for all 20 emotions of the PANAS questionnaire.	60
3.11	Technique rankings.	62

LIST OF FIGURES

3.12 Mean within-subject emotions' intensity differences between techniques.	62
3.13 Miniature view of an ARWAV.	67
3.14 Photomontage of a large-scale ARWAV.	68
4.1 Stages of an animated transition between an abstract and concrete visualization.	72
4.2 Abstract and concrete representations of crowd flows data.	77
4.3 Four visualizations of population densities in VR.	78
4.4 Our design space on concrete and abstract visualizations.1.	79
4.5 Our design space on concrete and abstract visualizations.2.	80
4.6 The seven animated transitions evaluated in the pilot.	85
4.7 Design guidelines.	86
4.8 A user performing the task.	87
4.9 Animated transitions directions.	87
4.10 Two speed transfer-functions.	88
4.11 Example of a Forward direction trial.	89
4.12 Example of a Backward direction trial.	90
4.13 One-staged (VS) vs. two-staged (V, S) transformation.	92
4.14 ItemsCloser for all conditions.	93
4.15 Subjective ratings by users.	94
4.16 ItemsCloser, broken down by direction.	94
4.17 The time (i.e., duration) of each trial.	95
4.18 Error rates.	95
4.19 Error rate by max angular speed of target.	95
5.1 Diagram summarizing our key terminology.	103
5.2 A WIM.	106
5.3 A virtual twin of a smart factory.	107
5.4 Tangible interface for visualizing the energy consumption of a university [126].	108
5.5 Example of a telexistence/telepresence system.	109
5.6 X-ray vision for surgery.	110
5.7 Out-of-view pointing in SidebARs [278].	111
5.8 The stereo panorama of the Luxor Temple.	112
5.9 Example of a virtual mirror.	113
5.10 Illustration of the non-visibility support dimension.	115
5.11 Illustration of the interactivity dimension.	115
5.12 Illustration of the interaction realism dimension.	116
5.13 Illustration of the actionnability dimension.	117

LIST OF FIGURES

5.14	Illustration of the representation fidelity dimension.	117
5.15	Illustration of the refresh rate dimension.	118
5.16	Illustration of the surroundings inclusion dimension.	119
5.17	Illustration of the medium of the representation/visualization dimension.	119
5.18	Illustration of the integration of the representation/visualization dimension.	120
5.19	Illustration of the location information dimension.	121
5.20	Illustration of the visualization richness dimension.	121
6.1	Be · Aware project process.	134
6.2	Two physicalizations to convey CO2 emissions. via http://dataphys.org/list/	135
6.3	PiperZy’s work on AR-art and digital fashion.	137
A.1	Experiment questionnaire from chapter 3, Part 1.	161
A.2	Experiment questionnaire from chapter 3, Part 2	162
A.3	Experiment questionnaire from chapter 3, Part 3.	163
A.4	Pre-questionnaire, filled by participants before performing the tracking task.	165
A.5	Post-questionnaire, filled by participants after performing the tracking task, Part 1/2.	166
A.6	Post-questionnaire, filled by participants after performing the tracking task, Part 2/2.	167
A.7	Techniques handling non-visible referents including visualizations, Part 1.	169
A.8	Techniques handling non-visible referents including visualizations, Part 2.	169
A.9	Techniques handling non-visible referents including visualizations, Part 3.	169
A.10	List of selected papers.	170

Résumé long en français

Contexte Les humains laissent continuellement des traces de leur passage sur le monde physique, en créant par la même occasion un espace riche en informations. Ce phénomène s'illustre autour de la notion de *ligne du désir* [13], qui fait référence à un chemin tracé progressivement par l'érosion provoquée par le passage répété de piétons, de cyclistes ou d'animaux. Ce genre d'informations est précieuse pour les urbanistes qui s'en inspirent pour produire des alternatives efficaces à certains chemins conçus en amont, sans se baser sur la pratique réelle des usagers, se révélants même parfois mortels pour ces derniers. Il existe bien d'autres exemples de traces laissées par les humains sur le monde comme les traces de magnésies sur les murs d'escalade qui informent de routes utilisées par les grimpeurs, potentiellement les plus adaptées, ou encore les traces d'usages présentes sur des touches de claviers d'ordinateurs qui peuvent jusqu'à informer sur la langue parlée par l'utilisateur. Cependant, bien souvent, ces traces sont imperceptibles. En effet, même en marchant un nombre de fois incalculable sur une route faite d'asphalt, un piéton aura très peu de chances d'informer de son passage. Face à cela, la technologie peut être utilisée pour visualiser des informations en contexte: par exemple, un coureur pourra observer le chemin parcouru lors de sa course sur une carte via une application comme Stava [11], sur son écran de smartphone. Les interfaces de Réalité Augmenté (RA), qui permettent d'afficher des informations numériques directement dans le monde physique, semblent être de bonnes candidates pour afficher des données à intégrées dans l'espace, pourquoi pas en utilisant des représentations qui iraient jusqu'à imiter les traces que nous y laissons. Par ailleurs, elles ont connu un essor considérable cette dernière décennie et se voient de plus en plus employées pour des usages non-experts. Cependant l'usage de celles-ci pose de nombreuses questions: Comment les informations peuvent-elles être visualisées d'une manière à favoriser leur compréhension par un large éventail d'utilisateurs · trices ? Est-il vraiment avantageux d'utiliser une telle interface à ces fins ? Quelles sont les limites à l'utilisation du monde physique comme support de visualisation et comment les surmonter ? Dans cette thèse, nous tenterons d'apporter des éléments de réponse à ces questions. Par exemple, nous explorerons de nouvelles manières de représenter les données à l'image de notre expérience du monde physique et nous investiguerons des manières de bénéficier de l'avantage de ces nouvelles visualisations sans perdre les avantages de visualisations plus classiques. Nous étudierons également comment certains travaux de recherche existants en Réalité Mixte (MR) permettraient d'aider

à surmonter certains obstacles liés à l'utilisation du monde physique comme support de visualisation.

Démarche L'objectif de cette thèse est de créer des connexions entre les mondes physique et numérique en utilisant des affichages immersifs pour la visualisation de données. Les effets des visualisations situées immersives sur les humains est encore relativement peu connue. A l'intersection des domaines de recherches de la Visualisation et de l'Interaction Humain-Machine, ce travail a pour but d'étendre les travaux de recherche existants selon les trois axes d'étude suivants :

- **Usage et effets d'encodages concrets en RA.** Je m'intéresse l'opportunité de concevoir des visualisations qui font appel à notre familiarité du monde physique. Par exemple, nous sommes habitués à manipuler des objets de certaines formes, tailles, présents dans certains contextes. Visualiser les données de cette manière pourrait aider à la compréhension de métriques difficiles à saisir. Je propose d'explorer comment ces représentations *imaginées* peuvent s'appliquer au domaine de l'éco-feedback. Je présente également certains effets de ces visualisations sur l'Humain.
- **Création et effets de transitions entre visualisations classiques et visualisations concrètes en environnement immersif.** Les représentations de données que nous visualisons via des interfaces classiques (comme sur un écran d'ordinateur) sont très différentes de représentations introduites dans l'axe de recherche précédent, alors même qu'elles peuvent représenter les mêmes données. Ces représentations que nous nommons abstraites (comme des graphs ou des diagrammes en barre) présentent de nombreux avantages comme la place qu'elles occupent dans l'espace ou la possibilité de lire des valeurs numériques précises. Je m'intéresse à la façon dont on peut connecter les représentations abstraites et concrètes par le biais de transitions animées qui permettent de profiter des avantages des deux types de visualisation.
- **Proposition de solutions RM permettant de surmonter les contraintes du monde physique pour concevoir des visualisations situées.** Les représentations mentionnées ci-dessus présentent des avantages à exister à l'endroit de l'espace auxquelles elles se rapportent. Les afficher de cette manière pourrait faciliter leur compréhension et potentiellement aider à la prise de décision. Les recherches passées en Réalité Mixte présentent des solutions qui permettraient d'outrepasser les obstacles qui viennent du monde physique (par exemple le fait qu'un objet puisse être temporairement ou de manière permanente occulté de la vue de l'utilisateur · trice). Pour aborder ce troisième axe, une approche est de répertorier ces techniques de XR, de les classifier, et d'expliquer comment elles s'adressent aux problématiques liées aux représentations situées ou intégrées dans l'espace physique.

Ces axes contribuent à la recherche sur les visualisations immersives situées selon les points suivants:

1. Étendre l'exploration de **nouveaux éléments de conceptions qui utilisent effectivement le monde physique comme espace de visualisation**. Ceci est nécessaire pour diversifier les directions de recherche sur le sujet. Cela permet de faire émerger de nouveaux usages, qui pourraient se révéler très puissants pour les utilisateurs concernés. Ces retombées positives sont importantes à identifier pour justifier l'usage de ces technologies qui présentent un impact écologique certain.
2. Permettre une **meilleure compréhension des effets de certains éléments de conception liés aux visualisations immersives**. Dans la même idée que le point précédent, il est impératif d'identifier les éléments de conception et les applications qui donnent effectivement plus de pouvoir à l'utilisateur · trice, qui leur permettrait de comprendre des données relatives au monde physique, prendre des décisions éclairés et potentiellement modifier leur comportement *dans le bon sens*. Arriver à cela demande de mener des études utilisateurs · trices pour évaluer les paramètres qui pavent le champs des possibles, déterminant ainsi les éléments de conception à pousser et ceux à écarter.
3. Avancer dans l'**implémentation pratique de visualisations [immersives] situées**. Même si la recherche avance sur l'utilisation des technologies immersives, les conditions expérimentales et les prototypes produits sont parfois en décalage avec les conditions réelles dans lesquelles ces systèmes pourront être déployés. En conséquence, il est essentiel de s'adresser à ce décalage et proposer des solutions pour l'implémentation pratique de ces scénarios.

Bien entendu, ces trois points offrent une perspective bien ambitieuse. Les travaux que je présente ici les adressent très partiellement et je les explicite dans le prochain paragraphe.

Résultats Dans une première partie, je présente une nouvelle forme de représentation holographique pensée pour exister dans le monde physique. En effet, ces représentations sont délibérément alignées sur l'échelle du monde physique et sont affichées dans l'environnement direct de l'utilisateur · trices. Cette idée est largement inspirée de travaux sur l'*Échelle Concrète*¹ [94] qui s'attaque à représenter des unités peu familières d'une manière plus compréhensible. Par exemple, la quantité de sucre dans une boisson gazeuse peut être représentée par une pile de morceaux de sucre à côté de la boisson, tandis que des quantités plus importantes peuvent être représentées par des contenants familiers tels que des camions, remplis de morceaux de sucre. Cette approche a été étendue par Lee et al. [201] en utilisant des scènes de Réalité Virtuelle pour représenter ces mesures de manière immersive et à l'échelle appropriée, à côté de l'utilisateur · trice. Cependant, aucun travail antérieur n'a envisagé l'utilisation de la réalité augmentée, où nous trouvons l'ajout distinctif de l'intégration d'échelles physiques et réelle dans l'équation. Cela signifie que les données sont représentées à l'échelle appropriée, dans l'environnement physique de l'utilisateur · trice qui contient des indices métriques tels que des objets et des espaces familiers

¹*traduit de l'anglais Concrete Scales

(par exemple, la salle de bain de l'utilisateur · trice, des chaises, des tables, etc.). Nous avons identifié une possibilité d'application dans le domaine de la visualisation de données à caractère écologique. En effet, la société moderne est organisée de telle manière que les déchets soient rapidement soustraits à nos sens pour des raisons de confort. Cela nous empêche d'être directement confrontés au volume réel des déchets que nous produisons, sauf lorsque les systèmes d'acheminement tombent en panne et créent des affichages d'information naturels (par exemple, une grève des ordures peut rendre visible le volume de plusieurs jours de déchets produit par un immeuble, une baignoire bouchée peut révéler la quantité d'eau utilisée pour une douche). Nous nous inspirons de ces affichages naturels et introduisons une classe de visualisations situées que nous appelons Visualisations d'Accumulation de Déchets en Réalité Augmentée² ou ARwavs, qui sont des représentations concrètes de données relatives à des déchets intégrées dans l'environnement familier de l'utilisateur · trice. Nous avons implémentés des exemples d'ARwavs et en avons fait la démonstration lors d'une première session de retour utilisateur avec des collègues ainsi que lors d'une exposition technologique grand public. Nous présentons des considérations générales relatives à la conception des ARwavs. Enfin, nous avons mené une étude auprès de 20 participants suggérant que les ARwavs suscitent des réactions émotionnelles plus fortes que les visualisations non immersives d'accumulation de déchets et les simples chiffres.

Nous avons montré que ces représentations que nous appelons concrètes présentent des avantages ; cependant, les représentations abstraites (telles que les diagrammes à barres ou les graphiques) rassemblent également des avantages, distincts. En effet, nous avons l'habitude de les utiliser pour communiquer et partager des données, ou pour lire des valeurs numériques précises par exemple. Nous avons donc adopté une approche pour bénéficier de leurs avantages complémentaires. Plus précisément, nous avons choisi d'explorer la manière de passer en douceur des représentations concrètes aux représentations abstraites et inversement dans des environnements immersifs. Nous avons établie un espace de conception pour les visualisations abstraites et concrètes, pour pouvoir les comparer selon un ensemble commun de dimensions. Les transitions animées sont des chemins à travers l'espace de conception. Ce dernier est défini par quatre axes, chacun correspondant à une transformation différente : Marques, Échelle, Intégration et Vue. Les marques peuvent être concrètes ou abstraites : dans le cas d'une visualisation concrète, chaque marque visuelle fait référence à un objet, tandis que dans le cas des visualisations abstraites, qui représentent généralement des agrégations, une marque visuelle indique une valeur quantitative pour un attribut. Ces marques peuvent exister à une échelle réduite ou physique. Par exemple, l'échelle physique d'une représentation concrète est celle de l'objet physique représenté (par exemple, si la marque concrète est un livre, la marque aura la taille d'un vrai livre). Une échelle réduite représenterait une fraction de l'échelle de l'objet physique. Si les positions des marques de la visualisation dérivent des caractéristiques du monde physique, l'intégration est concrète, sinon elle est abstraite. En outre, la vue peut être à la première ou à la troisième personne (respectivement, l'utilisateur · trice est entouré d'éléments graphiques

²Traduit de l'anglais *Augmented Reality Waste Accumulation Visualization*.

ou regarde une visualisation placée devant lui/elle). Nous envisageons différentes manières de concevoir des transitions animées en combinant et ordonnant les transformations le long de ces axes. Dans une expérience contrôlée menée en Réalité Virtuelle avec 16 participants, nous avons comparé quatre types de transitions animées et trouvé des preuves quantitatives et qualitatives de la supériorité d'une approche de combinaison de transformations spécifique par rapport à l'application simultanée de toutes les transformations. Nous proposons également un ensemble provisoire de lignes directrices en matière de conception pour de futurs travaux adressant les mêmes problématiques.

Enfin, nous étudions les techniques de Réalité Mixte existantes qui peuvent aider les concepteurs de visualisation à utiliser le monde physique pour créer des visualisations situées. Plus précisément, nous abordons la question de la gestion des objets d'intérêt non-visibles de manière temporaire ou permanente (par exemple, lorsqu'ils sont occultés de la vue des utilisateurs · trices). Nous avons ensuite classé ces travaux en huit familles de techniques capables de gérer la non-visibilité des objets physiques et de préciser où se situe leur potentiel usage dans le cadre des représentations intégrées (c'est-à-dire des représentations de données profondément intégrées dans l'espace physique). L'une d'entre elles, par exemple, est la vision à rayons X qui émule la capacité surhumaine fictive de voir à travers les surfaces occultantes. Lors de l'exécution d'une tâche impliquant des visualisations situées, si un obstacle bloque la vue de l'utilisateur · trice, une vision à rayons X pourrait s'activer pour maintenir le lien visuel entre l'utilisateur · trice et l'objet physique et la visualisation associée afin de permettre l'achèvement de la tâche sans interruption(s). Nous dérivons ensuite un espace de conception avec 14 dimensions divisées en trois catégories : une générale, une liée à la représentation de l'objet d'intérêt et une associée à la visualisation liée à l'objet. Les dimensions générales englobent des aspects tels que l'interactivité de la technique (par exemple, la mesure dans laquelle l'utilisateur · trice peut interagir avec la représentation ou la visualisation). Les dimensions de représentation concernent des éléments tels que la fidélité de la représentation (comme le niveau de détail avec lequel le référent physique est représenté) ou l'intégration de la représentation (par exemple, dans quelle mesure la représentation semble faire partie de l'environnement de l'utilisateur · trice). Les dimensions de la visualisation sont liées à des éléments tels que la richesse de la visualisation (par exemple, la quantité de données visualisées sur le référent physique). Ce travail permet d'examiner divers compromis, défis et possibilités de recherche future dans ce domaine, et principalement, nous envisageons qu'il sera utile aux concepteurs de visualisations situées dans leur prise en charge des référents non-visibles.

Conclusion Pour résumer, ces trois travaux sur lesquels nous avons travaillé contribuent à faire progresser la recherche sur les visualisations immersives situées. En particulier, nous étudions de nouvelles visualisations qui se calquent sur l'apparence d'objets réels, nous essayons de construire des ponts entre ces visualisations et les visualisations classiques et nous construisons un cadre théorique qui vise à soutenir la mise en œuvre de visualisations situées en conditions réelles. Ces travaux permettent également aux chercheurs et aux concepteurs en visualisation et en interaction humain-machine d'identifier des pistes de

recherche possibles dans ce domaine.

Par ailleurs, les expériences utilisateur que nous avons menées en laboratoire et les espaces conceptuels ainsi construits ont bien sûr leur limites. Par exemple, dans la première partie nous avons évalué la réponse émotionnelle des utilisateurs · trices mais ils serait judicieux de mener des études complémentaires pour préciser la cause de ces émotions (qui pourraient être dues au système d'affichage aussi bien qu'aux informations présentées par exemple). Aussi, les émotions sont reportés par les utilisateurs · trices, ce qui est un choix justifiable mais qui bénéficierait d'une analyse complémentaire menée via une expérience qui mesurerait ces émotions. Autrement, les prototypes que l'on a explorés et évalués pourraient largement être enrichis en terme d'interaction et de réalisme de rendu. De plus ils sont limitées au cas des données liées à l'impact environnemental des utilisateurs · trices et il serait intéressant d'implémenter et de tester des prototypes diversifiés. Dans la seconde partie, le fait de s'intéresser aux transitions animées est un choix en soit; en effet il serait envisageable de montrer des visualisations abstraites et concrètes côte à côte, dans le même espace immersif, et dessiner des liens cohérents entre les visualisations. Notre analyse prend en compte les deux directions de transitions à la fois (de concret à abstrait et abstrait à concret) alors que ces deux directions peuvent être envisagées indépendamment. Il existe également de nombreux éléments visuels qui bénéficierait à ce genre de transitions, comme l'utilisation de la transparence pour certains objets. Cette première étude sur les transitions animées entre représentations concrètes et abstraites en milieu immersif a été menée en Réalité Virtuelle pour des raisons pratiques. Comme les représentations concrètes présentent des avantages prometteurs à exister dans le monde physique, une prochaine étape de recherche assez immédiate serait l'étude de telles transitions en Réalité Augmentée. Aussi, nous avons choisi d'évaluer ces transitions via une tâche de suivi visuel, bas niveau (les participant · es devaient suivre du regard un élément de la visualisation concrète lors de sa transition vers une visualisation abstraite et inversement). Il serait très intéressant de proposer l'évaluation de transitions entre visualisations concrètes et abstraites via un tâche plus haut niveau qui demande l'étude d'un dataset particulier. Concernant la dernière partie, notre espace conceptuel est loin d'englober tous les aspects relatifs à la conception de visualisations situées qui peuvent gérer des référents non-visibles, et il ne couvre pas non plus toutes les conceptions possibles. De nombreuses considérations entrent en ligne de compte dans la conception de tels systèmes interactifs, et un seul espace de conception ne peut pas tous les saisir.

Les limites de nos travaux sont des opportunités de nouvelles recherches, maintenant identifiées. Les perspectives qui en découlent sont nombreuses autant pour concevoir des visualisations immersives situées que pour cerner leurs bénéfices pour l'utilisateur · trice.

Chapter 1

Introduction

1.1 Motivation

1.1.1 Visualize the unseen

Humans continuously leave traces in their physical environment, making it a significantly rich information display. The concept of *desire paths* shows an illustrative example of this phenomenon. A desire path refers to an unplanned trail that emerges due to mechanical erosion caused by human or animal traffic [2]. In a Ted Talk from 2016 [13], Tom Hulme presents a notable case in Brasilia, where an area initially presumed to be unused for pedestrian crossing became filled with desire paths over time (see Figure 1.1). The urban planners had not anticipated such foot traffic, but residents created shortcuts, and natural paths emerged. However, these paths pose significant risks as they intersect numerous roads, leading to a high number of pedestrian accidents. These are hints that this area was not designed according to the residents' needs. Hulme identifies cost and time savings as motivating factors that drive individuals to discover alternative, more efficient routes, often referred to as desired paths or shortcuts. He emphasizes the importance of designers observing and understanding these traces left by people in order to inform design choices effectively.

Numerous other examples illustrate how the physical world is a substrate of information, capturing and revealing human activities. For instance, magnesia traces on an indoor climbing wall indicate the most taken routes by climbers, while the wear and tear of a keyboard's keys highlight the most frequently used keys. In some cases, clues regarding human behavior extend beyond visual traces and engage other senses. For instance, the dry soil of a flower pot indicates that the plant has not been watered recently, while a foul smell emanating from the kitchen may suggest that someone has forgotten to remove the




Figure 1.1 Brasilia's area mentioned by Tom Hulme

trash. These examples demonstrate how the physical environment holds valuable information that can be interpreted and utilized to better understand human behaviors and inform decision-making processes.

However, in many instances, these traces may be invisible or imperceptible. While walking in a field might leave some visible traces, walking countless times on a floor made of asphalt will leave it unchanged. To address this limitation, humans have turned to technology as a means to capture and preserve these traces, enabling them to learn from their own behaviors. For instance, runners can utilize smartphone applications like Strava to record their runs and review them later on a map (see Figure 1.2). By incorporating additional technology such as a heart rate monitor, the app can establish correlations between heart rate and the route, providing deeper analytical insights.

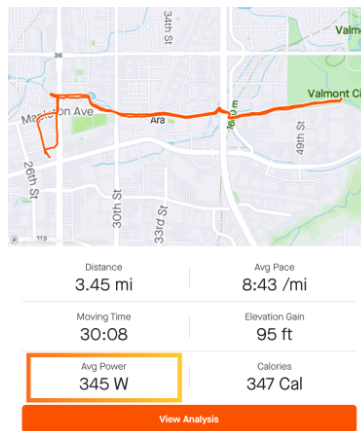


Figure 1.2 Strava user-interface. The orange line represents the run itinerary.

Augmented Reality (i.e., interfaces enabling the display of 3D virtual objects into a 3D real environment in real-time, also referred as AR) presents an exciting opportunity to integrate information directly into the physical environment. Data we record about ourselves can be displayed within the physical space,




Figure 1.3 Three interfaces displaying route information. (a) Paper map. (b) A GPS showing directions in Paris. The right side of the interface mimic physical direction panels that may be non-visible. Cutlery symbols indicate restorative areas. (c) Panasonic and Phiar simulate next-gen in-vehicle augmented navigation solutions.

resembling our natural traces. With the aid of AR glasses, for example, a floor made of asphalt could exhibit virtual erosion patterns corresponding to the areas we have walked on. In this thesis, we heavily drew inspiration from the idea of natural traces and consider *Mixed Reality* (mixing physical and digital objects, also referred as MR) as a medium in order to present instructive and prescient data where it makes sense in the real world, aligning with our experience of it.

1.1.2 Display within the world

We already enhance the physical world by adding physical and digital visualizations at specific locations. A good example is traffic signs: way-finding is facilitated by direction signs that originate from the early 20th century and consist of physical poles strategically positioned to ensure maximum visibility to drivers in the respective area. Their placement along the road is carefully selected, either preceding decision points or serving as location reminders for users. Furthermore, the road infrastructure is supplemented with traffic lights, which are also substantial physical poles but incorporate digital components, such as dynamic color lights. People rely on cues that can be combined by other information support, such as maps (Figure 1.3a).

In more recent times, fully digital interfaces, exemplified by GPS systems (see Figure 1.3b), have emerged to aid navigation and provide continuous guidance throughout journeys. These interfaces present information in distinctly different manners, taking into account its constraints (i.e., screen size). GPS also reuses and adapts existing representations that users might be familiar with. Consequently, features like map views, comprehensive overviews, and universally recognizable icons (some are new, such as police symbols denoting speed cameras; some are known from traffic panels, such as cutlery emblem signifying rest areas) can all exist on the same medium. However, an important point is that GPSs are not substitutes but complementary to traffic panels.

This Ph.D thesis is titled *Immersive Situated Data Representation: Using the Physical World*

as a Visualization Canvas. This idiomatic expression is very close to one of Nieman’s formulations from 2011 documented in the ADS-VIS proceedings. In the paper entitled “*Concrete vs abstract visualization: The real world as a canvas for data visualization*”, he argues that the real world and our *embodied* experiences (refer to section A.1) are often underused in data visualization. He illustrates this concept by introducing “Carbon Quilt” (see Figure 1.4), a novel visualization technique that co-opts the world itself in visualizing carbon emissions data. The latter is a great example of a *visceral representation*, i.e., a representation that restores the basic understanding of units and measures that are often abstracted away.




Figure 1.4 New York City’s daily carbon dioxide emissions as one-tonne spheres, represented among a 3D representation of the aforementioned city, from <http://www.carbonvisuals.com/>. This representation is a 3D model intended to be displayed in 2D on a screen.

More than ten years ago, he mentioned that Augmented Reality presents great potential into bringing real-world context in data representation to make them more engaging and informative. Interestingly, the last decade has seen significant advancements in head-mounted display development. Augmented Reality interfaces are not only used by experts anymore but also broadly invested by non-experts for various purposes. Companies like Meta, Microsoft, Magic Leap, and Apple recently developed sophisticated AR and MR devices, providing immersive experiences with high-resolution displays, advanced tracking, and spatial mapping capabilities. Along the way, development tools widen accessibility to build immersive experiences and to study them. Indeed, its advents pose analogous inquiries than with GPS birth in a world of maps (see some simulation Figure 1.3c), and numerous long-term questions arise: How can information be represented in a manner that promotes comprehension among a broad range of users, minimizing cognitive load? How can the limitations inherent to this new interface be identified and effectively managed? On top of that, is it really beneficial to use such an interface for that purpose? Would these interface and/or holographic representations replace or complement existing ones?

In this thesis, I will focus on some of them by exploring how to provide Augmented Reality experiences

that include our customs of visualizing data while harnessing some of the core benefits of AR. I will also present studies of some specific past AR developments that appear to be great strategies to cope with the challenge of visualizing data on the 3D dynamic canvas that is the real world. I will also consider *Virtual Reality* environments to facilitate investigations intended to exist in AR.

1.2 Contributions of this Ph.D

I will refer as *we* to include my supervisors/collaborators on these projects.

1.2.1 Research questions

Overall, the aim of this Ph.D thesis is to create connections between the physical and the digital world by using immersive technologies to visualize information. The effects of immersive situated visualizations (i.e., immersive visualizations displayed near the physical object, person, or space they refer to in the physical world [322]) on human performance are still relatively unknown, and although the use of immersive technologies is growing, there is an urgent need to understand their actual benefits. At the intersection of Visualization and Human-Computer Interactions, our work aims to enrich previous research on situated immersive visualizations following the three research axes:

- **Usage and effects of concrete visualizations in AR.** Using designs that leverage our customs of interacting with the physical world raises potential benefits. For instance, we are familiar with physical objects, shapes, scales, and contexts. Designs that resemble these objects may efficiently convey metrics that are difficult to grasp and that are usually abstractly represented. I propose to investigate how these *illustrative* representations may apply to a specific visualization field and certain of their effects.
- **Creation and effects of transitions between classic and concrete encodings in immersive environments.** Data representations we observe on classic interfaces (i.e., computer screen) are very different from the representations introduced in the previous research axis, yet they may convey the same data. These visualizations that we call abstract (e.g., graphs or bar charts) have advantages such as the space they occupy or their enabling of quick value judging. I present a way to connect concrete and abstract visualizations through animated transitions in order to harness the benefits of both types of visualizations.
- **Conveyance of MR-based solutions to overcome physical-world constraints for situated visualizations.** The different representations listed above would advantageously be displayed near the physical object they refer to. As a matter of fact, it may facilitate data understanding, decision

making and behavior change. Past development of Mixed Reality systems can provide solutions to some obstacles coming from the physical world (e.g., occlusion of the object of interest from the user's view). This third axis proposes to inventor them and clarify where their contribution falls in the scope of embedded representations, i.e., representations of data that are deeply integrated with the physical space.

The following components can capture the objective of our work on immersive situated visualizations:

1. Expanding the exploration of **novel designs that effectively utilize the real world as a platform for visual representation**. This is necessary to diversify research directions on the subject as a strategy to discover new and potentially powerful applications designed for non-expert usage. Trade-offs that accompany the deployment of Augmented Reality systems would be identified in a better way. In fact, the effects of AR on human cognition are relatively unknown, as well as their definite and little-discussed ecological impact.
2. Enabling a **better understanding of the effect of specific immersive visualization design parameters on users**. In line with the previous point, it is imperative to identify designs and applications that are meaningful and empower individuals to gain an understanding of data related to the physical world, make informed decisions, and potentially modify their behavior in a *good* way. Achieving this goal entails laboratory evaluations to explore the vast pool of possibilities, determining which designs to support and which to discard.
3. Advancing the **practical implementation of immersive situated visualization solutions**. Although considerable progress has been made in the MR research field, experimental conditions and implemented prototypes usually diverge from the real-world context within which these innovations will be deployed. Consequently, bridging this gap and proposing solutions for practical implementation scenarios is essential.

Of course, these driving elements are ambitious and our contributions are merely small elements that address it. I list them in the following section.

1.2.2 Contributions

- **chapter 2: Related work**. In this chapter, I introduce key concepts profusely used in this thesis manuscript (i.e., immersive analytics and situated visualization). I present existing systems and studies and enable hardware that may ground this theoretical background into more tangible scenarios. This chapter also presents a wide range of challenges faced by the research area to give a global context to the reader.

- **chapter 3: Designing Immersive Concrete Visualizations For Eco-Feedback.** In this chapter, I introduce a novel form of holographic representation designed to be attuned to the physical world. We took inspiration from the physical world and dived into a sub-category of interfaces technology that “*provides feedback on individual or group behaviors with a goal of reducing environmental impact*” [139]. We convey design recommendations based on lessons learned during early feedback sessions, one with colleagues and one with a large audience at a tech exhibition event. We also present the user study results suggesting that these representations yield more emotions than non-AR ones.
- **chapter 4: Binding Bridges Between Abstract and Concrete Immersive Visualizations With Animated Transitions:** In this chapter, I introduce our strategy to design animated transitions between abstract and concrete visualizations. To frame our approach, we build a design space for concrete and abstract visualizations with four axes (Marks, Embedding, View, and Scale), and animated transitions are different staging and ordering transformations along these axes. I also present the results of a controlled experiment in VR, yielding quantitative and qualitative evidence for the superiority of a specific staging approach over the simultaneous application of all transformations. We also propose design guidelines for immersive animated transitions between concrete and abstract representations.
- **chapter 5: Handling real-world constraints for situated visualizations using MR:** In this chapter, I introduce a framework that aims at supporting designers to build visualization related to an object, person, or location that may be temporarily or permanently non-visible. To do so, we reviewed existing MR techniques that support the non-visibility of 3D objects, categorize them, and show how they can complement existing systems. This leads to a 14-dimensional design space, useful for examining trade-offs, challenges, and opportunities in this research area.
- **chapter 6: Discussion, future work, and conclusion:** In this chapter, I review the main contributions of this thesis in the research field of immersive situated visualizations and highlight the significance of its findings. Among all, I clearly answer how we tackled the three research questions presented before. I also clarify its limitations and introduce very next steps for future research opportunities and challenges.

Along the thesis, some words are written in *italic bold* when they are first used. These are defined in the Glossary section A.1.

Chapter 2

Background

In this chapter, I give a brief and general overview of the context of our research. In subsequent chapters, I will present elements that are more specific to the contributions they hold.

As previously discussed, this thesis falls in the intersection of Human-Computer Interaction (HCI) and Visualization, explicitly focusing on the interaction between humans and Mixed-Reality interfaces providing situated visualizations. I will begin by giving context around mixed-reality visualization research, starting with visualization and then elaborating on immersive visualization concepts, enabling technologies, and general challenges. Then, I will go into more detail about situated visualizations, once more presenting general concepts, before moving on to immersive systems that offer this particular type of visualization. I will focus on presenting a representative range of immersive situated visualization systems and studies within the latter.

2.1 Mixed-Reality Visualizations

2.1.1 Information Visualization

Card, Mackinley, and Shneiderman describe *visualizations* as *the use of computer-supported interactive visual representation of data to amplify cognition* [9]. Cognition itself is defined as the acquisition or use of knowledge. People can transfer their loaded cognition to their perceptual system (in this thesis, I will focus on vision), especially for tasks that benefit from spatial organization, such as search or recognition. Munzner [231] highlights that visualization is helpful in many domains to augment human capabilities when they need to explore data (e.g., to understand the analysis requirements before developing an algorithm), analyze data (e.g., to know how the algorithm is affected by changes of parameters), make

decisions (e.g., to judge whether the algorithm is doing the right thing) or communicate (i.e., present outputs). The complex visualization process has been conceptualized and divided into several steps that I indicate at this thesis’s early age, forgetting the physical world (mainly because traditional displays do not rely on it).

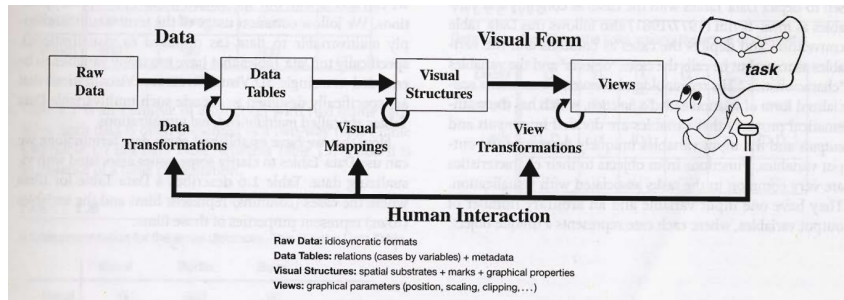


Figure 2.1 Visualization model presented in Mackinlay and Card’s work [9]

Figure 2.1 illustrates the visualization model in Mackinlay and Card’s work [9]. This model found several shapes and alternatives over the years, but the main structure remains unchanged. Roughly, **data** are mapped onto visual structures (i.e., a process known as **visual encoding**) that supports a **user** to perform a specific **task** through **interaction**. Each pipeline component has its declination of types: data could be items, links, or attributes that are part of tables or trees, attributes that can be categorical, ordered, etc.

To state if a visualization tool is *efficient*, researchers adopt strategies to evaluate it. Because it’s a complex system, each pipeline component may need a specific validation method. Plaisant [251] outlines four different thematic areas of evaluation for visualization tools including controlled experiments that compare design elements (e.g., speed of a transition between two visualizations), controlled experiments that compare different tools (e.g. a treemap and a spreadsheet, a computer screen and an AR headset), usability studies (i.e., focuses on the proficiency of users in acquiring and utilizing a tool to complete a task), and case studies (i.e., evaluation in real or realistic settings). One big difficulty is to guarantee that all these methods allow the applicability of findings in practical field settings.

Current research trends in visualization include immersive visualizations, presented in Section 2.1.2 and *situated and embedded visualizations*, introduced in Section 2.2.

2.1.2 Immersive Analytics

2.1.2.1 Definitions and concepts

In this thesis, I mainly explore Augmented Reality, when 3D virtual objects are integrated into a 3D real environment in real time [44]. Augmented Reality is a particular area on the continuum of Virtual Reality

presented by Milgram and Kishino in “A Taxonomy of Mixed Reality Visual Displays” [224], a keystone in the field. More broadly, all MR system exists in the continuum between completely real and completely virtual environments, bringing notions such as Reality (i.e., opposed to an artificial environment created by computers), immersion (i.e., users can feel immersed regardless of whether the environment is more real or virtual), and directness (i.e., whether primary world objects are viewed directly or using some electronic synthesis process).




Figure 2.2 Reality-Virtuality Continuum by Milgram and Kishino [224]

The field of study that combines visualization, immersive environments, and Human-Computer Interaction is known as **Immersive Analytics** [125]. Since its first goal identification half a decade ago by Chandler et al. as “*exploring the applicability and development of emerging user-interfaces technologies for creating more engaging and immersive experiences and seamless workflows for data analysis application*” [89] the research field has proposed many definitions for Immersive Analytics. A most recent one by Marriott et al. [218] clarifies its benefits and brings the concept of *embodied interaction* in the definition standing it is “*the use of engaging, embodied analysis tools to support data understanding and decision making*”.

2.1.2.2 Enabling Technologies

Here is a list and descriptions of some of the most commonly used AR displays:

- **Projection-based displays** (also known as Spatial Augmented Reality displays or SAR, Figure 2.3c): These systems detach the technology from the user and superimpose digital content onto physical objects or surfaces. By examining the surroundings and adapting the projection accordingly, these systems generate the impression of seamlessly changing the appearance of physical objects.
- **Handheld Displays** (Figure 2.3a): These displays are small computing devices that users can hold in hand and are sufficiently recent to be outfitted with integrated camera systems. They encompass sophisticated computer vision methods and algorithms (such as Simultaneous Localization and



Figure 2.3 Commonly-used AR-enabling devices. (a) Ikea app <https://apps.apple.com/fr/app/ikea/id1452164827>. (b) Hololens 2 ad <https://cdn-dynmedia-1.microsoft.com/is/image/microsoftcorp/gldn-HIGHLIGHT-hololens-2-4000x2100:VP2-859x540>.

Mapping methods or SLAM), enabling the real-time overlaying of 3D virtual objects. They include smartphones (Figure 2.3a) and tablets.

- **Head-Mounted Displays** (Figure 2.3b): These displays are worn on the head and typically cover the user’s eyes. We distinguish optical see-through devices that combine digital 3D objects with the direct view of the user (i.e., the real world is seen through semi-transparent mirrors, allowing the user to see holograms in the *real world*) and video see-through devices that, as handheld displays, present video feeds from cameras inside head-mounted devices with superimposed 3D virtual objects. Microsoft Hololens (Figure 2.3b) ¹, Magic Leap ², MetaQuest Pro ³ and the very recent Apple Vision Pro ⁴ are among the most known AR-HMDs.

To provide an Augmented-Reality experience, second and third types of devices are equipped with tracking systems (e.g., digital cameras and other optical sensors, GPS, accelerometers, and wireless sensors) that record and analyze the user’s posture and movements related to the physical world and some of the physical world attributes (e.g., planes identification, objects recognition, and tracking). They are usually furnished with robust CPUs and a critical amount of RAM to handle the processing of camera images [85]. These are elements that constitute a limitation in itself when building visualizations with immersive technologies. In the following section, I broaden the scope of challenges.

2.1.2.3 Challenges

In particular cases, such as when the data is three-dimensional, it appears beneficial to use immersive technologies for analytics [125]. However, the immersive analytics field faces many challenges. I list here the ones that are directly linked to this work.

¹<https://www.microsoft.com/en-us/hololens>

²<https://www.magicleap.com/magic-leap-2>

³<https://www.meta.com/ca/quest/quest-pro/>

⁴<https://www.apple.com/apple-vision-pro/>

- **Placing visualizations accurately in space** [125] which highlights the need to know the 3D position and rotation of the AR device and this relative to both other virtual and physical surrounding objects. This complexity augments with the scale of the visualization scene.
- **Extracting and representing semantic knowledge** [125] that underlines the need to recognize and track objects for scene understanding.
- **Designing guidelines for spatially situated visualizations** [125] that pinpoints the need for more controlled studies and the obstacles related to working with the dynamic physical world (discussed in subsection 2.2.4).
- **Understanding human perception and cognition in situated contexts** [125] exposing the need to understand how people’s cognition adapts to spatially distributed interaction in this new paradigm mixing real and virtual (I discuss existing work in subsection 2.2.3).

Alternatively, Bach et al. highlighted the limitations of the initial generation of HoloLens, including a restricted augmented field of view, device resolution, stability, and training requirements [47]. In a broader sense, interacting with immersive systems is inherently complex due to the need to potentially design for different human senses and employ various interactive technologies (e.g., tangibles, controllers, eye-gaze, speech) [125]. Additional challenges identified involve enabling collaboration within these systems (i.e., collaborative analytics) and addressing user scenarios (as discussed in subsection 2.2.4, particularly in the case of situated visualizations). Also, evaluating immersive systems is challenging due to the lack of clearly defined evaluation metrics and methods within immersive 3D contexts and the difficulty in understanding their effects on the user’s psychological state compared to traditional interfaces (e.g., fatigue, arousal). Experts argue that further research should include the establishment of a more precise evaluation framework [125].

2.2 [Immersive] Situated Data Representations

2.2.1 Definitions and Concepts

A situated representation is a data representation located close to the data’s *physical referent*, meaning close to the space, object, or person the data is about [322] (see Figure 2.4).

The relationships between data, referents, and data representations are multiple and complex. When designing situated representations for specific tasks, one approach is to initiate the process by selecting a set of relevant and informative data, with the subsequent goal of integrating this data into the physical environment (data-centric perspective). An alternative strategy entails first focusing on the




Figure 2.4 Visualization pipeline extended to the physical world [322].

referent of interest, which drives the task’s objectives, and subsequently deliberating on the pertinent data that would best represent around (referent-centric perspective). Moreover, this intricate process encompasses using either a simple data set or incorporating multiple diverse and heterogeneous data sets. Additionally, one may consider a unique referent or multiple referents (which makes it an *embedded representation* [322]⁵) dispersed in space. Further, we encounter multiple forms of so-called indirections when designing these connections between referents, data, and their representations. The most discussed are *spatial* indirection, which pertains to the spatial separation between the representation and the referent [322], and *temporal* indirection, which relates to the temporal misalignment between the time the representation is shown and the time the data refers to [322]. However, it is noteworthy that the significance of these indirections lies in their subjective perception by the users. Willett et al. thoroughly explored this critical aspect [322].

This type of visualization offers numerous potential benefits, including displaying information within the physical environment and enabling viewers to interpret data in context [322]. For example, a visualization of data about a house for sale is not situated, if shown on a desktop computer, from somewhere far from the house’s location. It becomes located is shown next to – or inside – the house. Visualization designers can also customize situated visualizations to highlight spatial relationships between data and the physical environment, facilitating decision-making and action. Moreover, they can integrate data into physical environments, ensuring that it remains visible over time and enabling ongoing monitoring of changes, identification of patterns, and collaboration with others [322].

A relevant research area is Situated Analytics (i.e., the science of analytical reasoning facilitated by visual interactive interfaces) focusing on information comprehension, decision-making, and support for pragmatic actions through situated visualizations [295]. These representations provide specific representations such as statistics. They are well suited for discovering, interpreting, and communicating meaningful

⁵along the thesis, I will use situated representation to refer to both situated and embedded representations. I use the term embedded when a distinction is needed.

patterns in data that are directly relevant to and integrated into the physical space around them [295]. While they provide rich contextual info, they have been triggered for ambient and casual use and specific data navigation and exploration [295].

2.2.2 Enabling technologies

Various techniques exist for displaying and situating data visualizations into the physical world. One of the most straightforward approaches is to place regular computer displays close to physical referents [322]. Alternatively, visualization designers can use physical data representations (also called data physicalizations [184]) instead of regular displays and place them next to referents (e.g., putting a robotic plant next to garbage bins to convey data about recycling behavior [170], Figure 2.5).



Figure 2.5 A robotic plant conveying data about recycling behavior, next to recycling bins [170].

Situated representations are often associated with AR and immersive technologies [295] (see Section 2.2.3.1) that enable the display visualizations in a way that makes them appear to be close to their physical referents or overlaid on top of them. In that way, information is seamlessly displayed, allowing immediate cohesive association of information with relevant physical objects, harmonizing contextual and overview information. As an example, AR-based situated visualizations have been explored in domains such as health (e.g., to show vital information about a patient during a surgery [37], Figure 2.6b), or aviation (e.g., to display flight information directly in front of the pilot [135], see Figure 2.6a).

In the next section, you will find precisions about the new perspectives offered by MR for situated visualizations.

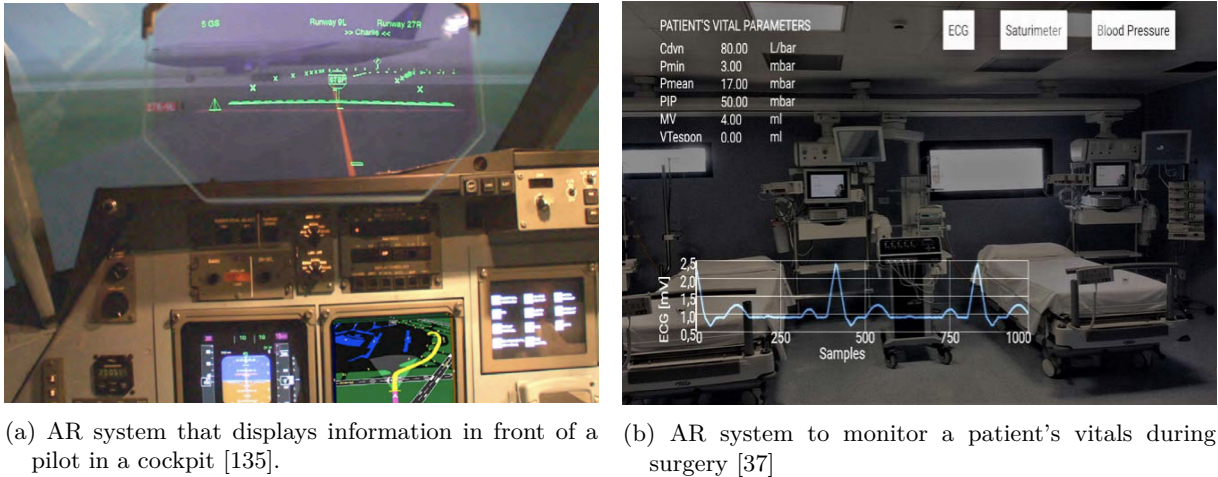


Figure 2.6 Two AR-based situated visualizations systems.

2.2.3 Existing Work in Immersive Situated Visualizations

2.2.3.1 Frameworks and Design Space Explorations

Willett et al. re-explored the visualization pipeline presented subsection 2.1.1 while considering the environment and covering both the logical and physical world. Figure 2.4 represents this extension of the visualization pipeline. The upside part shows the logical world visualization pipeline, which turns raw data into a visualization. This part is connected on both sides with the physical world; on one side, raw data relates to a physical referent, and on the other side, the visualization can exist as a physical presentation. I already discussed relations between representations and referents within the physical world (b) in Section 2.2.1.

Satriati et al. expanded the model by bringing nuances on the consideration of the physical world, introducing the concept of *environment*, which refers to the physical surroundings of the physical referent. They distinguish cases where situated visualization occurs within the environment from those in other locations, explaining that situated visualization can be achieved using a proxy of the physical referent, referred to as a prox-situated visualization.

Thomas et al. [295] and Bressa et al. [81] outlined several particularities on elements of the pipeline, when thought including the physical world and going further, the use of traditional screens. They underline **data** should be considered along with their temporal and spatial indirections (see Section 2.2.1), which can be supported by ubiquitous sensing. Also, prior research primarily examined the perception of **visual encodings** with a focus on two-dimensional representations suited for presentation on screens, whereas considering immersive displays, designers may need to incorporate the physical world into the visual encoding process (see Section 2.2.4). Taking the physical world into account brings a wide range

of potential new visual encodings. Moreover, how the user interacts with the visualization includes new various **interactions** modalities such as finger-worn sensors, digital pens, on-body interactivity, physical objects, etc. In particular, AR displays enable *reality-based* interactions [181]: users can interact with digital content using intuitive gestures, voice commands, or even direct manipulation of virtual objects within their physical environment. When this interaction mimics real-world affordances, we refer to *embodied interactions*. This natural interaction paradigm bridges the gap between the digital and physical worlds, enabling users to engage with visualizations more seamlessly and effectively. Additionally, the **task** shifts the focus from experts to a much broader population (i.e., considering InfoVis, see subsection 2.1.1). Thus, tasks pertain to a wide range of contexts, and the formulation of tasks should involve an iterative process that promotes a comprehensive exploration of space. Finally, unlike spatially and temporally detached visualizations, situated visualizations enable the immediate translation of analytical reasoning and decision-making into physical action, which pertains to a different but promising analytical process.

In conclusion, considering the physical world when implementing situated visualizations is complex. The configurations are diverse, and the research community strives to map out these issues better. While existing systems provide situated visualizations, it is also crucial to understand how they fit into this design area.

2.2.3.2 Situated Visualization Systems and Studies

First, I will introduce the most relevant and often used systems chosen when introducing situated visualizations using AR. The first one, Virtual Vouchers [316] (Figure 2.8), is a mobile augmented reality (AR) system designed to assist botanists in identifying and studying various plant species. It utilizes a vast dataset, including photographs of preserved leaves known as vouchers. These photographs are seamlessly presented within the user's field of view, allowing direct comparison with physical plant specimens. The second, SiteLens [317] (Figure 2.7b), is a tool that provides context-aware visualizations and information overlays (CO levels data) directly within the physical environment during urban site visits. The paper presents the results of user evaluation, where they led a usability field study demonstrating the effectiveness and positive impact of SiteLens in enhancing urban site visits. A most recent example is Corsican Twin [256] (Figure 2.7c), is a VR authoring tool that enables the remote design of AR visualizations in the context of building management using digital twins. It allows users to create situated and embedded experiences for building-related purposes. For instance, after designing the visualization in VR with a reproduction of an air handling unit, users can use them on-site and visualize digital information about fan speed or temperature next to the specific air handling unit. This application has been tested among experts who acknowledge its value for monitoring or documentation purposes. These three examples are just an entry point to the different experts and nonexperts fields immersive situated visualization can




Figure 2.7 Three AR systems providing situated visualizations.

support. Along this thesis, I will describe other examples covering various activities like surgery [200] or heat visualization [49].

However, a few works explicitly compare the efficiency of situated visualization to non-situated visualizations. Among them is the work of ElSayed et al. [122], who compared situated analytics in AR with a traditional reading manual approach for grocery shopping. Results show a significant improvement in the AR approach compared to the manual one: participants completed the tasks quicker and with fewer errors. Another example is HydrogenAR [318] (Figure 2.8), an illustrative instance where a data-driven storytelling experience is employed to elucidate the complexities surrounding hydrogen dispenser reliability. It is an exemplary case of collaborative design iterations involving a local hydrogen fuel research team and end users. HydrogenAR offers an innovative alternative to traditional PowerPoint presentations by effectively conveying data associated with dispenser liability close to the relevant equipment. The findings reveal a general inclination among participants towards HydrogenAR over the PowerPoint tour, with one participant’s feedback encapsulating the experience as the enjoyment of “identifying parts and observing data closely linked to them.”



Figure 2.8 HydrogenAR [318].

Finally Guarese et al.'s work [158] (Figure 2.9). propose the integration of data visualization with an augmented reality (AR) user interface to assist users in determining the optimal seating positions. They include attributes like temperature, WiFi signal, power outlets, and wheelchair accessibility path. They led a study comparing the usage of the AR interface (providing embedded visualizations) and a tablet device with a 2D map of the room (for situated visualizations). The study did not find statistically significant results in favor of the AR condition, but the tasks performed better and with fewer errors.



Figure 2.9 Augmented Situated Visualization for Spatial and Context-Aware Decision-Making by Guarese et al [158].

2.2.4 Challenges

Designing a visualization considering the physical referent can pose a significant challenge. One of the initial factors to account for is the geometry of the referent. It is crucial to ascertain whether a situated representation is necessary or if the spatial visual attributes of the referent's geometry can more effectively encode other data attributes (for example, data-driven jewelry [182])). When fitting a visualization into an environment, factors such as visibility, distance, and potential occlusion by referents need to be considered [81]. While most situated visualizations are crafted for scenarios where users can simultaneously view the visualization and the physical referent, as outlined by Ens et al. [125], one of the primary challenges faced is that the data is often presented in real-world environments that cannot be entirely controlled, characterized by cluttered spaces and continually moving objects and people. Consequently, in many instances, there is a high likelihood that the physical referent may not be visible when the user requires information about it. Previous studies have explored how to deal with non-visible objects of interest in mixed-reality environments (e.g., [64]), but these works have not delved into data visualization. Recent research endeavors explore methods for overcoming these limitations, including using proxies for physical referents, as introduced in subsection 2.2.3.1.

As previously discussed, it is crucial to envision new representations integrating the physical world into the design process beyond merely transposing our screen-suited representation methods. Likewise, it is essential to diversify and adapt evaluations for real-world scenarios (e.g., engage with communities to

conduct both laboratory and field studies). Additionally, we must continue to address technical challenges arising from implementation and address social and ethical questions associated with the use of AR. This entails considering the effects of situated analytics on human behavior, envisioning the potential for encountering unforeseen social acceptance, and examining data collection and privacy aspects, for instance [295].

This chapter unveils a multitude of exciting opportunities and research challenges that lie ahead. In this thesis, I will partially address some of these challenges to extend the presented state of the art. I will also reflect on how extensions of our work may approach some others of the above in chapter 6.

Our first contribution to representation design is the object of the next chapter: I present our exploration of Augmented-Reality *concrete representations* through the particular use-case of eco-feedback.

Chapter 3

Designing Immersive Concrete Visualizations for Eco-Feedback

Chapter summary

This chapter introduces a novel form of holographic representation designed to be attuned to the physical realm. Indeed, these representations are aligned with the physical world’s scale and displayed within the immediate user environment. We identified an applicative opportunity in the domain of eco-feedback. Indeed, modern society is organized so that waste is quickly taken away from our senses for comfort purposes. This prevents us from being directly confronted with the actual volume of waste we produce, except when disposal systems break down and create natural information displays.

We take inspiration from such natural displays and introduce a class of situated visualizations we call *augmented-reality waste accumulation visualizations* or ARWAVs, which are concrete representations of waste data embedded in users’ familiar environment.

We implemented examples of ARWAVs and demonstrated them in feedback sessions and during a large tech exhibition event. We discuss general design considerations for ARWAVs. Finally, we conducted a study with 20 participants suggesting that ARWAVs yield stronger emotional responses than non-immersive waste accumulation visualizations and plain numbers.

The main portions of this chapter are currently under peer-review for an academic journal and a conference. A portion of this work has been presented at the CHI’23 conference as an Interactivity demo [39] and an article in the Journal of Computing for Sustainable Societies [40]. Thus, any use of “we” in this chapter refers to Arnaud Prouzeau, Pierre Dragicevic, Martin Hachet, and me. You will find video material in subsection B.

3.1 Introduction and contributions

The adverse impact humans have on the environment (e.g., air pollution, plastic pollution, soil erosion, or damage to biodiversity) is one of the biggest challenges currently faced by our society. Although the causes are complex and numerous, individual behavior and lifestyles are among the key contributors [30, 219]. In particular, current consumption habits in industrialized countries¹ far exceed what humans can sustainably manage without damaging the environment. A factor that likely contributes to the excessive consumption and unnecessary waste production in those countries is the *invisibility* of waste there [53]: when we buy a new smartphone, we do not see the metal mines that went into its construction [329]; Even when we directly create waste, it quickly goes away and becomes invisible: dirty water moves to the water treatment plant right after we showered or flushed our toilet, and the garbage we throw away in our homes is conveniently taken away to a landfill site every few days.

Starting from this observation that modern society is organized so that our waste is hidden from our view, accidents sometimes happen that cause waste to accumulate and unexpectedly come back to our attention: bathtubs and toilets get clogged; Garbage gets scattered in the streets during large public gatherings (Figure 3.2a) or piles up during garbage collector strikes (Figure 3.2b). In such cases, waste piles serve as visualizations of our waste production that are physical [117], concrete [240], and situated [322]. Perhaps because such waste accumulations are so effective as displays of information, they are sometimes created and displayed on purpose. For example, some collective restaurants use transparent bread garbage cans to raise their customers' awareness about food waste (Figure 3.2c), and a designer has made a toilet with a transparent water tank to act as a preview of the amount of waste water that will be produced (Figure 3.2d). Waste accumulations like these act as situated information displays that are unique in their ability to convey waste production in a way that is immediately understandable by a large audience and can carry an emotional impact. However, they also have a very narrow range of applicability: creating physical waste accumulations can be difficult in practice, typically requires space, and can cause sanitary issues.

In this chapter, we propose to use Augmented Reality to re-create virtual versions of such information displays and embed them in users' physical environments (see Figure 3.1). We will refer to them as *augmented-reality waste accumulation visualizations*, or ARWAVs for short.

We use different use cases to illustrate how ARWAVs can be useful in various situations.

We also explain why ARWAVs go beyond what has been explored in research so far. In particular, they fill a gap in eco-feedback research [139], where most systems convey resource consumption and waste production using units and visual representations that are useful but often abstract and potentially

¹A developed country, or a sovereign state that has a high quality of life, developed economy and advanced technological infrastructure relative to other less industrialized nations (https://en.wikipedia.org/wiki/Developed_country)



Figure 3.1 A user looks at how much waste his corporate restaurant produces in a week. The waste is represented by virtual trash bags, which are displayed directly in the restaurant using Augmented Reality (AR).

difficult to grasp intuitively. ARWAVs are *concrete representations* of waste amounts (as they resemble their physical homologs, e.g., 300 liters of garbage can be represented with ten 30-litre trash bags) and are embedded in the user’s familiar surroundings; we expect them to give a more *visceral* sense of quantities and stand as more engaging representations. Also, using Augmented Reality to go beyond traditional representations, considering the physical world upstream of our reflections for designing situated visualizations, addresses one of the challenges discussed in the background section.

A user study with 20 participants that we will describe later in this chapter outlines that indeed ARWAVs tend to elicit higher emotional intensities than other more classic visualizations. As such, we see them as a useful potential complement to more standard information displays in eco-feedback systems, and we hope they will inspire future research in this area.

Our contributions are:

- (1) Identifying and naming a family of eco-feedback visualization techniques (ARWAVs) that have received very little attention so far and presenting scenarios illustrating the different situations in which they can be useful;
- (2) Describing prototypes that implement such visualization techniques and that were tested in an initial feedback session, both with colleagues and with a larger public during a tech exhibition event;
- (3) Introducing a terminology and a set design considerations for ARWAVs.
- (4) Reporting a user study suggesting that ARWAVs are emotionally more engaging compared to less



Figure 3.2 Examples of physical waste accumulations. (a) Street of London during the 2009 marathon (source: Paul Simpson, flickr.com) (b) Street of Marseille in 2010, during a garbage collector strike (credit: pxhere.com, public domain). (c) Transparent bread garbage can to raise waste awareness (source: <https://www.chef-eco.fr/>). (d) Toilet with transparent tank to raise awareness of water usage (photo taken in a Parisian restaurant).

immersive display modalities (3D on screen) and simpler information representations (numerals).

3.2 Background

In this section, we review previous research related to augmented-reality waste accumulation visualizations (ARWAVs). To this endeavor, we review related work in information visualization that falls into the scope of concrete visualizations as we define it. Additionally, we examine existing immersive systems that convey such data representations. Another highly relevant research area is the study of *eco-feedback* systems, i.e., “*technology that provides feedback on individual or group behaviors to reduce environmental impact*” [139]. Research on eco-feedback is vast and has been conducted both in psychology, focusing on cognition and behavior, and Human-Computer Interaction (HCI), focusing on technology [139]. We review them both in this section.

3.2.1 Concrete Visualizations

3.2.1.1 On-screen Concrete Visualizations

Visualization researchers have investigated and explored many different ways of representing data to promote understanding and engagement among large audiences; one of them is the use of collections of icons or objects to convey numbers, a common practice in visualization and infographic design.²

These are close to the idea of *Concrete visualization* that Nieman has defined [239] as “*the process of transforming entities that are usually represented abstractly as numbers in a table or lines on a chart into*

²See, e.g., [luizaugustomm.github.io/anthropographics](https://github.com/luizaugustomm/anthropographics) and dataphys.org/list/tags/single-datum/.




Figure 3.3 A concrete visualization of the CO₂ emitted by building a typical estate of 82 houses. Source: <http://www.carbonvisuals.com/projects/wood-for-good>.

objects in space (or a representation of real 3D space) available to our senses". Although this definition is vague, many examples from Nieman's work are pictorial unit charts rendered in 3D and embedded in familiar 3D environments. For example, in his short movie "Wood for Good - The Story of the Tree" (Figure 3.3), volumes of carbon emissions are conveyed using heaps of spheres rendered inside 3D environments.

Similarly, data videos [31, 104] have appeared on streaming platforms where familiar objects are embedded into realistic 3D scenes to help the audience understand numbers that are hard to grasp – such as large amounts of people, magnitudes of earthquakes [23], wind forces [20], or the gravity of different planets [24]. Nieman's concept of concrete visualization was refined by Chevalier et al. [94] who defined *concrete scales* as the "process of visually relating complex measures with familiar objects from the real world". Using familiar references, the concrete scale framework focuses on strategies to convey complex quantities and units. One such strategy is *unitization*, which breaks down abstract quantities into more relatable chunks. For example, to represent the amount of sugar in a soft drink, one may add a pile of sugar cubes right next to the drink [94]. It also discusses which object or container to choose to represent quantities. However, the paper does not discuss how concrete-scale visualizations can be used for eco-feedback, nor how they can be used in AR.

These have already been used in environmental communication in the web article "Drowning in Plastic" [271] (very similar to the example of Figure 1.4 presented in chapter 1 and to Figure 3.3), showing computer-graphics renderings of gigantic heaps of plastic bottles situated next to prominent landmarks such as the Eiffel tower. It conveys how many plastic bottles are purchased worldwide on different scales,

from a single day to the last ten years.

In all such examples, although the concrete visualizations occupy a 3D space, they are conveyed through pictures or videos shown on flat displays. This chapter is interested in concrete visualizations experienced in immersive environments, which I will review next.

3.2.1.2 Immersive Concrete Visualizations

I already discussed the recent surge of interest in using immersive displays, particularly for data visualization in subsection 2.1.2. While many immersive analytic prototypes proposed so far use abstract visualizations (e.g., 3D node-link diagrams, 3D scatterplots), a small number of them use concrete data representations in the sense of Nieman [239].

Inspired by the *Concrete Scales* framework, Lee et al. [201] introduced *data visceralization* as “*a data-driven experience which evokes visceral feelings within a user to facilitate an intuitive understanding of physical measurements and quantities*”. Many of their examples are concrete visualizations implemented in VR, such as a scene that conveys US debt through huge piles of dollar bills standing next to known landmarks or realistic reconstructions of protests that convey protest attendance numbers.

Similarly, Ivanov et al. [180] presented an immersive and interactive concrete visualization of mass-shooting data, where each victim is represented with a human silhouette in a VR environment. They speculate that immersing users in concrete visualizations can enhance their emotional connection with the data. Finally, Chen et al. [91] presented a tool for designing and creating 3D unit charts and concrete visualizations in AR.

I will now examine representation strategies used for eco-feedback, clarifying the current status of using concrete visualizations and immersive displays in this field.

3.2.2 Eco-Feedback Technologies

Simple eco-feedback displays have long been available. For example, in 2007, most UK residents with solar panels had a device that displays energy output and the total CO₂ mass saved since their installation [190]. Today, several eco-feedback apps are available on mobile devices to support self or group consumption awareness – for example, the MyImpact app³ helps users log their plastic reducing and recycling actions, and displays the total amount of plastic reduced and recycled, which they can compare with friends.

While researchers in environmental psychology have mostly focused on the efficacy of simple eco-feedback displays like these, HCI researchers have emphasized exploring richer and less conventional

³www.myimpact.com

information displays [139]. Examples include representing energy consumption as waves on a kitchen wall [55], or designing public trashcans that project their content and history of use on the sidewalk [246]. Researchers have also started to explore the use of virtual reality (VR) for eco-feedback, although mostly to test research hypotheses in experiments [48, 314]. We will discuss these experiments in subsection 3.2.3.

While there have been some explorations in VR, augmented reality (AR) has rarely been considered for eco-feedback. Several near-AR eco-feedback systems have been proposed or studied, including appliances (e.g., trashcans [246] and fridges [232]) augmented with sensors and displays, AR overlays of energy consumption on physical building models [127], and heat cameras to visualize thermal loss [74]. But as far as we know, enhancing real-world environments with AR for eco-feedback purposes has been only discussed mainly as opportunities (e.g., in [87, 136, 266]), and concrete implementations are very rare.

We only know of two examples of ARWAVs discussed in the past. The first one is an AR visualization of personal trash production designed during a research internship by Nic Stark [285]. In this prototype, which has directly inspired our work, users can use their smartphone to visualize a year’s worth of trash bags in front of them. The second piece of work is by Honee et al. [171]. In this work, conducted in parallel and independently from our own work, a specific type of ARWAV is proposed that focuses on food waste, and a usability study is reported. Besides these two examples, we do not know of any academic work that discusses the general concept of ARWAVs. This work fills this gap by conceptualizing ARWAVs, presenting a range of examples, and providing design considerations. As far as we know, our work is also the first to report a controlled user study of an ARWAV with non-AR comparison baselines.

3.2.3 Eco-Feedback Studies

Many studies have been carried out both in environmental psychology and in HCI to determine whether eco-feedback systems are effective and how to design them best. Directly relevant to this work is the question of how to display quantities, such as electricity consumed or amounts of water used. Pierce et al. proposed a taxonomy [250] that distinguishes three types of eco-visualizations: pragmatic eco-visualization that aims to effectively communicate resource processes using scientific visualization elements such as numbers, charts, and graphs, artistic visualization that employs persuasive imagery to create meaningful representations such as metaphorical metrics and informative art that occupies an intermediate position between pragmatic and artistic approaches, characterized by its decorative nature and initial lack of recognition as data visualization. ARWAVs can easily convey information when presented as data visualization and foster immediate understanding simultaneously. Therefore, they fall between two of the latter categories (i.e., pragmatic visualizations and informative art). Studies have found that units used by commercial eco-feedback monitors (e.g., kWh or tons of CO₂) are often not meaningful to users [287]. One study compared people’s preference for simple quantitative displays (i.e., abstract, like numbers or

charts) vs. qualitative displays employing metaphors (e.g., concrete as a polar bear on a melting iceberg), and found that STEM⁴ students tended to prefer the former, while non-technical students were more comfortable with the latter [51]. In contrast with metaphor-based or qualitative displays, the goal of ARWAVs is to offer faithful representations of quantities. However, such representations significantly differ from simple charts, focusing on efficiently communicating magnitude rather than precise numbers.

Eco-feedback displays can either use a negative framing (e.g., show energy consumed) or a positive framing (e.g., show energy saved). Studies overall suggest that negative emotions such as regret and guilt lead consumers to choose sustainable products and services [193] and to engage in sustainable behaviors like recycling and using public transportation [88]. This confirms that negative framing, which is used in many eco-feedback designs, is effective. At the same time, an overly negative framing can backfire. For example, Bao et al. [50] showed different eco-feedback designs to participants, including one where a polar bear was pictured decapitated when the light switch was on. While some participants liked the design, most of them found it inappropriate, and some found it manipulative. Similarly, a study found that people preferred to see their CO₂ consumption expressed in terms of trees necessary to offset it than the number of polar bears whose habitat it will destroy or the number of Earths needed to support humanity if everyone consumed the same (i.e., more negative framings) [249]. While the types of designs we discuss in this chapter (waste accumulations) naturally lend themselves to a negative framing, we also explore positive framing (e.g., waste avoided).

Two eco-feedback studies used a design similar to ARWAVs, but implemented in VR. In one study, participants were asked to take a virtual shower and were shown the amount of heating energy they used in terms of pieces of coal (each piece stood for 100 watts or 15 sec. of the shower) [48]. After the experiment, participants used less hot water to wash their hands when they saw virtual renditions of actual piles of coals than when they saw the number of pieces of coals expressed as text. Another VR study asked participants to repeatedly use a water bottle to fill a tank with enough water for a toilet flush or a 1-minute shower, after which they were shown multiple water bottles representing that quantity [314]. They found a positive effect on self-reported attitude towards water usage. Although both studies suggest that the kind of visual representations used in ARWAVs can be effective, they each only consider very specific instances and do not discuss how they can be employed in actual eco-feedback systems, for example, using augmented reality.

The natural waste accumulation displays we mentioned in the introduction are inherently situated, as are most of the eco-devices discussed in subsection 3.2.2. ARWAVs continue this stream of research by leveraging aspects of situated visualizations that have been recently presented as opportunities [81]; This includes *temporal relevance* – as ARWAVs represent an accumulation over time of recently produced waste, as well as *community* aspects, as ARWAVs can be used to reveal waste locally produced by a group.

⁴acronym of science, technology, engineering, and mathematics

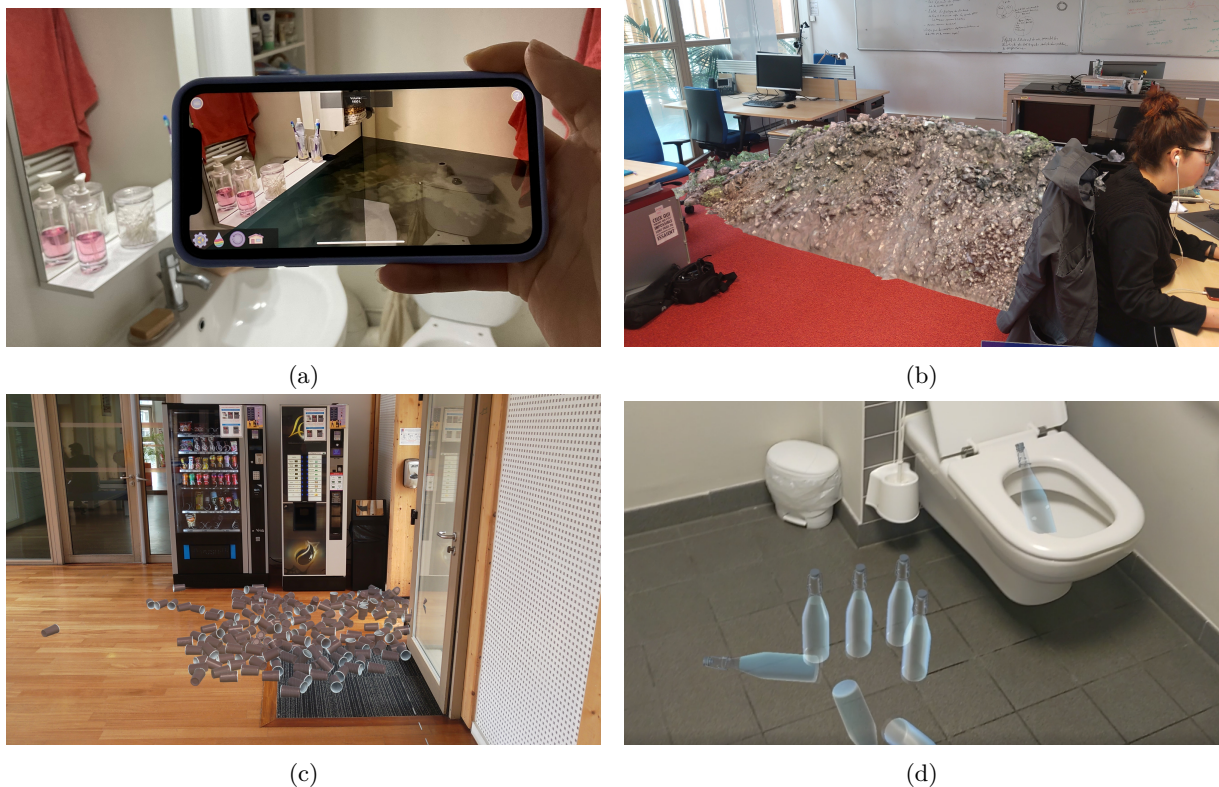


Figure 3.4 Illustrative examples of ARWAVs we prototyped. (a) The volume of consumed water visualized in a bathroom. (b) Approximate amount of material displaced and emitted to manufacture eight smartphones (captured from Hololens). (c) Quantity of single-use plastic cups accumulated for a given period of time (captured from Hololens). (d) 9 Litres of water represented using water bottle next to the toilets (average flush amount of water).

3.3 Illustrative scenarios

In this section, we present scenarios to illustrate and motivate ARWAVs. We will refer to these scenarios again in later sections, especially in our discussions of ARWAVs design considerations in section 3.5.

Augmented-reality waste accumulation visualizations or ARWAVs aim to help people better perceive the quantitative impact of their actions and decisions to inform them, raise their awareness, or encourage them to change their habits. The first one from Figure 3.1 illustrates a scenario where employees can perceive directly in their corporate restaurant a week of collective waste production represented with virtual trash bags. In the following sections, we present additional scenarios that illustrate the same concept. An illustrative video is available in the OSF project repository⁵.

Scenario 1: Personal motivation, water use Water is a resource that many of us abundantly use at home when we take a shower, wash our hands, flush our toilets, or use appliances like washing machines.

⁵https://osf.io/v4yxs/?view_only=cd8b973f9df54d7590f8421be6b11098

As water immediately disappears through the pipes, it can be hard to get a good impression of the amount we use over time. However, used water is waste that needs to be processed, and water can be scarce in lots of countries, where a drought can lead to a shortage with dramatic consequences. Knowing numbers is a good start, but it is likely not enough. For instance, it might not seem a lot to consume 5 m³ of water in a month, but it actually represents about 28 full bathtubs or an average bathroom filled with water up to 2 m high. ARWAVs can concretely show these volumes of water to make abstract numbers more concrete. A user could use their smartphone each time they get to a room of their apartment where water is used (e.g., bathroom, toilets, kitchen, garden). Figure 3.4a shows an example where they use it in their bathroom to see the amount of water they consume in a week, as a volume that fills the room. They can monitor this level occasionally when they get to the bathroom by looking at where the water stops on the wall, comparing their consumption across weeks, and trying to lower it little by little. Figure 3.2a shows another way of representing wasted water with water bottles.

Scenario 2: Group dynamics, ecological rucksack Computers and mobile devices are powerful tools with many useful applications, but they have a significant ecological footprint.⁶ One way to capture this footprint is the concept of *ecological rucksack*, which is the amount of matter displaced to build an object.⁷ Since the metal mines and the CO₂ emissions are typically situated far away from where the phone is purchased and used, it is hard for people to have a vivid picture of them in their mind, even when they hear about them. In addition, computing devices often look clean and beautiful, so looking at them does not bring about any association with mines, soil, or greenhouse gases. With ARWAVs, we can make those waste products more salient to users. For example, imagine a research team discussing whether to buy eight new smartphones for a project or use old ones. Here, an eco-conscious member of the team could book a room in their laboratory and set up a meeting in which they use a HoloLens 2 to show the rest of the team the size of the ecological rucksack necessary to build the new phones, as a pile of soil (Figure 3.4b). Having this heap of soil appear as if it was in the room (as opposed to, e.g., on a computer screen or a magazine) could give the team members a more visceral sense of quantity and size. It also adds a dramatic dimension, which can help make environmental issues more salient and influential in the discussion. Showing ARWAVs is a way for the meeting organizer to initiate dialogue and share their concern about their common purchase habits.

Scenario 3: Support for policies, plastic cups Single-use plastic cups have a negative repercussion on the planet. It is estimated that more than 50 billion are used in the US in one year.⁹ Yet manufacturing them is costly in energy and resources, they are hard to recycle, and they are one of ten most commonly

⁶The ecological impact of ARWAVs is an important issue that will be discussed in subsection 6.2.3.

⁷<https://www.gdrc.org/sustdev/concepts/27-rucksacks.html> A smartphone requires an ecological rucksack of about 70kg, which includes the amount of soil mined to obtain the metals that go into its fabrication, and the amount of CO₂ emitted during transportation and manufacturing.⁸

⁹<https://plastic.education/the-problem-with-disposable-cups/>

found waste items in European beaches¹⁰. However, plastic cups are still routinely given away in cafés, shops and companies, and thus it remains easier for many people to use them instead of bringing a reusable cup. In this third scenario, a company chooses to ban plastic cups as part of a program to reduce its ecological footprint but is afraid that not all its employees will agree. The company has already communicated about the issue through figures and charts, but they were largely ignored. Therefore, the company decided to organize an event with a booth in front of the coffee machine, where employees can observe through an AR-HMD (Head-Mounted Display) an ARWAV composed of all plastic cups typically used over a week (see Figure 3.4c). With this immersive experience, most employees get a much better sense of the amount of plastic cups accumulated over time, and many become more willing to make an effort to bring their own reusable cups. Afterward, some employees keep thinking of the virtual pile of cups every time they go to the coffee machine and are happy that the policy was adopted.

Process. To elaborate on these scenarios, we first designed and developed four ARWAV prototypes (further described in section 3.4). In designing these prototypes, we were inspired by the related concepts of concrete-scale visualization [94], and data visceralization [201], as well as by real waste accumulations and our own experience dealing with waste. We first discussed potential end-users of ARWAVs (e.g., individuals, small groups, communities) and their possible goals (e.g., self-motivation, decision-making, persuasion, support for policies). In imagining the prototypes, we sought to cover different types of visualizations (e.g., trash bags, cups, soil, water), different types of users, and different types of goals and tasks.

3.4 Prototypes and Initial Tests

To test ARWAVs and get initial user feedback, we implemented four ARWAV prototypes. We first review implementation details and then describe findings from our feedback sessions. We then list the main lessons learned from our feedback sessions and how they informed the subsequent stages of our research – i.e., the design considerations in 3.5 and the experiment in 3.6.

3.4.1 Implementation

The water use prototype (Scenario 1) was implemented with Unity 2020.3.2f1 on a smartphone and used the ARFoundation framework (handling most of ARKit and ARCore functionalities). We implemented mesh created from selected points on the real-world detected planes. This surface is extruded from the ground plane, matching the desired volume of water (predefined values include an average bath, shower, toilet flush, hand and dishwashing and washing machine). On rendering, we applied a 3D shader on

¹⁰https://ec.europa.eu/environment/pdf/waste/single-use_plastics_factsheet.pdf

each volume face with a material using the URP pipeline.¹¹ The final rendering takes the physical surroundings to render an adequate water color.

The trash bag prototype (shown in Figure 3.1) was previously implemented on a smartphone by Nic Stark¹² [285]. We re-implemented his prototype on a Microsoft Hololens 2 to be able to test and compare the two types of AR displays. The ARWAVs described in Scenario 2 and Scenario 3 (ecological rucksack, plastic cups) were also implemented on Hololens.

We developed with Unity 2020.3.18f1 and used the MRTK toolkit¹³. We enabled hands, head tracking, occlusion, and set the spatial awareness module to scan the environment at the start (update interval of 3.5 seconds). The objects making up the ARWAVs (trash, bottles, mud, and plastic cups) were either found on the Unity Asset Store or downloaded and adapted from 3D models browsers such as CGTrader, Turbosquid, or Free3D. For collision detection between the objects, to support numerous objects, we added an assembly of basic colliders that closely match the shape of the objects (instead of the more compute-intensive mesh colliders), and we set a discrete collision detection. We pre-created several prefabs (at run-time) to pile objects naturally, making them fall from a 1 to 2 meter height. Audio sources have been scripted to match collisions.

The source code for the ARWAV used in our controlled experiment (section 3.6) is available in our OSF repository.

3.4.2 Early User Feedback

In order to get initial feedback on our prototypes, gather impressions, and collect suggestions for improvement, we ran a testing session with five collaborators on a related research project. For all of them except one, who tested the prototypes three weeks prior, this was their first experience with AR-HMDs.

The testing session took place in a large meeting room during a workshop. We demonstrated the ARWAV prototypes (trash bags on HMD and smartphone, plastic cups on HMD, soil on HMD, and water with a smartphone).

Participants appeared engaged overall, qualifying the experience as *surprising* and *impressive*. When using the *Trash bags* HMD prototype, three participants stepped back when the trash bags started to fall due to the impression that they may fall on them. In the *Plastic cups* prototype, one participant commented on how real the cups looked when they happened to roll under a physical chair (and disappeared due to occlusion management). Participants commented on the differences between the HMD and the smartphone. Overall, the HMD provided a stronger immersion and elicited more surprise, although

¹¹<https://docs.unity3d.com/Packages/com.unity.render-pipelines.universal@11.0/manual/>

¹²We are grateful to Nic Stark for sending us his code.

¹³<https://docs.microsoft.com/fr-fr/windows/mixed-reality/develop/unreal/unreal-mrtk-introduction>

the surprise factor may be largely due to a novelty effect. Participants also felt they could manipulate items more naturally using mid-air gestures on the HMD than using 2D gestures on the smartphone. On the other hand, one participant mentioned that the smartphone is the only option available to them if they want to share a visual, which is already quite convincing. Among limitations, several participants mentioned the narrow field of view of the HMD (but also of the smartphone), especially for the *Trash bag* and the *Soil* prototypes. However, one participant emphasized that it was easier to overcome the narrow field of view by moving their arm with the smartphone than by moving their whole body and head with the HMD.

Finally, several comments had to do with emotions induced by the experience – *guilt* was mentioned several times, as well as other negative emotions such as *anxiety*, *stress*, and *oppression*. Positive feelings were also mentioned: one participant reported feeling *satisfaction* when seeing the *Plastic cups* prototype, as they do not use plastic cups anymore and could see the approximate amount of plastic they save.

We also had the opportunity to demonstrate ARWAVs in a large-scale tech exhibition (Vivatech ¹⁴, receiving more than 90K visitors). This allowed us to test the robustness of our prototypes and collect additional user feedback. Moreover, exhibitions are among possible real-world applications of ARWAVs, as deploying ARWAVs in such spaces can help raise the general public’s awareness of environmental issues. We demonstrated our Hololens prototypes and added a computer monitor for others to see. We met approximately 40 people in our demo booth, who were highly diverse in terms of age and professional background. Most of them have already tried VR systems before, and a few have tried AR-HMDs. Viewers were generally impressed and sometimes surprised by the quantities: “*All this for one person! It’s huge!*”. On several occasions, they mentioned that this type of tool can be useful to raise environmental awareness. A metal waste management professional was enthusiastic about the *ecological rucksack* demo: “*It’s really great to be able to show this! In my daily life, I am confronted with large piles of metal, and this has sharpened my environmental awareness. I have often told myself that it is a public necessity to show this kind of place; with what you propose, there is no need to go there!*”. Two science communicators specialized in environmental issues, and two teachers were also enthusiastic and argued that this tool could be useful in primary schools, at environmental awareness forums or fairs, or as a pedagogical tool for ecology courses in higher education. According to them, a novel user interface like this can be more playful and impactful. The teachers also commented that it would be interesting to embed the visualization experience in learning spaces (e.g., classroom or amphitheater), where physical or digital course materials already occupy walls, screens, and tables. Overall, no negative feelings were reported.

¹⁴<https://vivatechnology.com/>

3.4.3 Lessons Learned

Importance of affect. Both our testers and visitors of the tech exhibition mentioned emotions elicited by ARWAVs. This underlines the importance of the affective dimension and led us to reflect more on affect, such as whether ARWAVs can be designed to elicit positive or negative emotions, a concept to which we will return in our section on design considerations (section 3.5). This also prompted us to study affect experimentally, as we will report in section 3.6.

Importance of AR equipment. People involved in prototype testing were highly sensitive to the type of AR display device used (Hololens or smartphone) and commented on the strengths and limitations of those devices. In particular, those who never used a Hololens before were astonished by how immersive the experience was, but they were also disappointed by the device technical limitations. This underlines the importance of the type of AR device used, which led us to include this dimension in our design considerations of section 3.5. Although we do not explicitly manipulate this factor in our experiment (section 3.6), we stress that it is important to test it in future experiments.

Importance of animations and interactions. We implemented basic support for animations and interactions in our ARWAV prototypes but did not expect they would be so salient to users during our testing session. This led us to include animation and interaction as important design considerations (section 3.5) and to implement animation (falling trash bags) in our experiment. We relied on this set of suggestions to build our study prototype. However, we did not include interaction in our experiment in order not to distract participants and keep the experiment simple.

3.5 Designing ARWAVs

In this section, we present design recommendations that have emerged from the following process. Initially, we engaged in brainstorming sessions among the authors, drawing from existing literature, to generate multiple designs for testing in our early user feedback sessions. Our implementation of ARWAVs in alignment with specific scenarios provided valuable insights. Subsequently, we consolidated the lessons learned from our initial testers and revised our list of suggestions accordingly. We acknowledge that these recommendations would benefit from an evaluation among experts (e.g., eco-feedback designers) to testify of its generative power. As we know that design considerations for ARWAVs go way further than we suggested below, we choose to convey our observations here and leave a deeper analysis of potentialities that would be assessable for future work.

3.5.1 Type of representation

The visual representation of ARWAVs is concrete but it can vary from highly realistic depictions of actual waste products to slightly more symbolic representations. We categorize them into four categories.

1. Literal: The most straightforward way to represent waste material in ARWAVs is to use virtual representations that mimic the material itself. For example, waste water can be represented by water (Scenario 1, 3.4a), and used plastic cups can be represented by plastic cups (Scenario 3, Figure 3.4c). The main difference between the actual waste and its representation is in the layout: presumably, the plastic cups have never been arranged like in 3.4c, and the water has never taken the shape of the bathroom like in 3.4a. Nevertheless, the waste objects or material are represented concretely and remain easily recognizable.

2. Transformed: This category corresponds to waste material that is represented as if the material underwent some processing or transformation, although the real waste material did not. For example, flushed water can be shown as if it has been put in bottles (Figure 3.4d), or soil like in Figure 3.4b can be shown packed in bags. Such quantity-preserving transformations, also called “unitizations” [94], can help people make sense of large quantities.

3. Materialized: Some waste material is difficult or impossible to perceive. For example, CO₂ is a common waste material but an invisible gas. For the purpose of ARWAVs, it could be useful to materialize CO₂ as, for example, a heap of coal or black smoke (as in [79]). Although this goes toward metaphorical representations, it remains close to literal because coal is mostly made of carbon, and smoke can contain a lot of CO₂. However, it is unclear if absolute volumes or masses are meaningful in these cases, and perhaps such ARWAVs would be mainly useful to assist comparisons (e.g., CO₂ emitted this year vs. last year).

4. Metaphorical: waste can be communicated through symbols and metaphors in many ways. As an example, Ahn et al. [27] describe a VR application where trees are used to convey data about paper usage. However, using metaphors goes against the principle of ARWAVs, which aims to convey actual waste quantities with representations that resemble the waste itself. We include this category for completeness nonetheless, as the metaphorical distance is a continuum [331], and as we saw, *materialization* goes a bit towards metaphors.

3.5.2 Realism of the representation

Realism is on a continuum and concerns all four representations we mentioned. The representations (water, bottles, smoke, etc.) can be modeled and rendered more or less realistically in all types. Lee

et al. [201] discuss realism in the context of data visceralization and concluded that “*realism is not important to understand the underlying data, but is still important for engagement*”. We identify several potential advantages of realism when it comes to designing ARWAVs. For example, the plastic cups in Scenario 3 (Figure 3.4c) could be designed to be exact replica of the actual cups dispensed by the coffee machine. In that case, the cups would be immediately recognizable by people who regularly use them, which can contribute to making the ARWAV more plausible or more provoking. But in many cases, closely mimicking the physical waste material may be difficult and not necessarily useful. For example, in Scenario 2 (Figure 3.4b), the soil probably does not look like typical soil extracted from mines. Similarly, the trash bags in Figure 3.1 may not be of the same color and appearance as the trash bags actually thrown away by the person whose data is visualized. However, the objects are close enough to be recognized for what they are and to give a reliable intuition of the waste quantity.

3.5.3 Animation and Interaction

Visual representations of waste in ARWAVs can be dynamic and even interactive. We distinguish between three levels:

- 1. Static:** These are waste representations that maintain their shape and appearance over time. This is the case, for example, for our soil prototype (Scenario 2, 3.4b). Users can walk around it but cannot interact with it. The main advantage is ease of implementation, especially for material whose physical behavior is challenging to replicate.

- 2. Animated:** Adding physically plausible animations to waste representations can make them more realistic and increase the user’s sense of immersion. For example, trash bags can fall individually to form a pile, or plastic cups can roll on the floor. This approach is easier for waste representations that consist of small to medium-sized collections of objects. Animations can be enhanced with sounds of objects colliding or crashing on the floor.

- 3. Interactive:** Letting users interact with the waste representations may help enhance immersion and realism. For example, users can pick and move objects in our trash bags and plastic cup prototypes. However, this is currently done through the standard HoloLens pinch gesture, and objects can only be moved one by one, which limits realism. Allowing physically more realistic interactions (e.g., shooting in plastic cups) would likely increase realism and immersion [181]. Besides aimlessly moving objects around, interaction opens up a range of possibilities, including in terms of gamification: VR systems already exist where users are invited to fill water tanks manually [314] or cut trees with chainsaws [27] to raise their awareness about water and paper usage. Similarly, it is possible to imagine ARWAVs where users are asked to spend energy to create, rearrange, or destroy waste accumulations. Interaction also opens up endless opportunities for data exploration, some of which will be discussed in subsection 6.2.1.

3.5.4 Framing

As we already saw in subsection 3.2.3, eco-feedback displays can be framed either positively or negatively, with possibly different impacts on users' emotions, attitudes, and behavior. Which framing is preferable is highly context-dependent and difficult to predict without user studies. Overall, framing is important to consider when designing ARWAVs.

1. Negative: Guilt or fear-tripping visualizations showing the negative impact of users' behavior can effectively encourage changes in attitudes and habits. As ARWAVs show waste accumulations, they have a naturally negative framing, which AR can be used to enhance: because people instinctively want to get rid of waste, showing it in their personal spaces (e.g., their living room) can elicit a sense of discomfort, whether their actual waste has been already disposed of (e.g., trash bags), or has been produced in a remote place (e.g., a heap of soil from a mine).

2. Positive: It is possible to use ARWAVs for positive framing by having them represent waste saved rather than produced. For instance, our plastic cup prototype (Scenario 3, Figure 3.4c), could be used to show all plastic cups the user saved since they switched to a reusable cup. Such designs focusing on positive effects are more likely to elicit positive emotions such as pride and satisfaction.

Note that framing refers to the designer's intent to elicit specific interpretations or emotional responses, but the actual responses might not match the intent and might vary across people. For example, imagine an ARWAV showing plastic cups the user would hypothetically save if they changed their behavior. This visualization could be seen as either negative (especially if the behavior change is considered unlikely or too difficult) or positive (if the behavior change is easy to contemplate). Likewise, an ARWAV showing collective waste can be interpreted differently depending on the user's perceived contribution (e.g., a user might feel proud of having produced less or ashamed of having produced more), and how important fairness and collective action are to them (e.g., a user might feel annoyed that others are not trying as hard as them).

3.5.5 AR Display Technologies

Different display technologies can be used to implement ARWAVs. We review them here and the trade-offs involved.

1. Hand-held devices: Smartphones and tablets are widely used, and recent models provide sufficient capabilities to support AR applications. Hand-held devices are, therefore, a strong vector to provide ARWAVs to a vast majority of people nowadays.

For instance, Scenario 1 (3.4a) illustrates the use of a smartphone in a private bathroom and can

be extended to any place where water is consumed. Another important advantage of hand-held devices is that their content can be easily shared within small groups of colocated people to support collective experiences.

2. Head-mounted displays: Compared to smartphones, augmented-reality head-mounted displays (AR-HMDs) provide more immersive experiences [59]. Hence, they may enhance the visceral impression of quantities conveyed by ARWAVs. However, they are more expensive, only support a single observer, and current versions are still heavy with a small field of view. AR-HMDs are more appropriate for educative demonstrations, as the one described in Scenario 3. In the future, AR-HMDs may become much lighter and have a much larger field of view, extending the range of situations where they can be used.

3. Spatial AR [71]: ARWAVs can also be implemented with video-projectors, which can allow them to be smoothly integrated into physical environments as in, e.g., JETSAM [246], a trashcan that records trash people put in it and projects it on the floor. Like hand-held displays, spatial AR can support shared experiences but with even larger groups of people. However, virtual objects can only be displayed on 2D surfaces, greatly limiting possibilities. Nevertheless, research on 3D spatial AR is underway [206] and may open up future unprecedented possibilities for ARWAVs.

4. Tangible: The use of tangible and augmented objects can also be an interesting approach for creating ARWAVs experiences. For example, in the Erlen project [99], a physical object equipped with LEDs shows accumulated power consumption over the day. Such interfaces are, however, dedicated to very specific cases and currently lack versatility.

3.6 Experiment: are ARWAVs Emotionally More Engaging?

3.6.1 Experiment Design

There are many possible ways ARWAVs can be evaluated. One important question is the choice of evaluation metrics. Broadly speaking, potential studies can look at the perceptual aspects of ARWAVs (e.g., whether people get a better sense of the quantities compared to alternative representations), the affective aspects of ARWAVs (e.g., whether people feel more emotions or feel more engaged), and the decision-making and behavioral aspects of ARWAVs (e.g., whether people adopt a more eco-friendly lifestyle). The behavioral aspects are arguably the most important but also by far the most difficult to evaluate. Motivated by our previous feedback sessions highlighting the importance of affect (see Section 3.4), we focused on the affective aspects as a preliminary study.

Another important question to consider when evaluating ARWAVs is the choice of comparison baseline. Many baselines are possible, including the diverse representations of waste production used in eco-feedback

research, ranging from simple digits to metaphorical depictions (see 3.2). For this study, we chose two baselines of comparison: *(i)* a simple number conveyed in text form, and *(ii)* an interactive 3D representation similar to ARWAVS, but presented on a computer screen. The 3D condition is an intermediary condition that has some elements of ARWAVS but not all of them: like most concrete-scale visualizations [94], it uses a concrete and realistic representation of data, but it is not displayed in augmented-reality. We included those two baselines to make it easier for us to disentangle the effects of the presentation format (accumulation visualization vs. numerals) from the effects of the presentation medium (AR vs. screen). We now go through the experiment design in more detail.

3.6.1.1 Visualized Data

We chose to show waste production data of a company building to their employees. We wanted the visualized data to be meaningful and to feel personal, in the sense that participants must have contributed – at least in small part – to the waste represented. This idea gathers themes discussed in the scenarios (see section 3.3). Each company member contributes to corporate waste production without necessarily taking responsibility for it, while the waste is partly due to lunch, office equipment usage, or other employees practices direct output. Working on prevention in the office is a way for companies to regulate their waste production, which is more and more strictly regulated by the government [26]. In effect, companies can communicate data about waste production to help employees identify with them, start discussions around their reduction, and promote good practices such as using fewer disposable cups or setting printers to two-sided mode by default. However, it is still unclear if the way of showing this type of data can have more or less impact on the audience. To find out, we ran a study in our lab building (about 400 employees), using building employees as participants, and showed them the average amount of waste produced by all employees in a week (5 working days). We contacted the building management, which monitors waste output and estimates that an average of 2275 liters of waste is produced weekly. This amount includes household waste (mostly from the cafeteria catering activity), cardboard, recyclables, paper, and glass. 2275 liters make up approximately 46 trash bags of 50 liters each. This is the datum we convey in our study.

3.6.1.2 Physical setup

The experiment took place in the cafeteria of the building (visible in Figure 3.5b), outside meal hours. A desktop computer (Dell Precision 3640 with Intel Core i9 processor) was installed on a table to administer the two baseline conditions, which will be explained in more detail in the next subsection. The ARWAV condition was administered using a Hololens 2 AR headset.



(a) The **3D** condition, where a pile of trash bags is rendered in a generic 3D cafeteria and displayed on a computer screen.



(b) The **AR** condition, where the trash bags are rendered in the cafeteria where the experiment takes place and where most of the waste was produced.

Figure 3.5 Two of the three experiment conditions.

Condition	Description	Representation	Realism	Animation & Interaction	Framing	Display Technology
Text	Number shown on a screen	Abstract	n/a	Static	Negative	Screen
3D	Accumulation visualization shown on a screen	Literal	High	Animated	Negative	Screen
AR	Accumulation visualization shown in AR	Literal	High	Animated	Negative	AR-HMD

Table 3.1 The three experimental conditions according to the design considerations of section 3.5.

3.6.1.3 Presentation formats

We convey the waste output information in three different ways:

Number shown on a screen. The number of trash bags produced is conveyed as numerals in a text sentence displayed on the desktop computer. The sentence is (translated from French): “*Each week, employees working in the Inria building produce 46 50-liter trash bags worth of waste (2021 data)*”. This condition will be referred to as **Text**.

Accumulation visualization shown on a screen. The same number is represented as a pile of 3D trash bags on the desktop computer (see figure 3.5a). The environment is a 3D model of a generic cafeteria, which is comparable to but different from the actual cafeteria. Before the 3D representation is shown, the screen displays the following prompt: “*Each week, employees working in the Inria’s building produce the amount of waste you will see (2021 data)*”. A 30-second animation first shows the 46 trash bags falling one by one from 3 meters above the camera’s viewpoint. Participants are invited to navigate in the 3D scene with the keyboard and computer mouse. This condition will be referred to as **3D**.

Accumulation visualization shown in AR. The same pile of trash bags as in the previous condition is shown, but this time in augmented reality (see figure 3.5b). The same text prompt is displayed as before. Then, participants saw the 46 trash bags fall individually in the cafeteria. They are told that they are free to move in all the cafeteria areas with the head-mounted display on their head. This condition will be referred to as **AR**.

The three experimental conditions are summarized in Table 3.1 and characterized according to the design considerations of section 3.5. Since these considerations cover the design of ARWAVs and only one of our conditions is an ARWAV, some design options (shown in gray in Table 3.1) have not been discussed in section 3.5. Note that all three conditions are *situated* according to Willet et al.’s definition [322] because the representations are shown within the environment where the waste was produced. However, the AR condition is *more situated* than the other two because the representations (trash bags) appear to be integrated within the physical environment instead of being shown on a separate computer screen.

We opt for sufficiently distant conditions, but we could explore other interesting scenarios. For instance, reproducing the cafeteria in 3D on a screen or creating a scene in virtual reality, either a reproduction of the actual cafeteria or a generic one, inside or outside the real cafeteria. We note that in the case of a 3D reproduction of the cafeteria on a screen, where the user is already physically present, it could lead to confusing comparisons by the user between physical reality and its virtual representation and, therefore, demands a highly accurate 3D reproduction. In this scenario, a persisting confusing factor would be the perception of scale cues in 3D.

3.6.1.4 Design

We used a within-subject design where each participant sees all three conditions in order to reduce the impact of inter-participant variability and thus increase statistical power. The presentation order was fully randomized across participants, among four possible orderings: Text-AR-3D, Text-3D-AR, AR-3D-Text, and 3D-AR-Text. In other words, Text appeared either first or last. We wanted to capture situations where people knew the exact figure before they saw the visualizations (3D and AR), as well as situations where they learned the figure only after having seen the visualizations. We did not include the orders AR-Text-3D and 3D-Text-AR to avoid situations where participants have asymmetric information between the 3D and the AR condition.

3.6.1.5 Procedure and Measurements

After signing the consent form and indicating their age and gender, each participant was invited to read a text introducing them to the data¹⁵ they were about to see (waste output in their building over a week). The text provided context but without giving the actual number. The participant was then asked how often they eat at the cafeteria. This allowed us to assess the extent to which they were familiar with the place and were responsible for the waste output. Answering the question also likely prompted them to identify a bit more with the data and the place. The participant was then asked to provide a rough estimate of the number of 50-liter trash bags the building produces in a week. This question allowed us

¹⁵Find the full experiment questionnaire in subsection B

to assess the extent to which participants' prior beliefs were aligned with the information later shown to them and also likely increased their curiosity and engagement with the information.

The participant was then shown the three conditions in one of the four orderings mentioned before. After seeing each condition, the participant was asked to fill the PANAS questionnaire [313], a standard self-reported affect questionnaire consisting of 20 questions (10 about positive emotions, 10 about negative emotions). For example, the questionnaire asks participants to report how much they feel interested, distressed, upset, etc. on a scale from 1 to 5. The participant was given the option to justify their responses using free-form text. The participant was allowed to review their responses to the previous conditions if they wished to, which we expected would further reduce random variability.

At the end of the experiment, the participant was asked to rank the three techniques depending on *(i)* each technique's ability to elicit emotions, and *(ii)* how much each technique helped them to understand the data presented. Finally, they were allowed to provide general comments about the study. The ethics committee of Inria approved this study.

3.6.1.6 Hypotheses

We expected that participants would report feeling stronger emotions overall in **AR** than in the other conditions and would report feeling stronger emotions in **3D** than in **Text**. The overall strength of reported emotions was operationalized by averaging the responses to all 20 PANAS questions, yielding an aggregate measurement on a scale between 1 and 5.

3.6.1.7 Participants

We recruited 20 participants (15 male, 5 female, mean age 27, SD = 10) using the lab's mailing list and word of mouth. As a selection criterion, participants had to be working in the building for more than three weeks. This allowed us to make sure that participants were familiar enough with the building and were responsible for its waste output, at least to some extent.

3.6.2 Results

3.6.2.1 Planned Analysis

All analyses reported in this section were prepared and registered [100] with the Open Science Framework (OSF) before the data was collected (see link¹⁶). Experimental data and the final analyses are also

¹⁶https://osf.io/v78ay/?view_only=193d16fbde7641c896018024130c9063




Figure 3.6 *Left*: Participants’ reported attendance to the building cafeteria; *Right*: Participants’ estimates of the number of 50-liter trash bags produced by the building every week; The blue line indicates the actual number.




Figure 3.7 Overall emotion intensity for each of the three techniques.

available on OSF (see 5). We report and interpret all our inferential results using interval estimation [108, 114].

Complementary demographics. In addition to age and gender, we asked participants to report how often they eat at the building cafeteria, with responses shown in Figure 3.6a. The other question asked them to estimate the amount of produced waste, with responses reported in Figure 3.6b. Overall, the data suggests that most of our participants are regular cafeteria users (and therefore contribute to the building’s household waste), and many underestimate the amount of waste produced by the building employees (with the median response being half the true value).

Overall emotion intensity. Again, overall emotion intensity is the mean response to the 20 questions of the PANAS questionnaire. Figure 3.7 shows the value of this aggregate metric averaged across all participants for each of the three conditions, with 95% confidence intervals (CIs). Unless specified otherwise, all CIs reported here are BCa bootstrap confidence intervals [194].

Overall, participants report higher emotion intensities after seeing the AR condition than after seeing the 3D condition, followed by the Text condition. These trends are confirmed by estimating within-subject differences in overall emotion intensity, shown in Figure 3.8. All intervals are located far from




Figure 3.8 Within-subject differences in overall emotion intensity.




Figure 3.9 Overall intensity of positive and negative emotions.

zero, providing strong evidence and support for our hypothesis that people tend to report stronger emotions with AR than 3D, and with 3D than with Text.

Positive vs. negative emotions. We broke down overall emotion intensity into positive emotions (10 questions in the PANAS questionnaire) and negative emotions (10 questions). The results are shown in Figure 3.9. The trends are the same whether we consider only positive or negative emotions. However, negative emotions are overall less intense than positive emotions.

Individual emotions. Figure 3.10 shows mean responses for each of the 20 questions of the PANAS questionnaire. For context, possible responses were labelled: 1 – “very slightly or not at all”, 2 – “a little”, 3 – “moderately”, 4 – “quite a bit”, and 5 – “extremely”. If we only look at the point estimates, the trend we observed before $AR > 3D > Text$ seems to hold for each emotion separately, except for emotions that most participants did not feel, namely *irritable*, *hostile* and *proud*. Among the positive emotions, those that participants reported feeling the most intensely are (in decreasing value of point estimate for the AR condition):

- *Interest*. On average, participants reported feeling “quite a bit” of interest. Responses are high across all three conditions, suggesting that interest was mostly intrinsic to the information presented. This confirms that our participants identified with and engaged with the information presented to them.
- *Attentive*. Participants also reported being attentive (on average, between “moderately” and “quite a bit”). Attention is similar to interest but implies perhaps an even deeper cognitive engagement. There seem to be differences between conditions, which would mean that the degree of attention

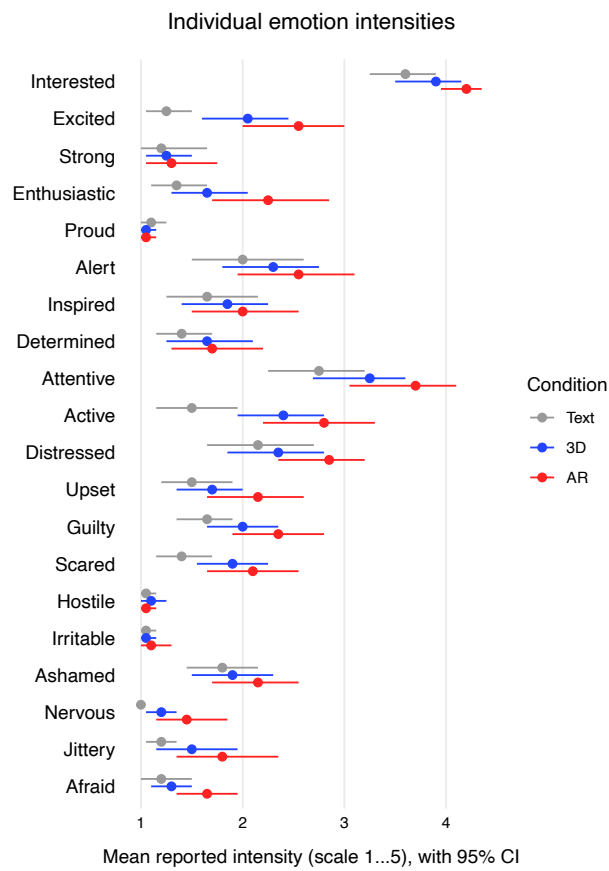


Figure 3.10 Reported intensity for all 20 emotions of the PANAS questionnaire.

may depend on the presentation technique.


- *Active.* Participants reported feeling reasonably active for 3D and AR (on average, between “a little” and “moderately”), but substantially less so for Text. This is likely because the Text condition is non-interactive, while in the 3D and AR conditions, participants were invited to move virtually or physically.
- *Excited.* Participants reported being reasonably excited but less so with Text, probably due in part to the interactive and entertaining nature of the 3D and AR conditions. In addition, there was likely a novelty effect at play with the AR condition.
- *Alert.* Similarly, participants reported feeling somehow alert but without large differences between conditions.
- *Enthusiastic.* Overall, participants felt slightly enthusiastic, but more so in the AR condition.

Among the negative emotions, those that participants reported feeling the most are (in decreasing value of point estimate for the AR condition):

- *Distressed.* No strong negative emotion was reported, but distress was clearly on top, with an average rating between “a little” and “moderately”. This is likely due to the negative nature of the information conveyed and the fact that most participants did not think the waste output was so high (see again Figure 3.6b).
- *Guilty.* Many participants reported feeling some degree of guilt (typically “a little”), with possible differences between techniques. This emotion is likely due to participants’ awareness of their contribution to the waste represented.
- *Upset.* This emotion yielded similar responses to the previous one, with likely higher responses for AR.
- *Ashamed.* Responses were comparable to the similar emotion *Guilty*.
- *Scared.* Finally, participants reported being a bit scared but substantially less so in the Text condition.

Technique ranking. Again, at the end of the experiment, participants were invited to rank techniques according to two criteria, the first one being their ability to elicit emotions. Responses are shown in Figure 3.11a: the three dots and CIs on top indicate the percentage of participants who ranked each technique first (the 95% CIs are computed using Wilson’s score method for independent proportions). The AR technique was ranked first by most participants (around 75%). The three CIs on the next row indicate that 3D was ranked second by most participants, while the last row indicates that Text was

CHAPTER 3. DESIGNING IMMERSIVE CONCRETE VISUALIZATIONS FOR ECO-FEEDBACK



(a) Ranking of techniques according to their ability to elicit emotions. (b) Ranking of techniques according to their ability to make information clear.

Figure 3.11 Technique rankings.




Figure 3.12 Mean within-subject differences between techniques for the three positive emotions (top) and the three negative emotions (bottom) for which evidence of an effect of technique is the strongest. The thin error bars are corrected for multiple comparisons.

ranked last by most. The second-ranking was according to the techniques' ability to make the information understandable, with results in Figure 3.11b. AR was ranked first more often than 3D, but Text was also ranked first by many participants, almost as many as AR. Overall, there is much more variability in the rankings for this criterion.

3.6.2.2 Additional Analyses

Analyses in this subsection were not preregistered. As part of our post-hoc analyses, we estimated pairwise differences between the three techniques for all of the 20 emotions reported in Figure 3.10. Doing so allowed us better to examine the strength of evidence of differences between techniques. For space reasons, we only report results for the three positive and negative emotions for which we found the strongest evidence of differences between techniques. We operationalized the strength of evidence as the p -value obtained by a one-way ANOVA omnibus test. The full results (20 plots) are available in the supplementary material on OSF.

Figure 3.12 shows the pairwise differences for our selected six emotions. Error bars are two-tier confidence intervals where the thick interval is a regular 95% CI, and the thin interval is a Bonferroni-corrected 95% CI for $3 \times 20 = 60$ comparisons.

Concerning positive emotions (top row on Figure 3.12), we can be reasonably confident that people report feeling more *excited* and *active* with 3D and AR than with Text. On top of that, people may report feeling more *excited* and *enthusiastic* with AR than with 3D, but the evidence is weaker.

The bottom row of Figure 3.12 shows the pairwise comparisons for the three negative emotions selected. There is some evidence that people report feeling more *nervous*, *scared*, and *guilty* with 3D and AR than with Text. It is possible that those emotions are slightly stronger in AR than in 3D, but the evidence in our data is weak.

3.6.3 Discussion and Qualitative Feedback

In this section, we put our results in perspective using open comments from participants. Participants are numbered P1 to P20 and all quotes are translated from french.

As suggested by our quantitative analysis (Figures 3.7–3.12), participants did not experience particularly strong emotions when shown waste output data but reported slightly more intense emotions with ARWAVs than with the two alternative formats. ARWAVs seem to have boosted positive emotions, such as engagement, and negative emotions, like guilt. This suggests that ARWAVs may be effective if the goal is to leverage people’s affect and that both the visual representation (literal visualizations of waste accumulation) and the medium (AR) may play a role.

The possible double role of visual representation and medium is consistent with participant comments. In terms of representation, P9 reported that plain text was insufficient for them to “*realize what [the number] meant*”, and P4 mentioned that it was “*hard to imagine this quantity*”. Participants commented on the benefits of having a 1:1 scale representation, with an environment (physical or virtual) that acts as a visual anchor: P5 and P9 mentioned that surrounding objects helped them get a sense of the size of the individual trash bags, while P1 and P4 reported they helped them get a sense of the total volume of waste. These comments are consistent with previous speculations about the benefits of concrete, object-based visualizations for conveying unfamiliar quantities [94, 201]. However, participants reported that the numeric format was useful too: P15 mentioned that it is “*more accurate on the precise amount of waste*”, P12 reported that it “*provides the clearest data, which I may be able to memorize*”, while P13 commented that it is the best tool to support understanding. Overall, participants had different opinions about which representation best-supported understanding, which can also be seen in the technique ranking data Figure 3.11b.

In terms of presentation medium, participants commented on the increased immersion and realism provided by AR. P4 and P18 referred to the trash bags in AR as if there were *“in real life”*, P2 mentioned that their amount was *“enormous”*, and P13 mentioned that the animation was impressive and gave a good feeling of the volume. Meanwhile, P4 reported that *“being immersed in the simulation enhances emotions”*. Participants also commented on the possible benefits of having a situated visualization, i.e., embedded in an environment relevant to the data presented [322]. Although all three conditions were administered in the cafeteria (where a lot of the waste originates from), only in the AR condition was the visualization directly embedded in the cafeteria itself. P15 reported that *“concerning emotions, AR is more striking because it is the most concrete, the most real, and in a familiar place. 3D on a screen is less striking because it is less immersive and represents a generic cafeteria that is not the one at Inria”*.

3.6.4 Limitations of this Experiment

This is only an initial experiment with many possible future work directions. One is to design experiments that disentangle the effects of immersion from the effects of situatedness – this would allow us to understand better if ARWAVs carry advantages compared to, e.g., data visualizations that are immersive but not situated [201]. Another possible direction is to design experiments that dive deeper into the causes of emotions. In particular, it is hard with the present experiments to know whether the reported feelings were reactions to the information presented, reactions to the technology itself, or a combination of both. For example, P2 reported being frustrated by the narrow field of view of AR and found that the natural light made it hard to see the trash bags. Such usability issues could have contributed to some of the negative feelings. Conversely, a novelty effect could have contributed to some of the positive emotions reported with AR. It remains to be seen whether the seemingly increased engagement provided by ARWAVs can be sustained or wear out as users become familiar with AR.

Other possible limitations lie in our design choices for the intermediary (desktop) condition. In particular, we chose to render the 3D trash bags within a virtual cafeteria for consistency with the AR condition (e.g., in both conditions, participants can use elements of the environment for visual size comparison), but we decided against replicating the real cafeteria to be able to test the difference between an ordinary concrete-scale visualization [94] (3D condition) and a visualization that is integrated within the relevant environment (AR condition). Alternative choices could have been to replicate the real cafeteria in the desktop condition (and possibly show it at a different location), or to show the waste accumulation on a neutral background. But all choices would have suffered from confounds – with only two baseline conditions, it is impossible to disentangle the effect of all factors, and more experiments are needed.

One more limitation of the experiments is that it measure self-reported emotions, not actual emotions. The two may differ for several reasons, including social desirability bias [133] and good-subject effects

[238]. Objective measurement methods such as skin conductance exist, but they are unlikely to be sufficiently sensitive to detect changes in low-intensity emotions. Another approach is to use a between-subject experiment design, which would require many more participants but would reduce possible good-subject effects. We also point out the limitation of the gender imbalance among our participants (15 males, 5 females). Indeed, some studies suggest that women may have a more pronounced inclination to openly express their emotions compared to men [197]. This tendency can be influenced by social and cultural factors that often encourage emotional expressiveness in women. Additionally, women may be more attentive to their emotions and possess a better ability to recognize and articulate them than men. This observation could be linked to social norms that promote the development of heightened emotional awareness in [132]. Conducting a study with a significantly larger and balanced number of participants would help us determine the influence of this factor. Finally, as we already mentioned, studying aspects other than affect, such as perception and behavior, is important.

3.7 General Discussion

In this section, we discuss other important aspects of ARWAVs to consider, as well as the limitations of our initial explorations and directions for future work.

3.7.1 Exploring Presentation Modalities for ARWAVs in Varied Contexts

When considering different ways to present ARWAVs based on the intended context of activity, it is pertinent to explore options tailored to the constraints related to device practicality. Notably, ambient visualization offers advantages by allowing augmented reality to be triggered by various stimuli. However, frequent use may lead to user habituation. It would be interesting to assess the impactful one-time effect while discussing ease of setup, varying levels of direct access to ARWAVs, and the degree of immersion, taking into account ecological trade-offs. Additionally, given the imminent commercialization potential of AR eyewear without necessarily resorting to a headset, it would be prudent to evaluate the costs associated with these devices and consider implications related to user accessibility in different usage contexts, especially when smartphone portability may present space constraints.

3.7.2 Exploring More Design Options for ARWAVs

We discussed design considerations for ARWAVs and started to explore them by building prototypes, but many areas remain to be explored. Regarding the type of visual representation for the waste material (see section 3.5), we mostly focused on literal representations and only started to explore transformed ones

(Figure 3.4d). It would be useful to explore transformed representations further, for example, by trying other unitization approaches [94]. More work is also needed on materialization, i.e., ways of representing quantities without an obvious concrete representation, such as electricity or greenhouse gases. Concerning metaphorical representations (e.g., using trees to convey paper saved), they are somehow outside the scope of ARWAVs as we defined them, but it could be interesting to see how they compare with literal representations. Another promising area is the physically realistic behavior of waste representations: our feedback session already suggested that animations and sounds (e.g., trash bags falling from the sky) can increase immersion and induce emotions. Exploring other physical behaviors and ways people can interact more realistically with the waste representations could be interesting.

In most of the ARWAV examples we considered, the waste accumulations are situated in the sense of Willett et al. [322] as they are displayed where the waste was produced, or in other terms, where the waste accumulation would have really taken place if the waste had not been carried away. Although this seems to be the most straightforward design choice, it is possible to imagine ARWAVs shown in places unrelated to the data, such as museums or exhibition halls (see subsection 3.4.2). Such ARWAVs could still help users understand waste data and make informed decisions, and could therefore be considered situated in the broader sense of Bressa et al. [81]. By the same reasoning, waste-accumulation visualizations could be displayed in virtual spaces while retaining some of their utility: such designs would be very similar to the data visceralizations proposed by Lee et al. [201].

In all these cases, there is still a lot of research to be done, not only in designing and implementing novel systems but also in empirical evaluation and comparison. Studies can be centered around psychophysics experiments focusing on waste quantity perception, or like we did with emotions, they can focus on other psychological impacts (e.g., attitude, engagement, or memorability), or they can be dedicated to the assessment of behavior change through the design of behavioral experiments.

3.7.3 Scale Issues with ARWAVs

The testers in our initial feedback session reported that viewing our larger ARWAVs was uncomfortable because they could only see a small portion of them at a time, and they had to scan them as if they were looking at it through a peephole. A combination of two factors likely caused this. One was the relatively small field of view of the AR displays (52° diagonal for the HoloLens 2, and about 20° diagonal for the smartphone if held at 40 cm), which is an important limitation of many current AR display technologies. The second factor was the relatively small distance of the ARWAVs to the user. To address this, one solution is scaling down. Based on the feedback session with colleagues, we implemented a miniature view, consisting of a 1:10 scale model of the ARWAV with a human silhouette and a car standing next to it (see Figure 3.13). As Lee et al. already discussed [201], it is impossible to get any direct notion of the true size of extreme-scale objects (e.g., planets), but a scaled-down version can show their relative sizes.



Figure 3.13 Miniature view of an ARWAV, with a car and standing man as reference points.

In the tech exhibition event, as before, when experiencing HMD ARWAVs, several observers commented on the narrow field of view and expressed the need to step back. With them, we activated the miniature view. They found it useful to be able to walk around the ARWAV and get a global overview of the volume. One observer made the point that the two objects of reference (the man and the car) were very important to get a correct idea of the size of the individual objects (trash bags and cups) and, therefore, of the entire volume. It was also mentioned that the two views could be complementary: the miniature gives a better overview, while the 1:1 visualization presents objects on a familiar scale. This kind of solution sure needs refinements and comparative evaluation.

Another solution could have been to invite our testers to move to a larger room or outside and position the waste accumulation visualizations further away from them. However, supporting AR in large spaces is difficult due to limitations in scanning technologies [38]. The Hololens 2, for example, needs to scan its physical surroundings to anchor virtual content and manage occlusion properly, and the depth spatial mapping is currently limited to 3.1 meters. More research is needed to develop techniques to realistically integrate virtual objects in the physical world far away from the user. But in principle, even very large-scale ARWAVs should be possible using computer vision, 3D maps, and/or spatial anchors (e.g., Microsoft Azure Spatial Anchors¹⁷). Once this becomes possible, we will be able to design ARWAVs that convey massive amounts of waste (see Figure 3.14). We can now explore fresh options in designing and implementing novel systems that explore trade-offs for adapting to these current technical limitations.

3.7.4 Generalizations of ARWAVs

Although ARWAVs focus on representing waste, similar techniques can be used for purposes other than waste representation. For example, ARWAVs could be repurposed to show the personal consumption of unhealthy substances that, like waste production, tend to disappear and leave no trace. For example, people may struggle to understand how much alcohol, salt, or sugar they consume over time. Although

¹⁷<https://docs.microsoft.com/fr-fr/azure/spatial-anchors/overview>



Figure 3.14 Photomontage giving a rough idea of the amount of material displaced and emitted to manufacture the smartphones for the entire population of Paris, placed next to the Notre-Dame cathedral. Credit : Léana Petiot <https://www.leanapetiot.com/>

mobile apps exist to log this information, showing cumulative consumption as numbers or bar charts cannot impact users seeking motivation to improve their behavior. Showing the same numbers using concrete AR accumulations (e.g., piles of alcohol bottles, sugar cubes or salt shakers in the kitchen) could help raise users' awareness of their consumption, make it more salient, and help them change their behavior. Similar ideas can be applied to other unhealthy habits, such as smoking (using cigarette packs). Alternatively, ARWAVs could be used to show the cumulative consumption of healthy substances like vegetables or water and provide positive feedback. Finally, ARWAVs could be used to improve people's understanding of personal material consumption more generally. For example, when buying clothes (leaving aside the environmental concerns associated with it), it can be difficult to assess the amount of clothes bought, thrown away, and sold over a period of time. Seeing them accumulated and organized in such a way as to have visibility of their life cycle may be interesting in the scope of personal behavior analysis.

3.8 Conclusion

In this chapter, we presented *augmented-reality waste accumulation visualizations* or ARWAVs, which are concrete literal representations of accumulated waste embedded in users' physical environments. We went through several examples of ARWAVs to illustrate the variety of situations where they can be useful, and we discussed general design principles and trade-offs. Initial feedback sessions with experts and exhibition visitors suggested that ARWAVs can successfully engage observers. A controlled user study involving 20 participants provided further evidence that ARWAVs can increase emotional response compared to standard ways of representing waste output. Further research is necessary to determine whether this could in turn, promote pro-environmental behavior. This research would benefit from a multidisciplinary collaboration with environmental scientists, behavioral economists, psychologists, sociologists, and science communicators. Minimizing the negative impact of humans on the environment is an incredibly complex problem, and the issue tackled by ARWAVs (giving individuals a better sense of the magnitude of their waste production) only represents a tiny subset of all problems that must be addressed. Nevertheless, ARWAVs could play a useful role if some designs have a high net ecological benefit or usefully inspire future research.

Otherwise, considering concrete representations globally, we hope to see them used in support of more abstract representations (such as numbers or graphs), which, as a participant mentioned, give a more precise and memorable idea of the data (see subsection 3.6.3). Identifying the advantages of these types of representation (abstract and concrete) led us to consider an approach that would allow them to coexist in the same immersive environment and switch from one to the other depending on users' needs. We then started to prototype transitions (as in the second video of subsection C) between holographic trash bags and bar charts and explored design implications for such transitions.

I present these explorations in the next chapter.

Chapter 4

Binding Bridges Between Abstract and Concrete Immersive Visualizations With Animated Transitions

Chapter summary

While data visualizations are typically abstract, there is a growing body of work around concrete visualizations, as tackled in the previous chapter. Concrete visualizations can complement abstract ones, especially in immersive analytics, but it is unclear how to design smooth animated transitions between these two kinds of representations. We investigate a design space of abstract and concrete visualizations, where animated transitions are pathways through the design space. The design space is defined with four axes, each corresponding to a different transformation. We consider different ways to design animated transitions by staging and ordering the transformations along these axes. In a controlled experiment conducted in virtual reality with 16 participants, we compared four types of animated transitions and found quantitative and qualitative evidence of the superiority of a specific staging approach over the simultaneous application of all transformations. We also propose a tentative set of design guidelines for future work. Our study pre-registration is available at https://osf.io/8mu73?view_only=f5ed74fd9e2346228214f5dec973b208

Main portions of this chapter have been executed during an internship at ÉTS (Montréal, Canada) under the supervision of Michael McGuffin. Thus, any use of “we” in this chapter refers to Michael McGuffin, Arnaud Prouzeau, Pierre Dragicevic, Martin Hachet, and me.




Figure 4.1 Stages of an animated transition between an abstract (a) and concrete (e) visualization. The scene shows three students’ desks and how many books borrowed by the three of them combined over the past three weeks, either with bars (a) or piles of books (e). To smoothly animate a transition from (a) to (e), the bars are decomposed into cubes (b) that are scaled and translated toward the appropriate desk (c,d) and then morphed (e) into books.

4.1 Introduction and contributions

In the previous chapter, we discussed some of the benefits of immersive concrete visualizations and demonstrated that they can sometimes evoke a stronger emotional response. More traditional *abstract visualizations*, like bar charts, have complementary advantages, such as being untethered to any physical space and enabling more precise quantitative comparisons. However, these visualizations can convey the same data as on Figure 4.1. In this example, the bar chart (a) shows clearly that most books were borrowed in the first week. However, when data about the borrowed books are shown by realistic piles of books on the desks of the respective students who borrowed them, it gives a clearer sense of who is the student who borrowed the most books and how much space the group needs to store all of them.

One strategy to combine these complementary information, and more broadly, the advantages of concrete and abstract visualizations, is to enable *animated transitions* [167] between them, so the user may freely change the way the data is shown according to their needs. The closest example we know of and that we already discussed in subsection 3.2.1.2 is Ivanov et al.’s work [180], who animate between different concrete visualizations in VR. We don’t know any previous work focusing on designing

or evaluating animated transitions between abstract and concrete immersive visualizations. Our work helps to fill this gap by investigating several variations of transitions between concrete and abstract visualizations in immersive environments. Specifically, we:

- (0) Present several scenarios illustrating the different situations in which these kind of animated transitions can be useful;
- (1) Present a design space with 4 axes (Marks, Embedding, View, and Scale), to frame our approach for designing these transitions;
- (2) Present a controlled experiment in VR, yielding quantitative and qualitative evidence that our [MEVS] condition (which displays changes in Marks and Embedding first, followed by View and Scale transformations) outperforms the [MEVS] condition (a transition showing all transformations simultaneously over the same total duration);
- (3) Propose design guidelines for immersive animated transitions between concrete and abstract representations.

4.2 Background

I already reviewed work related to [immersive] concrete visualizations beforehand (see subsection 3.2.1). In this chapter, we envision concrete visualizations as one end of a continuum from abstract to concrete. By abstract, we mean a conventional data visualization, where data values are mapped to elementary visual variables such as position, size, or color [231]. Unit charts (mentioned several times in subsection 3.2.1.1, section 3.5) lie in a particular place on that continuum as one can make bar charts – a common type of abstract visualization – slightly more concrete by turning it into a unit chart. Therefore, we start by presenting existing work on a unit chart.

4.2.1 Unit Charts

A *unit chart* is a chart that contains repetitions of visual symbols, whose number is proportional to the quantity represented [164, 245]. For example, a unit chart may represent 20,000 cars with 20 car icons, each one standing for 1,000 cars. Unit charts include *block charts* – where the symbols are abstract like circles or squares, and *pictorial unit charts* – where the symbols resemble real objects, like cars or people [164]. Block charts are slightly more concrete than bar charts because quantities appear explicitly instead of being aggregated and mapped to bar lengths. Pictorial unit charts are situated slightly further towards the concrete end of the abstract-concrete continuum because the pictograms reveal the semantics of the data being represented. IsoTypes are a prominent family of pictorial unit charts [163].

Even in pictorial unit charts, symbols can be more or less realistic, and the use of realistic symbols can make the chart more concrete. One example is an interactive web app that supports faceted browsing of the 500,000 coins of the Münzkabinett Berlin collection, where each coin is represented by its photograph [147]; Each data point thus conveys the coin’s size and appearance. Similarly, individuals can be depicted more or less realistically in unit charts about people, ranging from geometric symbols to icons to photographs to actual individuals [229].

As tackled in subsection 3.2.1.1 and subsection 3.2.1.2, a step towards more concrete visualizations can be taken by rendering realistic unit charts in 3D and embedding them in familiar 3D environments. We now review systems that consider the coexistence of abstract and concrete visualizations.

4.2.2 Hybrid Abstract/Concrete Systems

Researchers have argued that concrete visualizations – whether shown on regular computer displays or with immersive displays – can make it easier for audiences to engage with the data and to understand magnitudes and quantities on a visceral level [40, 94, 239]. However, concrete visualizations are only complements to abstract visualizations. In most situations, abstract visualizations are more suitable: they are well suited to audiences who wish to study a dataset in depth and are likely much more effective at conveying accurate values and revealing patterns in complex datasets such as multidimensional datasets and time series.

Because of the potential complementarity between abstract and concrete visualizations, it is important to understand how visualization systems can support both. Not only could such hybrid systems accommodate different audiences, but they could also support different tasks, help users answer a larger variety of questions, and give them more diverse perspectives on the same data.

Several examples of visualization systems already exist that integrate 3D models of real objects and environments with abstract visualizations, such as HANNAH [120], a system that blends 3D models of a waste water treatment plant with abstract visualizations of the plant’s process data, or Janiul et al.’s [185] system that combines 3D brain models with abstract representations of fiber tract data. However, such systems do not support abstract and concrete representations of the *same* data.

The immersive visualization of mass shooting data from Ivanov et al. [180], which we already discussed, comes close to being a hybrid abstract/concrete system, because it supports transitions from views where the user is immersed into a crowd to views where the user can examine the crowd from a distance and filter people by gender, age, or other attributes. Although all views are concrete in that they are made of 3D human silhouettes, the distant view is analogous to a unit chart.

Although the space of hybrid concrete-abstract visualizations has started to be explored, to our knowledge, no formal study investigates how to transition between concrete and abstract representations smoothly.

4.2.3 Animated transitions

Animated transitions are a type of user interface animation where abrupt visual changes are turned into rapid and smooth animations [93, 116]. Animated transitions have long been used in user interfaces [93] and have more recently been employed in data visualization applications, often to help users remain oriented while switching between different visual representations of the same data [45, 167, 245]. Animated transitions have also been used in visualization to convey dynamic data [46, 174] and to explain data analysis pipelines [160].

Recently, there has been interest in using animated transitions in immersive visualization systems, with researchers proposing conceptual frameworks and prototypes to illustrate what is possible in terms of transitions between 2D and 3D immersive visualizations [202, 203, 273, 326]. However, this work almost exclusively focuses on animated transitions between abstract (or mostly abstract) visualizations. On the other hand, the immersive visualization of mass shooting data by Ivanov et al. [180], which we discussed in subsection 4.2.2 demonstrates animated transitions between concrete (or mostly concrete) visualizations.

For now, we are not aware of systems supporting animated transitions between fully abstract and fully concrete visualizations, as we propose.

Animated transitions have been studied empirically, including in visualization. Many studies use visual tracking tasks as proxies to evaluate if the animated transitions are easy to follow. For example, Heer and Robertson [167] presented participants with animated transitions between scatter plots, bar charts, and donut charts. Each time, participants had to follow two objects and identify their position on the final visualization. The study found animated transitions to be effective, with preliminary evidence that staging (i.e., decomposing an animation in multiple steps) may help. However, a subsequent study found that some forms of staging can make object tracking more difficult [92]. Meanwhile, another study found that slowing down animated transitions at the beginning and end (sometimes called “slow-in slow-out”) facilitates visual tracking [116].


More recently, studies conducted in immersive visualization environments have started to appear. Yang et al. [326] compared animated transitions between 2D and 3D maps with a side-by-side condition, both in virtual reality. They gave participants elementary cartographic tasks. Participants were overall more accurate at the tasks and preferred the transition condition. Schwajda et al. [273] studied animated transitions between 2D node-link diagrams shown on a regular screen and mid-air 3D node-link diagrams

shown in AR. They asked participants to track clusters between the two representations and found some evidence that the animated transitions helped them perform the task. As far as we know, there is no study on how to design animated transitions between abstract and concrete visualizations in immersive environments.


4.3 Illustrative Scenarios

During repeated meetings among authors, we discussed exciting scenarios to motivate the study of transitions between concrete and abstract representations in immersive environments. All these examples present use cases that would grant users the power to manipulate abstract representations tied to their real-world existence. Here are three possible applicative scenarios we identified:

- **Scenario 1: Objects in space.** This situation pertains to data representation that refers to physical objects in the real-world space. It includes physical attributes of the object (e.g., distance, height, volume) and data related to their spatial organization (e.g., density, Body Mass Index, Biomass). For instance, representing data such as population density with standard visualization tools may be challenging. An inspiring desktop tool is MapChecking [6], which helps estimate and fact-check the maximum number of people standing in a given area. After delimiting a given area, one can increase and decrease the number of people and get an indication through a color map corresponding to either "light," "crowded," or "packed" human presence. Interestingly, the developers included three static 3D pre-computed images showing more concretely what two people per square meter look like, with a gathering of humanoids on an esplanade. Lee et al. [201] already built this kind of scene in VR to represent a fixed density of people present at a specific protest. In both cases, the user can access the concrete visual representation of one specific situation. However, none of them enables a dynamic manipulation of density values while having the concrete representation associated. This may enable non-experts to understand density metrics and experts (for instance, event planners) to parameterize venues with concrete representations. We implemented a mockup of this idea, showing a *physical scale* and a *reduced scale* (qualificatives presented in the next section), available at: <https://youtube.com/shorts/xAAuaoUaj04?feature=share>.
- **Scenario 2: Objects in movement in space.** Our following example concerns the movements of objects in space. For non-experts, it can be very challenging to associate the force of a river, the magnitude of an earthquake, or the gravity of a tornado with physical reality, even though the communicated units appear in Sverdrup via the Richter or Fujita scale. The YouTube videos discussed in subsection 3.2.1.1 are a great starting point for better understanding these data. Indeed, a representation at a *physical scale* (see subsection 4.4.1) could amplify this effect. Another highly illustrative example is crowd movements. They have been extensively studied, and their complexity



(a) Patterns made by walkers exiting a door of size $W=4\text{m}$ with an initial density of $c=0.3$. On the left: $t=100\text{s}$, on the right: $t=200\text{s}$.



(b) An example of a 3D animated on-screen concrete representation of crowd flows.

Figure 4.2 (a) is an abstract representations of simulated crowd flow data [292]. (b) is a more concrete visualization [300].

is challenging to grasp, to the extent that there are still deadly events linked to uncontrollable flows of people by causes difficult to identify. Researchers use abstract representations [292] such as Figure 4.2a and sometimes 3D on-screen more concrete representations [300] such as on Figure 4.2a. A tool providing access to concrete representations, offering multiple perspectives in real-time, and enabling navigation within abstract datasets to serve the validation of predictive models (e.g., crowd speed, movement patterns) seems promising.

- **Scenario 3: Own body's experiences.** Another applicative scenario we have identified concerns our own movements. When we perform movements as part of engaging in sports activities (e.g., martial arts, weightlifting, tennis), we can access specific abstract performance data (e.g., kick amplitude, bar trajectory, position of winning shots, or weak zones on the field). These data can be challenging to manipulate and correlate with the actual training: only skilled practitioners and experts in these sports manage to master the manipulation of such data to improve their performance, often with the assistance of a coach. At the same time, sportsmen and women sometimes use cameras or mirrors to monitor their movements and improve on them. These cameras and mirrors could easily be replaced by 3D avatars of themselves embedded in their gym, allowing them to turn around it and get a 360° view of their workouts. An app that would bind abstract metrics to avatar movements may support and empower non-experts by enabling them to understand their body performance.

4.4 Designing animated transitions between concrete and abstract visualizations

Throughout approximately 10 group meetings, we discussed examples of abstract and concrete visualizations, and realized that these could be mixed together (for example, the upper right image of Figure 4.3 shows abstract bars that are concretely positioned on a geographic map) and possibly placed on a continuum. We also discussed the transformations involved in designing animated transitions between these variations. These discussions, which were informed by an analysis of existing literature, culminated in a design space involving 4 axes, illustrated in Figure 4.4 and Figure 4.5. In those figures, to remove the details of any particular visualization, we use small cubes as proxies for the units in a unit visualization [245].

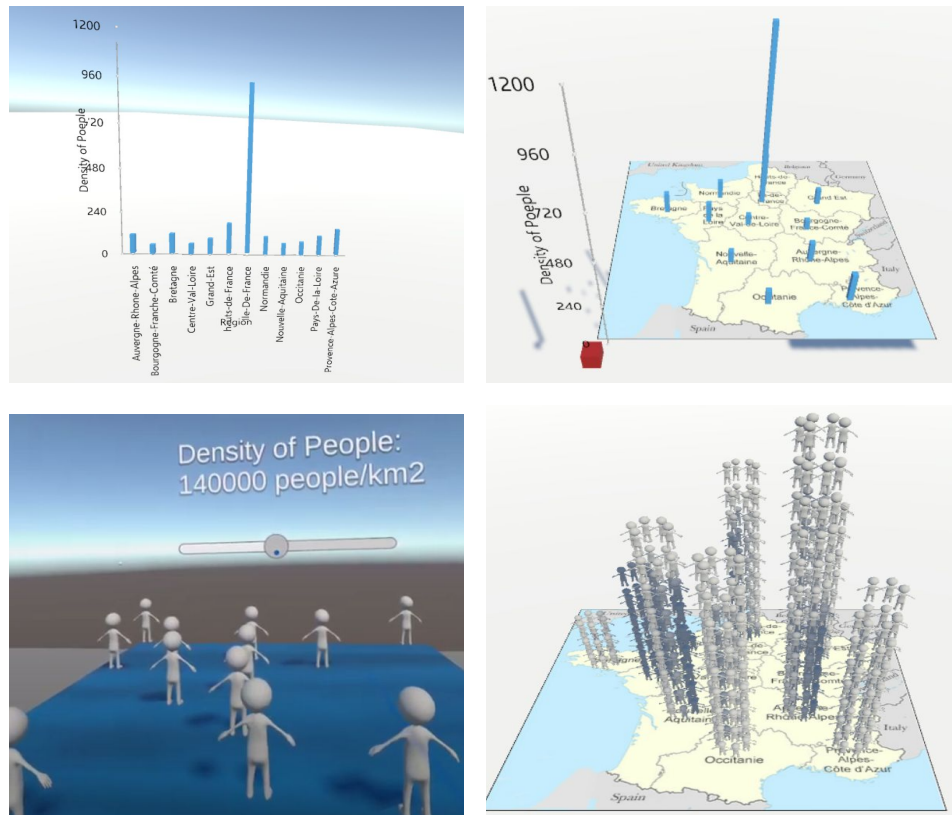


Figure 4.3 Four visualizations of population densities in VR. Clockwise from upper left, the visualization is increasingly concrete. Top left: a bar chart shows the population densities of regions of France. Top right: the same bars are positioned (embedded) on a map of France. Bottom right: the bars are replaced by stacked stickmen units (1 unit = 10^5 people). Bottom left: people matching the density selected with the slider are displayed on a ground plane (these human figures could be shown at a reduced scale to provide an overview of a unit of geographic area or scaled up to full size to more closely simulate a viewer’s real-world experience).


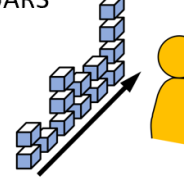
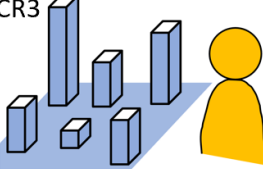
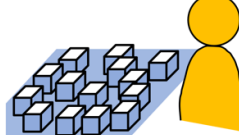
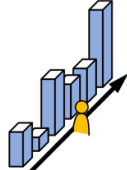
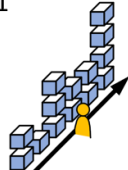

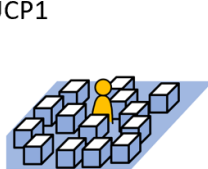
		Marks	
		Aggregated	Unit
Scale is Reduced, View is 3rd person	Embedding is Abstract	AAR3 	UAR3 
	Embedding is Concrete	ACR3 	UCR3 
Scale is Physical, View is 1st person	Embedding is Abstract	AAP1 	UAP1 
	Embedding is Concrete	ACP1 	UCP1 

Figure 4.4 Our design space. Unit Marks show individual items (in this example, cubes) whereas Aggregated Marks show collections (in this example, bars); these are positioned according to some physical or geographic space (Concrete Embedding) or by some algorithm (Abstract Embedding; in this case, a bar chart axis); View may be 3rd person (in front of the user) or 1st person (surrounding the user); The scale may be Reduced (small enough for comfortable viewing) or Physical (e.g., life-size). Each 4-letter acronym identifies a variant within Marks \times Embedding \times Scale \times View = $\{U,A\} \times \{C,A\} \times \{R,P\} \times \{1,3\}$.

4.4.1 A Design Space for Concrete and Abstract Visualizations

To design animated transitions between abstract and concrete visualizations, we define a space that enables comparisons along a common set of dimensions. Then, we will be able to define a set of transformations that must be performed on the visualizations to create seamless animated transitions. We picture our Design Space through Figure 4.5 on which each vertex is a visualization referred to with an acronym of three letters and a digit corresponding to their value for the four dimensions we introduce in this section (e.g., the first letter refers to the Marks, if a visualization present aggregated marks, its acronym will start with an A).

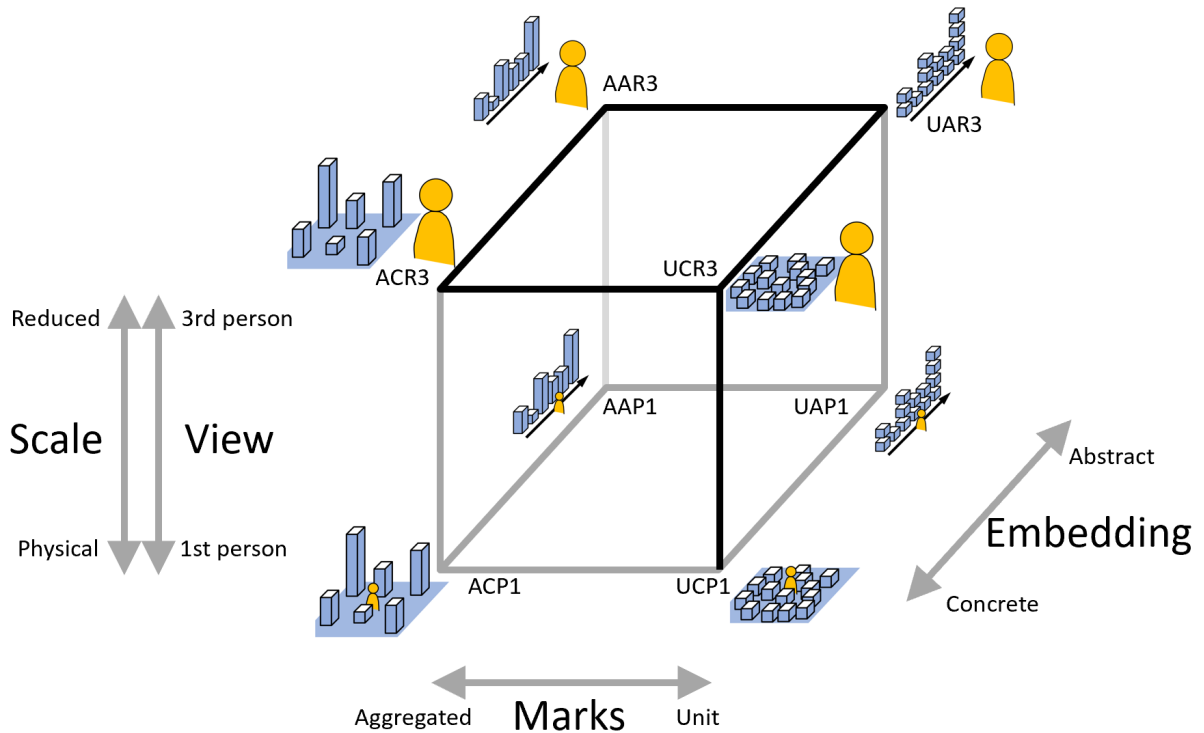


Figure 4.5 A visualization of our design space. Because the V and S axes normally vary together, they are shown as a single spatial dimension.

- **Scale:** Abstract visualizations such as Scatterplots or Barcharts are usually displayed at **reduced** scale (e.g., AAR3 in Figure 4.4 and Figure 4.5, Figure 4.3). However, Danyluk et al. [110] evaluated what they call *room-scaled* visualizations (see ACP1 and AAP1 on Figure 4.4 and Figure 4.5) that *allow viewers to explore data by physically walking through, under, and around it* through range, order and comparison tasks. Unfortunately, they found that these performed worse than with smaller scales. Still, such larger scale visualizations, which we call **physical** scale, are relevant for concrete visualizations (UAP1 and UCP1 on Figure 4.4 and Figure 4.5) as discussed in subsection 3.2.1. As a result, we consider both scales in our design space.
- **View:** when a user visualizes data, the visualization is usually located in front of them, through a **3rd-person view** (i.e., on a screen as AAR3 on Figure Figure 4.4 and Figure 4.5). In cases similar to the one described just above for *room-scaled* visualizations, the user can be surrounded by graphical elements (see UCP1 on Figure 4.4 and Figure 4.5, Figure 4.1.e.). In this case, we refer to as **1st-person view**.

Two planes on the cube have not been depicted in our representation. These encompass four more visualizations that present visualizations at a human scale from a 3rd-person view and an additional four that pertain to a reduced scale observed from a 1st-person view. However, **Scale** and **View** often vary together.

- **Marks:** Munzner defines marks as “*basic geometric elements that depict items or links*” (e.g., point, line, area) [231]. In the case of a concrete visualization, each visual **mark** refers to an object (UCR3, UAR3, UAP1 and UCP1 on Figure 4.4 and Figure 4.5), whereas in the case of abstract visualizations, usually depicting aggregations, a visual mark shows a quantitative value for an attribute (ACR3, AAR3, AAP1 and ACP1 on Figure 4.4 and Figure 4.5). Aggregates and unit marks may come in different colors and shapes (e.g., the representation of a physical object and a scatter plot data point, respectively) and go beyond what we depict in Figure 4.5. For simplicity, we kept variables such as color, texture, and shape constant and do not discuss transformations such as morphing or color changes.
- **Embedding:** If the visualization’s mark positions derive from the physical world characteristics, the embedding is **concrete**; if not, it is **abstract**. Indeed, an abstract visualization is relatively confined in a restricted area of space (as (a) on Figure 4.1 or AAR3 on Figure 4.4 and Figure 4.5), whereas a concrete one tend to have greater disposition to exist at multiple specific locations (see section 2.2, (b) on Figure 4.1 or UCP1 on Figure 4.4 and Figure 4.5).

We clarify that what is referred to as “*the spatial substrate in which marks are embedded.*” by Heer and Robertson [167] also changes regarding our axes. In both [167] and [208], they mention that scaling the visualization may require a different scale for the substrate, which is often a graph axis. In our case, the substrate differs when the marks are abstractly or concretely **embedded** as the latter demands that the marks exist at a specific spatial location. For instance, in Figure 4.1, marks exist in a 2D space on (a) and in a 3D space of a specific shape on (d), which results in different substrate dimensions and shape (respectively a graph axis and the virtual ground).

4.4.2 Initial Design Guidelines

We use the term transition to refer to a pathway through the design space of the previous section (Figure 4.5). Each transition involves one or more transformations, which may be changes to the **Marks**, **Embedding**, **View** or **Space** (denoted with the first letter of these words). More generally, a transformation may be a translation, rotation, morphing, or fading in/out of elements. The transformations in a transition may be simultaneous or may be separated into sequential stages, causing the transition to pass through intermediate vertices in the design space.

Since there are many possible transitions (i.e., pathways) through the design space to get from one vertex to another, it would be helpful to have some rules for choosing the best transitions. We propose some tentative design guidelines (**DG**) for these, some of which are supported by previous work. Some of these guidelines conflict with each other and cannot all be complied with. As we will see later, different transitions constitute different tradeoffs with these guidelines.

▷ **DG1** Use stages: divide transitions into stages to show different kinds of transformations at different times.

SpaceTrees [252] are the earliest example we know using animated transitions broken into stages in an information visualization. Stages can help the user understand a transition by showing a simpler set of transformations at each moment.

▷ **DG2** Move slower: avoid high-speed translations over long distances. Supported by [92]. (This conflicts with DG1.)

The user cannot follow very fast translations, making it difficult to understand the relationship between items before and after a transition. If transitions have a fixed total time (as is the case in our experiment in section 4.5), there is a tradeoff between DG1 and DG2: having more stages leaves less time for each stage, which means that translations will be at a higher speed. We compared transitions comprising one, two, and three stages in our experiment, enabling an indirect exploration of this tradeoff.

▷ **DG3** Move fewer: reduce the number of visible items during translations. For example, items to be removed by the end of a transition can be removed before translating other items, and items to be added by the end of a transition can be added after others are translated. This idea is used in [221, 252].

▷ **DG4** Move each subset as one: If many items must undergo individual translations, and these items form a meaningful subset, it is preferable to first translate the entire subset as a single rigid body before applying individual translations to the items in the subset.

An example of DG4 is seen in the “hybrid” transitions of [159] when a parent node must be translated, and its children must also be permuted to new locations within the parent: the parent and its children are first translated as a single object, and then the children undergo smaller translations to their new locations. This results in one coarse-grain motion first, followed by finer-grain motion of individual items, allowing the user to progressively narrow their attention, as desired.

▷ **DG5** Move smaller: if items must be translated and must also change size, it is preferable to order stages so items are smaller during translations (e.g., shrinking before translating or translating before growing). Supported by [167] in the vain of limiting occlusion.

▷ **DG6** Move aggregations: if items must be translated and must also be aggregated into a whole (or disaggregated into parts), it is preferable to translate aggregations rather than parts. For example, aggregate many items into a whole before translating or translating before disaggregating into parts. This guideline can be seen as a special case of DG3, but may sometimes conflict with DG5 if the aggregations are visually larger than their parts.

DG3-DG6 are all ways to simplify the appearance of motions and/or reduce inter-item occlusion by reducing the size or the number of moving items.

▷ **DG7** Use exocentrism: when transformations involve all items, it is preferable to perform them within the user’s field of view.

DG7 recommends using an exocentric view (a 3rd-person view, in front of the user) rather than egocentric (a 1st person view, where the user is surrounded by items). This reduces the amount of head and eye rotation necessary to perceive the items. Exocentricity can lead to higher accuracy in judgment of target location [319] and is often the best frame in terms of visualization task accuracy in 3D virtual environments [220, 310].

▷ **DG8** Slow down complexity: slow down translations when items are closely spaced and/or when items will be moving in different, unpredictable directions.

DG8 is motivated by reasoning similar to that behind the ‘adaptive’ animation technique of [116], where the animation is slower when the scene is more complex.

The next two guidelines were formulated after seeing the results of our pilot study (subsection 4.5.7). They were not initially obvious to us, as they directly contradict some of the previous guidelines, but they came out of wondering why certain transitions seemed to perform better than others.

▷ **DG9** Transform all as one: if items must be translated and scaled to extend beyond the user’s field of view, items can first be translated to their new relative positions while still within the user’s field of view during a first stage, and then in a 2nd stage, the entire collection of items can be scaled up as a single object, maintaining the relative positions of items during this 2nd stage. In the reverse direction, the entire collection can first be scaled down as a single object to fit in the user’s field of view (1st stage), and then items can be individually translated to new positions (2nd stage). This conflicts with DG5.

DG9 contradicts DG5 in deliberately varying the size of items as they are translated, rather than only translating items at their smallest size. The potential advantage of DG9 is that the user may be able to remember an item’s placement with respect to nearby items before a transition, and hence even if the user loses track temporarily of an item during the translation, the user might more easily find the item again by looking for its remembered location with respect to nearby items, e.g., by recognizing the ‘shape’ of a subset of items.

▷ **DG10** Don’t move aggregations: if items must be translated and must also be aggregated into a whole (or disaggregated into parts), it is preferable to translate parts rather than aggregations. For example, disaggregate into parts before translating or translate items before aggregating them into a whole. This guideline may support DG5 if aggregates are visually larger than their parts but conflict with DG3 and especially DG6.

The potential advantage of DG10 is that, by breaking up large aggregates before translating individual items, this might make it easier for the user to identify an individual item during motion, and/or reduce

occlusion because the items are smaller than their aggregations.

As we will see in the next section, not all of these design guidelines are directly applicable to the choice of transitions that we made. We nevertheless enumerated the set of ten above to make it clearer how some guidelines support or conflict with other guidelines.

4.5 Study

In chapter 3, we discussed the potential of using AR to create visualizations in the real world. However, for practical reasons and because this type of study on immersive animated transitions hasn't been explored much in VR, we decided to start with a VR environment. We describe our study and results in this section.

4.5.1 Choice of transitions for study

Of the eight vertices in the design space, there are three that are of less interest: (ACP1), (UAP1), and (AAP1). These are at the back and bottom of Figure 4.5, and correspond to large-scale visualizations with either aggregated marks or an abstract embedding, which seem to us to offer little theoretical advantage. [110] evaluated a condition similar to (ACP1) and found that it performed poorly compared to smaller scale visualizations.

We, therefore, focus on transitions involving the 5 remaining vertices. The most challenging transitions involve all 4 transformations (M, E, V, and S), i.e., transitions between diagonally opposing vertices. The only transitions like this within the 5 vertices of interest are transitions between (AAR3) and (UCP1). Figure 4.6 shows the seven transitions that we chose for evaluation.

The first transition, [MEVS], shows all transformations simultaneously in a single stage. The others involve two or three stages. In choosing transitions to evaluate, we eliminated transitions involving four stages, and also eliminated transitions passing through any of the three vertices ((ACP1), (UAP1), and (AAP1)) on the bottom that seems less useful. This results in transitions that never perform the V or S transformations before M or E, which makes sense if we want to avoid changes happening outside the user's field of view (DG7). The [ME-V-S] transition is the only one that separates the V and S transformations (DG5); all other transitions we chose perform V and S within the same stage.

Figure 4.7 shows the extent to which each of the seven transitions implements the design guidelines in subsection 4.4.2.




Figure 4.6 The seven animated transitions evaluated in the pilot. Of these, the four highlighted in green ([MEVS], [M-EVS], [ME-VS], [M-E-VS]) were retained for the full experiment.

4.5.2 Choice of the task

As we saw in subsection 4.2.3, animated transitions can be evaluated with low-level tasks (e.g., tracking individual objects) or with high-level tasks (e.g., asking users questions about a dataset). In this experiment, we chose a low-level tracking task in order to achieve more experimental control (we can fine-tune the characteristics of the trials) and more statistical power (the short duration of the task allows us to give participants more trials and thus record more observations).

The justification for using object tracking tasks to evaluate animated transitions is two-fold: *i)* such tasks are likely to be performed in many hybrid-concrete visualization systems: for example, in the scenario of Figure 4.1, the user may want to know which student borrowed the more books during the first week and where their desk is; *ii)* Object tracking is likely a low-level component of many higher-level tasks: if a user cannot even follow individual objects during animated transitions, it is unlikely that they will be able to perform higher-level tasks. Consequently, object tracking has been used in many previous evaluations of animated transitions [92, 116, 167, 223].

CHAPTER 4. BINDING BRIDGES BETWEEN ABSTRACT AND CONCRETE
IMMERSIVE VISUALIZATIONS WITH ANIMATED TRANSITIONS

Relevant Design Guideline		1 stage	2 stages			3 stages		
		[MEVS]	[M-EVS]	[E-MVS]	[ME-VS]	[M-E-VS]	[E-M-VS]	[ME-V-S]
DG1	Has stages, with different transformations in each stage, possibly making each stage easier to understand		✓	✓	✓	✓	✓	✓
DG2	Has fewer stages, which reduces the speed of long-distance translations	✓	✓	✓	✓			
DG5	The V iew is changed before increasing the S cale from R educed to P hysical (and in the backward direction, the S cale is decreased from P hysical to R educed before changing the V iew), avoiding occlusion during the change in V							✓
DG6	Changes the E mbedding of bars before disaggregating them into cubes (M) (and in the backward direction, aggregates cubes into bars (M) before changing their E mbedding), so that fewer items are moving during the change in E			✓			✓	
DG7	Changes M arks and/or their E mbedding at R educed S cale and in 3 rd person, so that changes to M and E may be easier to see		✓	✓	✓	✓	✓	✓
DG9	Cubes have the same relative positions and same relative sizes before and after the change in S cale				✓	✓	✓	
DG10	Disaggregates bars into cubes (M) before changing their E mbedding (and in the backward direction, changes the E mbedding of cubes before aggregating them into bars (M)), to reduce occlusion and make units more distinct		✓			✓		
	Retained after pilot study with 4 participants	✓	✓		✓	✓		
	Best in the final experiment with 16 participants				✓			

Figure 4.7 Each transition (column) that complies with a design guideline (row) is indicated with a checkmark. Partial compliance is shown with a smaller checkmark.

In our task, a single item (i.e., a cube, possibly within an aggregated bar) is designated with a visual highlight (Figure 4.8). The highlight is then removed, an animated transition plays out (Figure 4.9), and the user must indicate where the item is now located by selecting it. We did some early testing with simultaneously tracking of 2 items, but found this too difficult (especially when two targets moved to opposite sides of the user at the same time), and thus limited our experimental task to single-target tracking.

4.5.3 Transfer function and duration

We define the transfer function as the function that maps a fraction of the time of a stage to the fraction of the progress of a transformation. Both fractions are numbers in the interval [0,1].

During early testing, our transfer function was linear, and we found it subjectively difficult to track




Figure 4.8 A user at the start of a Forward-direction trial seeing the target unit cube highlighted.




Figure 4.9 In the pilot and full experiment, all transitions were between these two visualizations, with half of the trials in each direction. In the case of 2- and 3-stage transitions, there were also intermediate states not shown here.

items when they were small and densely packed, and also when they were moving at high speed over a long distance. For instance, in a Forward-direction transition (Figure 4.9), after items are disaggregated, the items are small and closely spaced, before traveling at relatively high speed to reach their destinations, which could be to the side of, or even behind the user, requiring them to rotate their head. We reasoned that slowing down the initial motion would facilitate tracking (DG8), even at the cost of making the subsequent motion faster. Indeed, the user must be able to initially distinguish the target from other small items nearby. Then, once the user has seen an item’s general direction of motion, they could estimate where it might end up with less difficulty since the cubes are more spaced out later in the transition.

To determine how to modify the transfer function, our software was instrumented to allow the user to interactively adjust two control points that define a piecewise linear transfer function. Two of us (co-authors), each independently, viewed several repetitions of a transition and adjusted the control points to our satisfaction, for a stage showing translations and lasting a total of 1 second. The resulting transfer functions are shown in Figure 4.10. We finally chose the S-shaped function (Figure 4.10, right) for our experiment because it is more similar to commonly-used ‘slow in, slow out’ animations [116]. It is also symmetrical and, therefore, easier to implement in software without regard to the direction of a transition. It results in an aesthetically pleasing deceleration at the end of the transition.




Figure 4.10 Two transfer functions independently tuned by two co-authors. The one on the right was chosen for the experiment.

To determine the duration of our transitions, we looked to previous work where transitions lasted between 1.25 seconds [167] to 10-15 seconds [273]. We also analyzed screen-captured videos to measure the time it took for tooltips to appear in a variety of software applications on Microsoft Windows, and found that the time varied from 450ms to 850ms. Informal testing with our software revealed that if a stage lasted less than 1 second, it was often too difficult to follow. We, therefore, fixed the total duration of our transitions to 3 seconds. Hence, transitions with one, two, or three stages allocate 3.0, 1.5, or 1.0 seconds to each stage, respectively.

4.5.4 Hardware

Users wore an untethered Meta Quest Pro headset in VR mode (i.e., without video pass-through). The headset was adjusted to their interpupillary distance. Users could see some of the physical world through peripheral vision, as the headset does not completely block the physical world from view. Users also had two handheld controllers to point at objects and interact with our software.

We measured a frame rate of 71 fps when using our experimental stimuli.

During the experiment, participants were free to walk around a physical space measuring 10×10 feet (Figure 4.8, right), although there was a starting position marked on the floor that users were asked to return to at the start of each trial.

4.5.5 Details of the task

Each Forward-direction trial involved showing the user a transition from a bar chart in front of the user (AAR3) to a set of large cubes surrounding the user (UCP1), as shown in Figure 4.9 and Figure 4.11. In Backward-direction trials, the transition happened in the opposite direction (Figure 4.12).




Figure 4.11 Example of a Forward direction trial. (a) Pre-animation phase: the target cube to track is highlighted on the bar chart. (b) The animated transition from abstract bar chart to concrete visualization occurs. (Not shown) After the transition, the user selected the wrong cube and confirmed their choice. (c) Post-animation phase: a red cross indicates the answer is wrong, and a blue arrow indicates that the correct cube is out of the participant’s view. (d) The user turns back and sees their selected cube in blue and the correct answer in green.

Each bar chart was composed of 6 bars, made of a total of 35 unit cubes, each corresponding to one of the 35 large cubes that would surround the user in the (UCP1) state. Each bar of the bar chart was between 1 and 10 unit cubes tall, with the unit cubes distributed randomly between bars and each unit cube measuring 6cm tall. The bar chart was positioned 1.2m in front of the user’s starting position, and 1.3m above the ground plane.

Each large cube in the (UCP1) state was randomly positioned in a 10×10m square on the ground plane, and each measured 60cm tall since this is of the same order as many immersive visualizations and it requires a change in size during the transition of 10×, making the design of the transition more challenging.

At the start of each trial, a pre-animation phase lasted 5 seconds. During this phase, one of the unit cubes in a bar or one of the large cubes was designated with a highlight to indicate that it would be the target. In the case of (UCP1), sometimes the designated target cube could be behind the user, so an arrow positioned in front of the user, pointed to where the target cube was, helping the user find it.




Figure 4.12 Example of a Backward direction trial. (a) Pre-animation phase: the blue arrow points toward the cube to track, situated out of the user’s field of view. (b) The participant turns around to see the cube to track, highlighted in blue. (c) The transition to bar chart occurs. (Not shown) After the transition, the user selected the correct unit cube in the bar chart. (d) Post-animation phase: a green check mark above the bar chart indicates that the answer is correct.

At the end of the pre-animation phase, the highlight was removed from the target, and the transition was played out, lasting a total of 3 seconds.

Next, the user had to select the large cube or (in the case of a Backward trial) select the small unit cube, which they believed was the target. They performed this selection by pointing a ray cast from a handheld controller at the cube and pressing a button. In case of a mistake, they could optionally select again to redo their selection. They confirmed their choice by selecting a special “Validate” button in the 3D scene.

Next, there was a post-animation phase lasting 5 seconds, during which time visual feedback indicated to the user if their selection was correct or not, and if not, a highlight showed the correct target cube. This completed the trial.

4.5.6 Variables Measured

The time for each trial was measured from the end of the animated transition to the time of the last selection the user performs before selecting the “validate” button.

For each trial, we recorded whether the user selected the correct target or not (as a boolean flag), to compute an error rate for each user in each condition. We also recorded the magnitude of the error in each trial in a few different ways. For our pilot study, we used the distance between the correct target and the cube selected by the user and recorded this as both a euclidean error distance (in meters) and an angular error (where the angle is calculated with respect to the user). In addition, the euclidean error distance was divided by the minimum distance between two cubes to obtain a normalized error distance.

After the pilot study, we borrowed an idea from [138] and used a different way to quantify the magnitude of an error: for each trial in the full experiment, we recorded the number of cubes closer to the correct target (including the correct target itself) than the cube selected by the user, and we called this quantity **ItemsCloser**. If the user selects the correct target T , then **ItemsCloser** is zero; if the user selects a cube R that is close to T , with no other cube closer to T , then **ItemsCloser** is one. In other words, **ItemsCloser** is the number of responses that would have been better than the user’s response. We propose this as a better way of quantifying errors when users are asked to select from a set of discrete possibilities.

4.5.7 Pilot Study

Our pilot compared all 7 conditions in figure Figure 4.6, with 4 users.

The normalized error distances indicated that three of the multi-stage transitions might be better than the [MEVS] transition, while the other three multi-stage transitions appeared worse than [MEVS]. Since the [MEVS] transition would be a useful baseline for the final experiment, we kept it along with the three most promising multi-stage transitions, namely [M-EVS], [ME-VS], [M-E-VS], for the final experiment.

We wondered why these three particular multi-stage transitions might perform better than the others, and realized that each pair of them had some theoretical advantage. The first two transitions, [M-EVS] and [ME-VS], have the advantage of only having 2 stages, making translations slower than in 3-stage transitions (DG2). The last two transitions, [ME-VS] and [M-E-VS], transform View and Scale in the same stage, meaning that even if that stage involves some rapid translations and the user loses the target temporarily, the user might be able to remember a target’s position relative to other elements in the previous stage, and find the target again later based on the relative positions of elements. This follows DG9 and is illustrated in Figure 4.13, top. Finally, the first and third transitions, [M-EVS] and [M-E-VS],

follow DG10: disaggregating bars before translating units may reduce occlusion and make an individual unit easier to follow. The other design guidelines (Figure 4.7), some of which conflict with these, may be less important in this context.




Figure 4.13 In the [ME-VS], [M-E-VS], and [E-M-VS] conditions, one stage of the transition transformed View and Scale together (top), maintaining the relative positions and distances of elements, as if the entire scene was a single object being scaled up. In the [ME-V-S], by contrast, View and Scale were transformed in sequence (bottom), making elements appear more sparse in the middle of these transformations.

4.5.8 Full Experiment

Our final, full experiment was preregistered ¹, declaring our sample size, hypothesis, and analysis.

A sample size of 16 users was chosen in the preregistration, not including the 4 pilot participants. Of the 16 users, 10 were women, 6 men; 14 right-handed, 2 left handed, but all with a habit of using the mouse with their right hand; age 20 to 32 years (average 25.6); inter-pupillary distance (IPD) 58 to 66mm (average 61.8).

Within each condition, the user experienced an equal number of trials covering 2 directions \times 8 repetitions in random order, for a total of 4 conditions ([MEVS], [M-EVS], [ME-VS], [M-E-VS]) in random order \times 2 directions (Forward, Backward) \times 8 repetitions \times 16 users = 1024 trials, not counting warm-up trials.

We **hypothesized** that the [MEVS] condition will result in a greater error (in terms of ItemsCloser, defined in subsection 4.5.6) than each of the three other animated transition conditions.

¹https://osf.io/8mu73?view_only=f5ed74fd9e2346228214f5dec973b208

4.5.9 Results

4.5.9.1 Planned Analysis

For these results, advice was adapted from Dragicevic [114]. We present effect sizes visually with confidence intervals (CIs), following Tips 15 and 16 in [114]. Our CIs are computed using one (averaged) value for each (user, condition) pair (Tip 9). The use of null hypothesis significance testing (NHST) can cause misleading, dichotomous thinking [33, 107, 115, 204] and we decided to forego any NHST in our analysis (Tip 25 in [114]).

Figure 4.14 shows the result of our preregistered analysis. Of the three multi-stage transitions that were compared, we find evidence that [ME-VS] is better than [MEVS] in terms of ItemsCloser.




Figure 4.14 ItemsCloser quantifies the size of errors (subsection 4.5.6). The upper part of the chart shows the four conditions, where each dot shows the average for one user, and each bar is a 95% CI computed with bootstrapping. The lower part of the chart shows paired differences, revealing evidence that [ME-VS] yields smaller errors than [MEVS].

4.5.9.2 Subjective Feedback and Exploratory Data Analysis

Users were asked which condition they liked the most and the least, and were asked to answer four questions on a 7-point scale²: “What mental effort was required?”, “What physical effort was required?”, “How much frustration did you feel?” (1 for “little”, 7 for “a lot”) and “How did you find the speed of the transition?” (1 for “too slow”, 7 for “too fast”). Figure 4.15 shows the results. In the pilot, the 7 conditions were grouped into sets of 1-, 2-, and 3-stage transitions, which users were asked to compare. In the final experiment, [ME-VS] was preferred by most users, and received the most favorable scores for mental effort and frustration.

²Find the full experiment questionnaire in subsection C

CHAPTER 4. BINDING BRIDGES BETWEEN ABSTRACT AND CONCRETE IMMERSIVE VISUALIZATIONS WITH ANIMATED TRANSITIONS

	Conditions	Number of users		Likert score (1 to 7)			
		Most favorite	Least favorite	Mental effort	Physical effort	Speed	Frustration
Pilot (4 users)	[MEVS]	0	1	4.0 3.8	1.5 2.3	3.0 3.0	2.0 2.5
	[M-EVS], [E-MVS], [ME-VS]	4	0	3.5 3.5	2.0 2.0	4.0 4.0	2.0 2.5
	[M-E-VS], [E-M-VS], [ME-V-S]	0	3	5.5 5.5	2.5 3.3	5.5 5.5	4.5 4.0
Final (16 users)	[MEVS]	3	3	3.5 3.6	1.0 1.8	4.0 4.4	2.5 3.0
	[M-EVS]	0	6	5.0 4.7	1.5 2.4	5.0 4.9	4.5 3.8
	[ME-VS]	10	0	2.0 3.0	1.0 1.9	4.0 3.9	1.0 1.4
	[M-E-VS]	3	7	5.0 4.8	2.0 2.3	5.0 5.1	3.5 3.8

Figure 4.15 Subjective ratings by users. Blue shows counts, green shows averages, black shows medians, and grey shows interquartile ranges.

7 out of the 16 users said that trials were more difficult when the target moved behind them because they had to turn around quickly. 5/16 users said they easily lost track of the target when items crossed trajectories, creating occlusion.

However, 8/16 users remarked that when items re-arranged themselves in front of the user (i.e., with a 3rd person view), this helped them to track the correct target. This echoes DG7 and was more the case with [ME-VS] and [M-E-VS]. Of these two conditions, [ME-VS] only has 2 stages, thus displaying each stage more slowly than [M-E-VS]. This helps explain the success of [ME-VS].

In Figures 4.16 to 4.18, as with the top part of Figure 4.14, each dot shows the average for one user, and each bar is a 95% CI computed with bootstrapping.




Figure 4.16 ItemsCloser, broken down by direction.

Figure 4.19 shows that, unsurprisingly, error rates generally increased when the target was moving at a higher speed, except for an apparent drop in error rate at the highest speeds achieved, which occurred with the only 3-stage transition, [M-E-VS].

8/16 users stated that trials were more difficult when the target traveled very close to them (these were trials where the target traveled through a nearly 180 degree angle relative to the user, and achieved the highest angular speed). Interestingly, one user stated that it was easier to track the target when it




Figure 4.17 The time (i.e., duration) of each trial.




Figure 4.18 Error rates, computed from the boolean result of each trial either being successful or not.

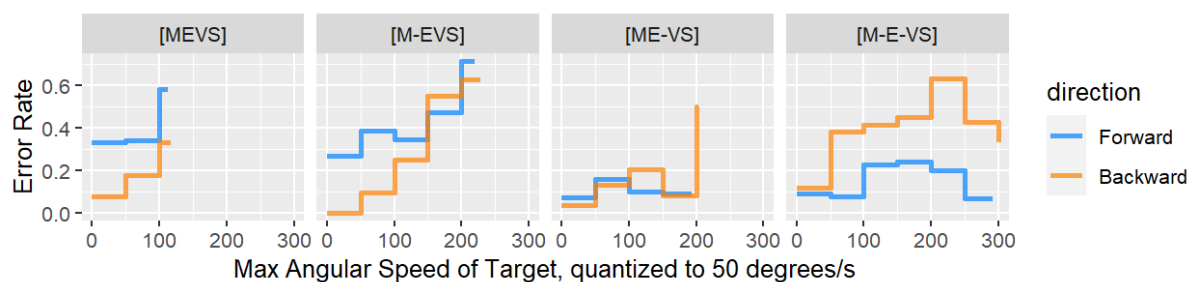


Figure 4.19 Trials were binned by max angular speed of the target (a bin size of $50^\circ/s$ was chosen using the Freedman-Diaconis rule). The average error rate is shown on the vertical axis for each bin. The right extremity of each curve indicates the highest max angular speed achieved.

went through their body, despite its high speed, because this gave them a precise indication of the target’s trajectory and how to find it after turning around. This may explain the drop in error rate observed at the highest angular speed in the [M-E-VS] condition in Figure 4.19.

4.5.10 Discussion

4.5.10.1 On the design of animated transitions

The apparent success of the [ME-VS] transition may be due to the importance of DG2 (i.e., not having 3 stages), DG7 (i.e., M and E transformations occurring in 3rd person), and/or DG9 (“Transform all as one”) for the experimental conditions that were tested, despite the conflict between DG9 and DG5 (“move smaller”) and not supporting DG10 (“Don’t move aggregations”).

Rather than eliminate or claim to have confirmed any particular guidelines, we propose that our guidelines must still be viewed as tentative and would require additional evaluations of various kinds to understand better. The rationale behind each guideline may nevertheless be useful for generating ideas to design future transitions.

We furthermore propose a few additional guidelines that could be investigated in future work:

▷ **DG11** Make units distinct: make units inside aggregations visually distinct. (This could be a way to combine the intended advantages of DG6 and DG10 by making individual units easier to track while still keeping the number of moving objects small.)

▷ **DG12** Use transparency: make items and aggregations semi-transparent.

▷ **DG13** Show tracks: display trails or line segments showing where elements will move and/or where they came from after a translation. (This is comparable to ideas in 2D in [58, 65], and comparable to the visual feedback in [179].)

We observe several ways to group our tentative guidelines according to the underlying motivation behind them: they simplify the complexity of transformations to make them easier to interpret (DG1), reduce the angular speed of translations and/or the need for large rotations of the head or eyes (DG2, DG7, DG8), reduce the density of items (DG3, DG6) simplify motions so they are easier to interpret (DG4, DG6, DG9), reduce inter-item occlusion (DG5, DG10, DG12), or otherwise facilitate tracking of individual items (DG10, DG11, DG13).

Our designs and analyses stem from the conceptual framework we established earlier. A different and complementary approach would be to study ad-hoc choices that influence our results. We could, for example, test different speed functions or selection/pointing methods for each transition/representation before comparing the animated transitions. All of this raises further research questions.

4.5.10.2 On the implementation of animated transitions in AR

Implementing these transitions in Augmented Reality brings new technical challenges. Indeed, the visualization support becomes the physical world, much more cluttered, diverse, and rich in objects than our virtual reality scene. The system must track physical objects in real time to integrate concrete representations into space. Unlike our scene where objects are arranged in 2D, elements can be placed on certain real objects (such as a table in the teaser example) in a third dimension. Furthermore, the user's perception is complicated by the presence of physical objects that he/she will have to differentiate from virtual objects in the scene.

By using VR and overcoming the technical limitations outlined above, we do not ensure the transferability of the study results. However, the technique recognized as the most effective is the one that offers a reduced third-person view during the animated transition. This view seems beneficial due to the scale and dimensions of the visualization space, analogous to a real scene. The user can benefit from a transitional state that allows them to predict where objects will be positioned in the extended space during an animated transition that requires significant marks movements.

4.6 Conclusion

We have presented a design space with four axes that characterize concrete and abstract visualizations. This enables us to identify a set of transformations necessary to design animated transitions between these two types of visualizations. Within the pool of possible transitions, we compared four of them in a controlled study in VR and found quantitative and qualitative evidence in favor of one particular transition. We also proposed a set of tentative design guidelines that may inform future design strategies.

4.7 Future Directions

Although we investigated how to enable the coexistence of abstract and concrete visualizations through animated transitions, other exciting strategies can be envisioned. For instance, concrete and abstract visualizations of the same data could exist side by side and provide visual *links* between them in the way of Prouzeau et al.'s work [255]. Considering our animated transitions design choices, many other options were not tested. For instance, using cube proxies lets us put aside questions associated with abstract marks morphing into more complex objects. Future research could focus on identifying these undiscussed transformations and how and when they could happen within an animated transition. This also pertains to how those parameters (e.g., shape, color, material) could fit in our design space.

Otherwise, our transitions could involve supplementary design elements. For instance, the display of marks trajectories' predictions [179] or traces in space may help track their back-and-forth movements (supported by DG13). Additionally, playing with the transparency (DG12) of non-essential (or chosen) objects could prevent occlusion between marks during transitions. Potentially, people can miss portions of the animations when an object goes outside their field of view; future work can explore techniques similar to Phosphor [58] or Mnemonic Rendering [65] to show or replay what has been missed. Also, we argue that animated transitions are not necessarily symmetrical [202]. It would be interesting to test if the use of the forward and backward transition of the same type is more efficient than a combination of the most efficient forward and the most efficient backward transition. Besides, considering a specific layout, a particular type of data, or characteristics of the concrete visualization, the most efficient animated transition may change.

Furthermore, even if our choice of task is a good starting point to evaluate these animated transitions as it is common practice among the community, this could surely benefit from further analysis with higher-level task evaluations. In our experiment, the task is performed in a very cleaned-up virtual reality scene, whereas transitions may exist in much more busy scenes like the teaser's one Figure 4.1. The proposed task could have been performed with cubes' positions that vary along the vertical axis, closer to a real-world setup. Indeed, as discussed in chapter 3, displaying concrete visualizations in AR holds great promise; very next steps encompass conducting a similar study in AR, utilizing the real-world environment.

Also, we scratch the surface of an interesting research topic by defining concrete and abstract visualizations based on the rare existing designs of concrete visualizations. Concrete and abstract visualizations exist on a continuum that would benefit from being neatly defined and would open a wide range of creations, especially for concrete visualizations.

Within the next chapter, I leave representations explorations and get to reflect on more practical issues related to immersive situated visualizations. Indeed, we discuss the real-world scenario in which the physical referent might be temporarily or permanently non-visible.

Chapter 5

Handling real-world constraints for situated visualizations using MR

Chapter summary

In this chapter, we focus on the challenge of providing situated visualizations in realistic conditions. Indeed, while situated visualizations can be beneficial in various contexts and have received research attention, they are typically designed with the assumption that the physical referent is visible. However, in practice, a physical referent may be obscured by another object, such as a wall, or may be outside the user’s visual field. In this chapter, we propose a conceptual framework and a design space to help researchers and user interface designers handle non-visible referents in situated visualizations. We first provide an overview of techniques proposed in the past for dealing with non-visible objects in the areas of 3D user interfaces, 3D visualization, and mixed reality. From this overview, we derive a design space that applies to situated visualizations and employ it to examine various trade-offs, challenges, and opportunities for future research in this area.

Main portions of this chapter are published in the journal of Transactions on Visualization and Computer Graphics under [41]. Thus, any use of “we” in this chapter refers to Arnaud Prouzeau, Martin Hachet, Pierre Dragicevic, and me.

5.1 Introduction and contributions

Consider a search-and-rescue scenario similar to the one envisioned by Willett et al. [322], where a fire team needs to rescue people from a burning house. Imagine the firefighters are equipped with Augmented Reality head-mounted displays that show them useful information about the house (e.g., location of emergency exits), the fire (e.g., temperature and CO₂ level in different rooms), and people (e.g., location of victims and their vitals). All this information would be directly overlaid on the physical objects it refers to (e.g., rooms, victims). The firefighters could use this digital information in combination with perceptual cues informing them about their surroundings (e.g., did parts of the house collapse? Do victims appear distressed or wounded?). A situated information visualization system like this would likely help the fire team reach faster and more reliable decisions than if they had to examine the information on separate displays. Although the scenario we mentioned is futuristic, AR-based situated visualizations, have been demonstrated in various domains. As tackled in chapter 2, situated visualizations are becoming easier and easier to implement, and we will likely see more and more of them. However, designing effective situated visualizations must be informed by a good understanding of their design space. This space is vast and complex, and despite recent advances in exploring and understanding it [81, 256, 296, 322], large areas remain under-discussed and poorly understood. In particular, previous work has almost always considered situated visualizations where the physical referent is in sight. However, in the real world, objects are often hidden or out of sight; thus, we cannot always assume that physical referents are visible. Back to the firefighting team example, we should expect some victims to be invisible to the rescue team because they are in another room, occluded by fragments, flames, or smoke. Yet some of them may need assistance, so knowing that they exist, where they are, and what information is available about them is important for the rescue team. The same applies to elements of the environment, such as emergency exits. However, it is unclear how to visualize information about hidden physical referents like these and what interaction techniques to use. Since most previous situated visualization systems have been designed with visible referents in mind, they cannot provide us with much guidance for such questions. In this chapter, we fill this gap by laying out a conceptual framework consisting of a terminology, a typology of techniques, and a design space for handling non-visible physical referents in situated visualizations. Although visual occlusion has been rarely discussed in the literature on situated visualizations, it is a known problem in the areas of 3D user interfaces and 3D visualization [121], as well as in AR and VR [22]. We survey the techniques proposed in these research areas and classify them into eight families. We then use those families – together with other previously published conceptual frameworks – to build a design space. We ended up with 14 design dimensions broken down into three categories, depending on whether the dimension is relative to the representation of the object of interest, to the visualization, or to the general situated visualization system. We illustrate the use of our conceptual framework in three application areas, discuss different design trade-offs and the limitations of our conceptual framework, and

suggest several directions for future work.

5.2 Related Work

In this section, we discuss related work, mostly on 3D occlusion management.

Taxonomies of occlusion management. As in the physical world, occlusion is a major concern in virtual worlds and in 3D user interfaces and visualizations. As a result, various techniques have been proposed for handling occlusion, either by moving the user’s point of view, or by altering the rendering, position or shape of the occluding objects. Elmqvist and Tzigas [121] proposed a taxonomy of occlusion management techniques in 3D visualizations. They classify techniques in terms of their purpose, the arrangement of the visualization, and the user’s status (active/passive). Examples of techniques include X-ray vision (i.e., making occluding content transparent) and multiple views. The paper discusses a few examples of AR visualizations but they are not situated, and therefore the paper does not discuss how to handle the occlusion of physical referents. More recently, Macedo and Apolinario [22] surveyed AR techniques for handling objects that are physically hidden in real environments. In particular, they reviewed existing X-ray vision techniques that overlay hidden objects in front of occluding surfaces as if the user could see through them. Much of their paper focuses on how to fine-tune the rendering techniques to make the X-ray vision experience as believable as possible. The paper does not discuss data visualization. We borrow from this previous work to discuss how to handle non-visible physical referents in situated visualizations.

Situated visualization frameworks addressing non-visible referents. Closer to our own work, two conceptual frameworks have been recently published that address specific aspects of handling non-visible referents in situated visualizations. Lin et al. [210] reviewed labeling techniques to support visual search for situated visualizations where physical referents are out-of-view. We build on this survey and cover labeling approaches, but we also cover a range of other techniques for addressing non-visible referents, including techniques that convey information about referents in a way that does not require people to find them. Another closely related contribution is the work on ProxSituating Visualizations [267], which focuses on using proxies for physical referents – i.e., virtual representations replicating physical referents. The paper extends Willett et al.’s [322] model to account for proxies and observes that a visualization can be considered situated if it is close to a proxy, even if the physical referent is far away – a situation we will often encounter in this chapter. Again, this chapter focuses on specific techniques for handling non-visible referents. We do not build on this work as this is parallel work, but we will contrast their design space with ours in detail in subsection 5.6.2. Despite these two prior conceptual frameworks, to our knowledge, no research has offered a broad and comprehensive survey of different ways non-visible physical referents can be handled in situated visualizations. This chapter fills this gap

by taking inspiration from the many occlusion management techniques proposed for VR, AR, and 3D user interfaces and generalizing them to situated visualizations to build a comprehensive conceptual framework.

5.3 Terminology

5.3.1 Illustrative Scenarios

We now describe two short scenarios we will use throughout this chapter in order to illustrate our conceptual framework. The first scenario elaborates on the fire rescue example we mentioned in the introduction, while the second scenario covers a more casual situation.

Scenario 1. Alice is part of a firefighter team. Just before dinner time, the team receives a fire alert involving a nearby house. They know from the caller that the household consists of two parents, two children, and an elder and that one parent and the firstborn are out of the house and on their way to the incident. Soon, the fire truck speeds on the road, and inside, the team prepares for the operation: they check the house layout on a 2D map, focusing especially on fire exits. They call the safe parent to gather more information and then assign each other tasks. Alice is in charge of rescuing the younger child, and she identified on the map that his bedroom is on the second floor. The parent confirmed that the child is sleeping in his bedroom because the room is equipped with a video baby phone connected to his smartphone. Some of Alice’s teammates have been assigned to rescue the other family members, while others are in charge of extinguishing the fire. Two members are in charge of staying near the truck to coordinate the whole operation, give instructions and monitor information such as the amount of water left and the vitals of the team members captured in real-time by wearable physiological sensors. Once the team arrives, the rescue operation goes relatively smoothly despite a few unexpected events. For safety reasons, the rescue team first enters the basement and cuts the electricity. From there, they use stairs to reach the rest of the house. They manage to save everyone but not without difficulty due to the obscurity, the presence of heavy smoke, and the important number of rooms making it hard to see anything.

Scenario 2. Jessie gets home where her husband Max has been waiting for her. These days, they are trying to change their cooking and shopping habits to waste less. They took the habit of regularly logging their groceries and the meals they prepare in a phone app. Tonight, Max is thinking about what he will prepare for dinner. He opens the fridge and sees that the bottom drawer is full of vegetables and fruits and that the leeks must be used soon. Max opens the cooking book and chooses a leek pie. He leaves the book open on the counter and looks at the ingredients listed with small illustrations and the cooking instructions. He goes back and forth between the kitchen worktop and the counter to read the different steps. He realizes two items are missing: one they thought they had but forgot to buy and another one




Figure 5.1 Diagram summarizing the key terminology introduced in Section 2.3. The pictures refer to Scenario 1, where a firefighter needs to evaluate the state of the victim hidden behind a wall.

that has expired. However, they are not essential, and Max decides to continue. He puts the pie in the oven and starts to wash the dishes, and after a while, he smells a slight burnt odor. He turns around and sees it is high time to take the leek pie out of the oven. Fortunately, the pie turned out pretty good at the end.

These two scenarios are very different, but they both involve situations where people are engaged in tasks that require access to information about objects in the physical environment, and these objects may be hidden from sight. This information may be something that can be directly seen as long as the object is visible (e.g., whether a vegetable is ripe), or something that needs to be conveyed by a visualization (e.g., the firefighters’ vitals) In this chapter, we are interested in cases where data is visualized and where seeing objects in the environment may provide complementary information [322].

5.3.2 Basic Terminology

In this section, we introduce the basic terminology that underlies our conceptual framework and which will be useful for describing the design space. Figure 5.1 gives a visual overview of our terminology.

An **object of interest** is an entity about which the user wants to learn. For example, in Scenario 1, a person to rescue, an emergency exit, or another firefighter can all be objects of interest for a firefighter. In Scenario 2, an object of interest can be a grocery item necessary for the recipe. Whether an object is of interest depends on the user’s task and intent, which can change over time. However, we will be using the term “object of interest” liberally to refer to objects that are currently of interest but also to objects that are expected to become of interest at some point in time. Note that an object of interest can be a physical referent, but it is a more general term as it does not indicate that any data visualization is attached to the object.

A **representation** of an object of interest is a visual depiction that stands for this object. It can reproduce the object’s visual appearance or not. For example, in Scenario 1, the image of the child that appears on the parent’s smartphone is a relatively detailed and faithful representation of the actual child. In contrast, elements depicted on the house map, like emergency exit icons, are purely symbolic representations. The drawings of the ingredients in the cooking book from Scenario 2 are an intermediary example: while the drawings approximate the appearance of the actual ingredients, they are generic drawings that do not show, for example, whether the actual leeks in the fridge are overripe. The design space section will further discuss this notion of representation faithfulness.

Data is information about an object of interest that is not necessarily conveyed by its visual appearance. For example, in our Scenario 1, information like the composition of the household, the firefighters’ vitals, or the amount of water left exists but cannot be accessed through mere visual inspection. In Scenario 2, examples of data include the ingredients’ expiry date, the list and quantity of ingredients in a recipe, or the history of Jessie and Max’s meals. The location of an object of interest is also data. Data can be conveyed through data **visualizations**. Following previous terminology [322], whenever data is visualized about an object of interest, the data’s **physical referent** is this object of interest. In addition, when the visualization is near (or appears to be near) the object of interest, it is **situated**.

We refer to an object of interest as **non-visible** if it is visually imperceptible by the user at a given time. For example, in Scenario 1, while Alice is in the basement, the child located on the second floor is non-visible. In Scenario 2, most cooking ingredients are non-visible unless the fridge and kitchen cupboards are opened. We identify three main reasons why an object of interest may be non-visible:

1. Following previous terminology [210], an object is **out-of-view** if it is not within the user’s visual field but can be seen if the observer turns their head. For example, in Scenario 2, the pie in the oven is out of view while Max is washing the dishes.

2. An object is **occluded** if its sight is blocked by another object. For example, in Scenario 1, the child is occluded by walls. We more specifically refer to an object as **hidden** if it is in the user’s vicinity but occluded from all reachable locations and viewpoints. For example, in Scenario 2, as long as the

fridge’s door is closed, there is no viewpoint from which Max can see the food it contains.

3. An object is **illegible** if none of its useful details is visible to the user for reasons other than being out-of-view or occluded. For example, the object can be too small, too far, moving too fast, or insufficiently illuminated. In Scenario 2, the cooking book is too far for Max to be able to read, forcing him to move back and forth between the book and the kitchen worktop.

5.4 Handling Non-Visible Objects in Immersive Environments: A Survey

Although there has been little work on handling non-visible referents in situated visualizations, there has been a substantial body of research in VR and AR on techniques for handling non-visible objects of interest. In this section, we present a survey of such techniques and classify them into eight categories. This survey will inform the design space we lay out in the next section.

5.4.1 Methodology and Scope

In order to collect possibly relevant papers, we conducted multiple keyword searches on several scientific search engines: Google Scholar, Scopus, ACM DL, and IEEE Xplore. Our search terms consisted of occlusion, off-view, distant, hidden, and awareness combined with either augmented reality, mixed reality, or immersive. In addition, we added papers cited in the 2008 taxonomy of 3D occlusion management techniques from Elmqvist and Tsigas [121].

In order to keep the size and complexity of the survey and the resulting design space within reasonable limits, we iterated over a set of criteria to restrict the scope of our survey. We converged towards three major criteria. For a paper to be included, the technique it presents had to meet the following requirements:

1. **Visual only.** The technique needs to use visual representations. Techniques that increase the user’s awareness about non-visible objects of interest through non-visual channels (e.g., haptic or auditory) were not included in this survey.

2. **Object-aware.** The technique requires that an object of interest is defined. Object-agnostic techniques- e.g., techniques that broaden the field of view or magnify vision- were not included. This choice is consistent with our focus on situated visualizations, where the notion of object of interest is central (objects of interest are physical referents, i.e., objects about which data is visualized).

3. No movement required. The technique should be usable from any location in space. Techniques that do not provide information about the object of interest but instead help the user access it by locomotion or body movements were excluded.

We ended up with a total of 107 papers, which we divided into eight families of techniques. We give an overview of the eight families here. The full list of papers is available as supplementary material section A.4.

5.4.2 Worlds-In-Miniature




Figure 5.2 A world-in-miniature showing a scaled-down version of the physical room where the user is located, overlaid in their visual field using augmented reality [63]

A World-In-Miniature (or WIM) is a virtual scaled-down replica of the environment the user is in, embedded within the users' view of their immediate surroundings. The scaled-down environment is often virtual. Figure 5.2 shows an example of a WIM of a real environment (the laboratory room the user stands in) displayed in AR. A WIM aims to help users discover and learn about the environment around them, including parts that may be occluded, out-of-view, or illegible.

WIMs have been used in various domains, including for medical imaging [102], building exploration for tourist guides or situated urban planning [294], air traffic visualization [118], US Forest Service's forest inventory data visualization [235], and interacting with smart home and IoT appliances [29, 274]. In both AR and VR, WIMs have been mainly explored for supporting navigation [247], but they can also be used to get information about hidden objects of interest.

Depending on the implementation, WIMs are more or less faithful depictions of the full-scale environment. In VR, WIMs can be exact copies. In AR, WIMs are typically simplified versions of the real scene (e.g., Figure 5.2), but some WIMs are photorealistic copies [75](here, a minimap).

For a WIM to be usable, it is important that it can be freely manipulated. Although the first implementations only allowed users to move their head closer or further away from the WIM [62], further implementations allowed them to grab the WIM and rotate it [34] and even re-scale it [109]. Direct manipulation of elements inside the WIM has also been explored – for example, one application allows users to move furniture in a virtual room by moving their replica inside the WIM [73]. In another application showing building consumption on a university campus, users can get additional info about buildings by grabbing them within the WIM [126]. Some variations of WIMs even address occlusion problems with the WIM itself, for example, by allowing users to make walls transparent [302] or to select non-visible regions of interest [34].

Whenever she needs, Alice can visualize a 3D replica of the endangered house. She can localize herself, the people to rescue, and the safety exits. She can get additional information with or without input (eye glance, virtual touch...): teammates' and victims' vitals, any room CO2 level and temperature, water supplies, and localized density of fire. She can see a trace of her route in the representation and better understand it through the house layout. She can manipulate and rotate it in the free space if occlusion blocks any information.

5.4.3 Virtual Twins




Figure 5.3 A virtual twin of a smart factory shown in augmented reality, allowing the user to see the state the factory in real-time and remote-control it [199].

A virtual twin¹ is a dynamic 3D model that conveys the visual appearance and key properties (e.g., state, features, physical measurements) of a physical system or object. Figure 5.3 shows an example of the virtual twin of an illegible factory.

¹There are many definitions of virtual and digital twins in the literature. Here, we consider that a virtual twin is a digital twin with a 3D visual representation.

Conceptually, a virtual twin is similar to a WIM, except it replicates an object instead of a spatial environment and focuses on replicating its key properties. Although a virtual twin can faithfully convey the visual appearance of the object of interest [293], often the focus is on conveying its important properties and associated data. Virtual twins have been suggested for a range of applications, including for surgery [200], radiology [283], assembly tutoring [258] and for sharing content during teleconferencing [188]. They can help support maintenance [199], learning [258], and collaboration [188] tasks.

Alice can see a symbolic replica of each person to rescue in her field of view that transmits information about their global state. She can interact with the virtual replica to get different levels of information about the victims' vitals. She also can see a metaphoric representation of the water supplies that indicates the amount left.

5.4.4 Physical Replica




Figure 5.4 Tangible interface for visualizing the energy consumption of a university [126].

A physical replica is a physical representation of an object of interest. As an example, Figure 5.4 shows a tangible model of a campus augmented with energy consumption data. The model is installed in a room within the campus itself, and thus, it allows users to get an overview of the whole campus, which is out of view for the most part. Physical replicas are conceptually similar to WIMs and to virtual twins, except they are physical. A benefit of physical replicas is that they can be designed to be easy to grab and manipulate, as in the Uplift system [126] (Figure 5.4). However, accurate and dynamic physical replicas are difficult to build. As a consequence, the purpose of most physical replicas that have been developed

so far is not to convey detailed and faithful information about the object of interest but rather to serve as support surfaces on which to display digital information about the object or as tangible controllers for navigating this information. Examples include the Tangible Globe [269] which users can use as an input device to browse information about the Earth, the head probe from Hinckley et al. [168] which users can use to control a virtual cutting plane on a separate screen, or the humanoid puppet Teegi which has a display integrated showing EEG data [137].

Alice is in the firetruck waiting for the right time to come as support with water supply. Inside, they have a physical replica of the house in which they take action. This replica is augmented; she can see how her teammates are evolving in the house, where the victims are, and the safety exit. She has additional information about the fire progression and the vitals of the endangered people.

5.4.5 Telepresence Systems




Figure 5.5 Example of a telexistence/telepresence system. User A on the left sees a virtual representation of the persons B and C they want to interact with. The user can also remotely control a robot as if they were physically in the same location [291].

Telepresence or telexistence broadly refers to techniques that give the user the impression of being in a different physical location, typically by providing an immersive view of the remote environment or object and/or by allowing users to directly interact with it from a first-person perspective [276, 291]. Figure 5.5 shows a system allowing users to see and interact with remote people as if the user was in the same room. Telepresence systems sometimes let users physically interact with remote objects of interest using

physical actuators or robots – for example, physically operating a remote cockpit through a robot [290]. In many systems focusing on social telepresence, the objects of interest are other people, who are rendered virtually as if they were near the end user [276]. Often, these representations do not resemble the people themselves, but research has started to explore automatic 3D reconstruction [141, 216]. Alternatively, the representation of remote people can be physical, like in the case of mutual telexistence systems where robots represent remote users [198]. Overall, telepresence covers situations of teleconferencing [216], visualization sharing at a distance [312], tutoring [297]. School and Healthcare have also been identified as promising application areas [198].

Alice is in the fire truck, ready to get in the house to deploy the water supplies. She has a small window on the top right-hand corner that gives her what her teammates have in view, and she can switch from one to another. She can see that one of her teammates is carrying a child and sees overlaying vitals about the child and their location. She sees he is struggling with a dense fire shelter, so she knows where to head.

5.4.6 X-Ray Vision



Figure 5.6 X-ray vision for surgery, projection of anatomical structure on the patient's body [309].

X-ray vision techniques emulate the fictional superhuman ability to see through occluding surfaces [230, 321]. Figure 5.6 presents an overlay of the patient's hidden anatomical structure on their body.

Applications of x-ray vision range from visualizing the heat distribution of buildings [49] to assisting underground exploration [129].

A range of rendering techniques has been proposed to emulate x-ray vision in augmented reality (for reviews, see [214, 230]). In x-ray vision systems, objects of interest may be reconstructed based on

3D models [121, 187] or on video images, for example, as captured by a drone [128]. Although most implementations of x-ray vision do not focus on interaction aspects, two systems have been proposed that use a torch metaphor to specify the region to reveal: one allows users to explore virtual annotations left in different rooms of a workspace [83], while the other one allows users to control home appliances such as lights and TVs even when they are in other rooms [144]. Explosion diagrams are closely related to x-ray vision, but they displace occluding surfaces instead of removing them or rendering them transparent [121, 187].

Alice is searching for a specific person she knows is on the other side of the house; while moving forward the person, Alice can see through the wall and understand the layout and obstacles that separate her from the victim. She can see her teammates acting through the different layers, and the safety exits are highlighted. She gets additional information, such as her teammates and victims' vital, directly overlaying their representation. She can better understand her global route through the house layout.

5.4.7 Out-Of-View Pointing




Figure 5.7 Out-of-view pointing in SidebARs [278]. On the top right, arrows point toward locations of fire companies.

Out-of-view pointing techniques give information about the presence and location of an object of interest that is outside the user's field of view, typically by indicating the object's direction on the edge of the display. Figure 5.7 shows an example in AR.

Out-of-view pointing techniques were originally developed for regular computer displays [57, 76, 161], and later adapted to VR [95, 154] and AR [149, 157] displays. Some techniques augment the display hardware, e.g., with LEDs [215].

In most cases, no information is given about the object of interest apart from its existence and approximate location. Most techniques only give direction clues [154, 157] while others also convey the distance of the object of interest [186]. A few applications show data visualizations about the object of interest for notifications [178], for fire companies [278] (figure 5.7) or to discover a city [5]. But in general, the amount of information shown with this technique is limited.

Alice, while searching for the people to rescue directly in her field-of-view, sees arrows that point toward her different teammates and victims relative to her orientation. At the bottom right-hand corner she can see a schematic overall representation of the area with dots, one for her and others for the teammates and victims to get a global idea about the victims' location relatively to her and to each other. She can interact with the arrows and have further details about their vitals, current tasks for the firefighters, and distress state. It is the same thing for resource indications such as safety exits and water supplies. She can manipulate a cursor to choose the level of details she wants.

5.4.8 Labeling




Figure 5.8 The stereo panorama of the Luxor Temple was provided by Tom DeFanti, Greg Wickham, and Adel Saad, with stereo rendering by Dick Ainsworth for AR labeling [261].

Labeling in AR applications consists of displaying additional textual information about objects of interest, which can be visible or not [210]. Figure 5.8 shows an AR application that displays information boxes linked to the physical objects they belong to.

Labels are often displayed near the objects they refer to, and sometimes visual links are drawn between the labels and the objects to highlight their association [28]. In order to position the labels at the right place, a specific layout tracking or processing system is often used [148] along with an adequate view management system [62] [148]. Recent work has introduced the use of labels in the specific case of out-of-view objects and highlighted their benefits [210]. In terms of interaction, the user generally does not have explicit control over the position of the labels. However, to make sure the labels are visible and readable,

their position is often automatically adjusted to the position and orientation of the users' heads [210]. Although labels can sometimes provide rich information, as in 5.8, in most cases, labels only show very basic information about objects of interest. This makes them similar to out-of-view pointing techniques.

Alice, while searching for the people to rescue, can see lines coming from the cluttered areas that end up in tags with written information about the item it refers to a family member to rescue and their vitals, a safety exit and the different obstacles between them, a teammate and where he is heading at.

5.4.9 Virtual Mirrors




Figure 5.9 Example of a virtual mirror used to see the hidden side of a spine [66].

A virtual mirror is a virtual tool that behaves like a physical mirror, i.e., it reflects nearby objects. Figure 5.9 shows an example.

Users can generally manipulate a virtual mirror as they would manipulate any real mirror [67]. The mirror's surface shows the simulated reflected image that allows the user to see non-visible content, such as hidden faces of objects of interest or objects outside one's field of view. A variation of the virtual mirror are physical mirrors whose image is digitally generated and therefore, virtual. For example, in the Mind-Mirror project [222], users are equipped with an AR display and a real mirror and can see their own electroencephalography data through the real mirror. One of the major benefits of mirror-based techniques is that the mirror metaphor is very familiar to most users.

While getting to the person she has been assigned to rescue, Alice has a virtual mirror at hand. She enters a room, spots the victim in a corner and a door with flames at the other side of the room. She runs to the victim, and, with a quick hand gesture, puts a virtual mirror to monitor the door from her point of view while performing first aid actions.

CHAPTER 5. HANDLING REAL-WORLD CONSTRAINTS FOR SITUATED
VISUALIZATIONS USING MR

Table 5.1 Summary of all 14 design space dimensions.

Category	Dimension	Description	Values		
General	Non-visibility support	The types of non-visible physical referents that are supported.	Out-of-view	Occluded	Illegible
	Interactivity	The extent to which the user can interact with the representation or the visualization.	Low	Medium	High
	Interaction realism	The extent to which the interaction techniques are reality-based.	Low	Medium	High
	Actionability	The extent to which the user can act on the non-visible physical referent.	Low	Medium	High
Representation	Representation fidelity	The level of detail with which the physical referent is represented.	Low	Medium	High
	Representation refresh rate	How often the representation of the referent is updated to reflect the real referent.	Low	Medium	High
	Surroundings inclusion	How much of the environment around the referent is included in the representation.	Low	Medium	High
	Representation medium	Whether the representation of the physical referent is digital or physical.	Digital	Mixed	Physical
Visualization	Representation integration	To what extent the representation appears to be integrated in the user's environment.	Low	Medium	High
	Location information	The amount of information shown about the physical referent's location.	Low	Medium	High
	Visualization richness	The amount of data visualized about the physical referent.	Low	Medium	High
	Visualization refresh rate	How often the visualizations are updated to reflect data about the physical referent.	Low	Medium	High
	Visualization medium	Whether the visualizations are digital or physical.	Digital	Mixed	Physical
	Visualization integration	To what extent the visualizations appear to be integrated in the user's environment.	Low	Medium	High

5.5 Design Space

To lay out possible techniques that can handle non-visible referents in situated visualizations, we describe 14 design dimensions broken down into three categories, depending on whether they are relative to the representation, to the visualization, or to the general situated visualization system (see Table 5.1 for an overview).

We will reuse the two scenarios introduced in chapter 5.3.1 to illustrate the different values the dimensions can take. In contrast to the initial scenarios, we will assume that Alice and Bob are now equipped with lightweight AR head-up displays. We will ignore technological limitations and implementation issues – for example, if Alice uses X-ray vision to see behind a wall, we will not specify how the hidden environment is visually reconstructed: it could be from an array of surveillance cameras, from drones equipped with depth sensors, etc.

5.5.1 Category: General

The dimensions in this category capture general characteristics of the situated visualization system that do not specifically apply to the representation of the physical referent or to its visualization.

Non-visibility support: Refers to the types of non-visible physical referents that are supported by the situated visualization system (see subsection 5.3.2 for our terminology on the subject). In contrast to the other dimensions, the values here are not mutually exclusive.

- *Out-of-view:* The technique allows the user to see information about physical referents that are outside their field of view. Example: In scenario 1, Alice stands in front of the bedroom door where the child is stranded, thinking about what to do next. Meanwhile, she talks to her partner behind her. Even though she is not facing him, a virtual rearview mirror allows her to see him.
- *Occluded:* The technique allows the user to see information about occluded referents. Example: Alice can see the child and information about her through the bedroom door using an X-ray vision display.



Figure 5.10 Illustration of the non-visibility support dimension. The red cube represents the object of interest and the green one the visualization and/or the representation.

- *Illegible*: The technique allows the user to see information about illegible referents. Example: Alice zooms into the child’s face to see if she is still conscious.

Interactivity: The extent to which the situated visualization system allows the user to interact with the representation of the physical referent or the data visualization (see again Figure 5.1 for our terminology). A technique like World-In-Miniature is typically highly interactive because users can rotate and rescale the representations. However, the same technique can come with different levels of interactivity.




Figure 5.11 Illustration of the interactivity dimension. The green cube represents the visualization and/or the representation. The low value is represented by a user who does not manipulates the visualization and/or representation. The medium value shows a user who interacts in a limited way with the cube. High interactivity is represented by an individual manipulating the representation/visualization as if it was a physical object.

- *Low*: The representation and the visualization are non-interactive. Example: In scenario 1, Alice searches for the stranded child. She sees a video of the child with basic information about her vitals and an arrow pointing to her direction, but there is no control to change the camera settings or get further details about the child’s state.
- *Medium*: The user can interact with the representation or the visualization to a limited degree. Example: Alice can turn the camera image and the visualization on or off.
- *High*: The representation and the visualization are highly interactive. Example: Alice can finely tune the camera image and the arrow representation and explore many different types of data dimensions

and visualizations for the vitals.

Interaction realism: The extent to which the interaction techniques are reality-based, i.e., “increase the realism of interface objects and allow users to interact even more directly with them—using actions that correspond to daily practices within the non-digital world” [181]. In the literature on immersive analytics, such interactions are also called embodied [106]. Examples of reality-based interactions are virtual mirrors, as they are used very similarly to real mirrors.

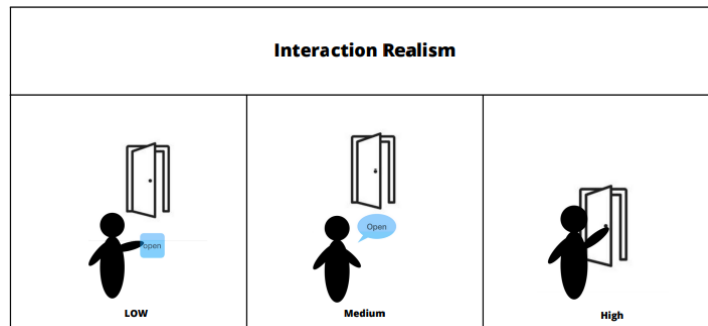


Figure 5.12 Illustration of the interaction realism dimension. The low value is represented by a user who pushes a button to open the door. The medium value shows a user who uses a voice command to open the door. High interaction realism is represented by our custom to open a door, with a handle

- *Low:* The interactions are not reality-based. Example: Alice is looking for the child’s bedroom. She opens a virtual menu and selects an option that displays a 3D arrow pointing at the bedroom. Once in front of the bedroom door, she uses another menu to activate an X-ray display that reveals the room’s interior. She then activates visualizations of the room’s temperature and air quality and of the child’s vitals.
- *Medium:* The interactions are weakly reality-based or combine reality-based elements and others that are not.
- *High:* The interactions are reality-based. Example: Alice looks around, which activates the 3D arrow. Once she sees the door, she grabs a physical controller and uses it like a torch to activate the X-ray display (as in [144]). The closer she moves to the door, the more detailed information she gets about the room and the child.

Actionability: The extent to which the user can act on the non-visible physical referent. For example, IoT and smart home applications can give users control over lights or home appliances (e.g., [144]), thus supporting actionability. In contrast, techniques such as out-of-view pointing and labeling rarely support actionability.

- *Low:* The situated visualization system does not let the user act on the physical referent. Example: in scenario 2, while Bob washes the dishes, his heads-up display shows a heatmap visualization of the pie’s temperature. The visualization indicates that the pie is about to burn, but Bob has to walk to




Figure 5.13 Illustration of the actionability dimension. The red cube represents the object of interest. The low value is represented by a user who doesn't act on the red cube, whose light seems to be turned off. The medium value shows a user who lights the cube with a switch. High actionability is represented by the user commanding the cube with a joystick.

the oven to turn it off.

- *Medium*: The technique supports basic actions on the physical referent, such as flipping electronic states on and off, and other actions typically supported by remote controllers.

Example: Bob can turn the oven off without having to walk to the oven.

- *High*: The technique allows rich manipulation of the physical referent, which includes mechanical actions. Maximum actionability is achieved when the user can manipulate the object as if they were standing in front of it. Example: Bob can open the oven door and pull the oven rack out from a distance.

5.5.2 Category: Representation

The dimensions here capture characteristics of the representation of the physical referent.

Representation fidelity: Refers to the level of detail with which the physical referent is represented. It can go from a simple symbol to a highly detailed model of the physical referent.




Figure 5.14 Illustration of the representation fidelity dimension. The red cube represents the object of interest. A green square represents the low value. The medium value shows a green cube. High representation fidelity is the actual visual appearance of the object of interest.

- *Low*: The representation of the physical referent is either non-existent or does not convey any information on the visual appearance of the physical referent. Example: Alice sees a “plus” symbol indicating the child’s location behind the wall.
- *Medium*: The representation conveys some elements of the visual appearance of the physical referent. Example: Alice sees the icon of a girl standing for the child.
- *High*: The representation is detailed and reasonably matches the physical referent’s visual appearance. Examples: Alice sees a photo or a 3D model of the child.

Representation refresh rate: How often the representation of the physical referent is updated to reflect the real referent.




Figure 5.15 Illustration of the refresh rate dimension. The red cube represents the object of interest, and the green one the representation/visualization. A picture of the cube on the wall represents the low value. A live broadcast represents a high refresh rate.

- *Low*: The representation of the physical referent is generated once and never updated. Example: in Scenario 2, Bob asks the situated visualization system to show him the fruits and vegetables he bought a week ago at the market. The system shows him photos of the produce taken before they were put in the fridge, which tells about their number, type and size, but not whether they are about to go bad.
- *Medium*: The representation of the physical referent is occasionally updated. Example: Bob asks the situated visualization system to show him the fridge interior. The system takes a photo every morning and shows him the last one. Bob can see the fruits and vegetables and which of them should be eaten soon. However, he does not see that Alice ate all the remaining grapes just an hour ago.
- *High*: The representation of the referent is updated in real-time. Example: Bob sees a real-time video of the interior of the fridge.

Surroundings inclusion: How much of the physical environment surrounding the referent is included in the representation.

- *Low*: The surroundings of the physical referent are not represented. Example: Thanks to the x-ray feature, Alice sees the child through the bedroom door, but only the child is represented. She cannot see whether there are hazardous elements around the child.
- *Medium*: Several elements of the physical referent’s surroundings are included in the representation.



Figure 5.16 Illustration of the surroundings inclusion dimension. The green cube stands for the representation and/or visualization. The low value shows a cube without any surroundings depiction. The medium value shows a limited surrounding of the cube. High surrounding representation shows the entire room the cube is in.

Example: The X-ray feature is able to capture and render the entire bedroom, but Alice can only see the portion that is visible through the virtual hole in the door.

- *High*: The representation includes many of the referent’s surroundings. Example: Alice switches from the x-ray representation to a world-in-miniature replica of the bedroom, where she can see the child as well as all the details of her physical environment.

Representation medium: Whether the representation of the physical referent is digital or physical. It can go from completely digital to completely physical or can be anywhere in-between [184, 224].




Figure 5.17 Illustration of the medium of the representation/visualization dimension. The physical value shows a real pot of grass. The mixed value is represented by a physical pot of grass augmented by computer-generated leaves. The virtual value shows a completely computer-generated plant.

- *Digital*: The representation is a computer-generated immaterial image. Example: While the truck is speeding to the house on fire, the firefighting team downloads a floor plan of the house from a central database and video-projects it on the truck’s floor.
- *Mixed*: The representation combines physical with digital elements. Example: The firefighting team uploads a floor plan of the house on a shape display, i.e., a device with motorized pins [184]. The shape displays morphs into a coarse physical replica of the house’s floors. Other details (doors, fire exits, etc.) are video-projected.
- *Physical*: The representation is purely physical. Example: a high-resolution shape display whose pins can change color morph into a detailed physical replica of the house’s floors.

Representation integration: The extent to which the representation appears to be integrated into the user’s physical environment. If the representation is physical (see above), its integration with the physical world is necessarily high. If it is digital, it depends on the visual metaphors and rendering techniques used to display it. One key factor is whether the representation is egocentric (it follows the user) or exocentric (it is attached to the physical referent or its environment).




Figure 5.18 Illustration of the integration of the representation/visualization dimension. The green cube stands for the representation and/or visualization. The low value shows a cube floating in space. The medium value shows a cube geometrically well integrated in space but that does not occlude elements behind it. High integration shows a cube well integrated in space, as a physical cube would have been.

- *Low:* The representation of the referent does not appear integrated in the physical environment at all, and it is very clear that this representation is a virtual image. Example: Alice sees a 3D model of the child floating in the air.
- *Medium:* There is some sense that the representation is integrated into the physical environment, but this integration is not fully believable. Example: Alice sees the child through the door, but it is clear that the representation has been added because of imperfections in the AR rendering and because it is impossible to see through physical objects in reality.
- *High:* The representation appears to be an integral part of the physical environment. Example: The situated visualization system renders a perfect full-size replica of the child’s bedroom in front of Alice, as if the door and the walls have been completely removed.

5.5.3 Category: Visualization

The dimensions in this category capture the characteristics of data visualizations about the physical referent.

Location information: The extent to which the situated visualization system visualizes information to help the user know where the referent is located. This dimension can be low or high depending on how precise and complete the information is.

- *Low:* The situated visualization system gives no information about the location of the physical referent.




Figure 5.19 Illustration of the location information dimension. The green cube stands for the representation. The low value shows a cube without any location information. The medium value shows a cube with an indication that the cube is in Montreal. High location information is represented by the exact geographic coordinates the cube is located at.

Example: Bob needs to assemble the ingredients listed in the cooking book. The system shows him his current ingredients but not where to find them.

- *Medium*: The system gives hints about the location of the physical referent, for example, its distance or its direction. Example: The system adds an arrow next to each ingredient listed in the cooking book, that tells Bob in which direction to go to find the ingredient.
- *High*: The system gives enough information to be able to find the physical referent quickly. Example: a combination of arrows, labels, and x-ray representations tell Bob exactly where to find each ingredient.

Visualization richness: The amount of information shown about the physical referent. It can range from low (e.g., as in simple labeling systems) to high (e.g., as in many virtual twins). This dimension is analogous to representation fidelity, except representation fidelity is about the visual appearance of the physical referent, whereas visualization richness is about its associated data. In addition, there is an upper limit to representation fidelity but no such limit for visualization richness.

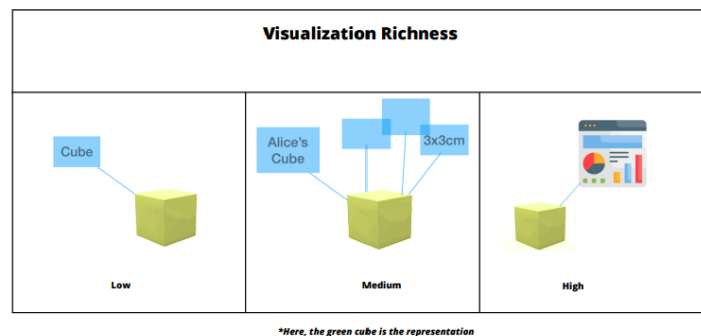


Figure 5.20 Illustration of the visualization richness dimension. The green cube stands for the representation. The low value shows a cube labeled “cube”. The medium value shows a cube with multiple labels. High information richness shows a cube labeled with advanced analytics.

- *Low*: The situated visualization system shows minimal information, such as the physical referent’s name or location (see before). Example: Alice sees a label next to the child behind the door, indicating her

name and age.

- *Medium*: The system visualizes some data about the physical referent. Example: next to the child, the system displays her body temperature, breathing rate, and heart rate.
- *High*: The system visualizes lots of data about the physical referent, allowing the user to perform analytical tasks. Example: Alice sees detailed health data about the child, allowing Alice to estimate how much time is left before the child starts suffocating.

Visualization refresh rate: How often is the visualization updated to reflect data about the physical referent. Refresh rate can be low either because the visualization is not updated frequently enough or because the data is not updated or transmitted frequently enough. This dimension is analogous to the representation refresh rate but applied to the referent's data. See Figure 5.15.

- *Low*: The data is collected, visualized once, and never updated. Example: Alice sees a visualization of the planned trajectories of all her firefighting partners but does not see their current location or any new information collected since they entered the house.
- *Medium*: The data and its visualizations are occasionally updated.
- *High*: The data and its visualizations are updated in real-time. Example: Alice sees the real-time location of her partners, which she can compare with the trajectories to see who deviates from the plan.

Visualization medium: Whether the visualizations are digital or physical. This dimension is analogous to the representation medium but applied to the visualizations. See Figure 5.17.

- *Digital*: The visualizations are computer-generated images that are immaterial. Example: In their truck, the firefighting team video-projects a floor plan of the house on the truck's floor, together with visualizations of temperature and air quality data.
- *Mixed*: The visualizations combine physical with digital elements. Example: A shape display morphs into a physical replica of the house floors, and temperature and air quality data are video-projected on top of this replica.
- *Physical*: The visualizations are purely physical. Example: on the shape display, the walls and doors on the scale model change their texture and color to reflect temperature and air quality data.

Visualization integration: The extent to which the visualizations appear to be integrated into the user's physical environment. This dimension is analogous to the representation integration but applied to the visualizations. It is generally more difficult for visualizations to be perceived as highly integrated because visualizations tend to use abstract encodings that do not resemble objects in the real world. See Figure 5.20.

- *Low*: The visualizations do not appear integrated into the physical environment at all. Example: While facing the bedroom door, Alice sees data about the air quality inside, as 2D charts floating in the air.
- *Medium*: There is some sense that the visualizations are integrated in the physical environment, but

this integration is not fully believable. Example: With her x-ray vision, Alice sees situated volumetric visualizations of air quality data through the door, but it is clear that those visualizations are not part of the physical world.

- *High*: The visualizations appear to be part of the physical environment. Example: the situated visualization system replicates the interior of the bedroom perfectly and at full scale, giving the illusion that walls and doors have been removed. Smoke of different colors can be seen in different areas of the bedroom to convey high concentrations of invisible but harmful gases and particles.

5.6 Discussion

In this section, we illustrate the use of our design space in different application areas, discuss trade-offs and limitations, and propose several directions for future work.

5.6.1 Use Cases and Design Trade-Offs

To illustrate the different concepts from our framework (terms highlighted in bold), we discuss three examples of application areas for situated visualizations. For each of them, we cite an existing research prototype from the visualization literature and discuss different ways in which this prototype could handle non-visible referents.

1. Building Management.

Some buildings like factories and power plants contain many sensors producing large quantities of **data** [256]. Technicians need to understand this data to be able to deal with failures or optimize the building's behavior (e.g., its energy consumption). It can be useful to have AR systems that can position data **visualizations** directly near their **physical referents**, but often the referents are **non-visible**. Consider an air conditioning system: **objects of interest** may include air handling units, air return fans, dampers, pipes, and vents. Most of these are non-visible due to **occlusion**: they are either behind walls, in the ceiling, or in the floor. Suppose we want to 1) help technicians diagnose air conditioning problems from anywhere in the building 2) help them locate and fix the problem once they are in the room. For task 1), the AR system could use a virtual twin, as in the Corsican Twin [256], which uses virtual twins for building management. In such a case, it is important to provide as much useful information as possible about objects of interest, which could imply supporting a high **representation fidelity**, a high **visualization richness**, and **surroundings inclusion**. Now, if we turn to task 2), surroundings inclusion is much less important since the technicians can now see the physical environment. However, choosing a technique with high **integration** like x-ray vision is desirable, as technicians may need to test and fix broken equipment quickly. While equipment repair often requires physical presence, other building

maintenance tasks can be done through the remote control of electrical and mechanical equipment. In these cases, integrating remote control features directly in the situated visualization system itself can help support **actionability**. Furthermore, the system could let users control remote objects of interest directly through their representation as if they were actually manipulating the objects themselves, and thus support high **interaction realism**.

2. Studying Plants. Situated visualizations can help study plants: for example, Virtual Vouchers [316], already mentioned in chapter 2, is a mobile AR system that helps botanists recognize and study plant species using a large dataset consisting of photographs of dead leaves called vouchers. Commercial mobile apps (e.g., <https://plantnet.org>) can automatically recognize plants and display information about them. In all such cases, the **objects of interest** are plants, and AR can be used to display **data** about those plants through **visualizations** next to the plants themselves. Examples of data include geo-temporal records of the species [316], or the place and time the same species were identified. While current apps assume users are always facing the plant they want to learn about, there are possible scenarios where the plants of interest cannot be seen. For example, a botanist, a forest conservationist, or a nature enthusiast may want to go to a location known for a rare plant species or go on a hike and wish to learn about interesting plants nearby. In such cases, a situated visualization may use a geographical database of plant species distributions [77] or individual specimens [72] and show data about plants that are in proximity but cannot be seen. The system could additionally help users orient themselves using techniques with high **representation integration** and low **surroundings inclusion** (e.g., off-view pointing techniques or labels) or techniques making the opposite trade-off (e.g., worlds-in-miniature). For plants that do not change rapidly, using a high **representation refresh rate** or **visualization refresh rate** is not necessary. Having a high **representation fidelity** may be useful, however, but is very difficult to do technically unless the geographical database of plant species is regularly updated by volunteers taking photos or scans of plants.

3. Caregiving. Situated visualizations can be used to assist caregivers, nurses, and doctors by providing them with information about patients in a medical facility or a care home. For example, Situated Glyphs [307] provide caregivers with situated medical information about patients, such as records of their recent incidents (e.g., a fall) or their schedule (e.g., lunchtime, recreational activities, family visitation). In this context, the **objects of interest** are the patients, and the patients may be **non-visible** if they are in a different room than the caregiver. Yet caregivers may still need to access the data, especially for monitoring purposes or in the presence of a suspected emergency. Despite all their advantages, situated physical displays like Situated Glyphs have a poor **non-visibility support**. Medical facilities often have control rooms for monitoring patient vitals, but when caregivers are neither in the control room nor in the patient's room (e.g., they are moving between patients), they lose access to the information. Therefore, an alternative could be to equip caregivers with lightweight AR displays where they could see both a **representation** of the patients and **visualizations** conveying important information about them. In

the context of caregiving, a high **visualization refresh rate** is important for patients under intensive care or who are at risk of an accident. If keeping an overview of all patients at all times is a desirable option, then the system would need to use a low **visualization richness** to prevent clutter and offer the user the option to increase visualization richness for specific patients of interest. In this case, having a high **representation fidelity** as well as a high **representation refresh rate** could potentially convey important qualitative information that complements the quantitative information captured by sensors.

5.6.2 Link to Other Conceptual Frameworks

In this section, we discuss the differences and similarities between our terminology and design space, and other conceptual frameworks.

Non-visibility support. In the ProxSituating Visualization paper [267] we already mentioned in Section 4.2, the authors relate the visibility of the referent to *spatial indirection*, i.e., the extent to which the physical referent is far from the visualization and the user. Our terminology on non-visibility (Section 5.3.2) is more precise and considers whether the physical referent can actually be seen irrespective of spatial indirection (see subsection 2.2.1) and, if not, why. Our *non-visibility support* dimension also relates to the *scope* dimension of Willett et al.’s [321] visualization-as-superpower framework: the more the kinds of non-visibility supported, the larger the scope of the situated visualization system.

Interactivity and interaction realism. Those dimensions were initially inspired by the *interaction model* dimension in Elmqvist’s [121] taxonomy of 3D occlusion management for visualization. Like their taxonomy, our dimensions aim to capture how the user interacts with the system: are they active or passive, what kind of gestures do they use, are they in control? We first reused Elmqvist’s *active/passive* dimension. However, in XR environments, users are rarely completely passive – we capture this with our *interactivity* dimension, which lies on a continuum. This dimension also maps to the *degree of control* dimension in Willett et al.’s [321] visualization-as-superpower framework. In addition to the level of interactivity, users can be active in very different ways; In XR research, interactions that go beyond classical interaction styles are often qualified as embodied [106] or reality-based [181]. The two terms are similar, but the former emphasizes the use of the body, while the latter emphasizes interactions that are close to real-world actions. We chose the latter term for our dimension interaction realism because we found that it made it easier to classify techniques.

Actionability. In contrast with Elmqvist’s model, we make the distinction between interacting with visualizations (*interactivity*) and interacting with physical referents (*actionability*). A similar distinction is made in Thomas et al.’s [296] characterization of situated visualizations, which distinguishes between three levels of interaction: altering the visualization pipeline, directly altering the visualization’s physical presentation, and directly altering the visualization’s physical referent. Actionability also relates to the

notion of *pragmatic superpower* in Willett et al.’s [321] visualization-as-superpower framework.

Representation fidelity and representation / visualization medium. These three dimensions were initially inspired by a design space of worlds-in-miniature introduced by Danyluk and al. [109], and in particular by its *abstraction* dimension. The ProxSituating framework [267], despite having been developed independently from ours, also includes the related dimension of *representation form* and in particular *tangibility*, which maps to our *representation/visualization medium* dimensions. The dimensions in our category *visualization* were also broadly inspired by Ens and al.’s [125] paper on Grand Challenges in Immersive Analytics, in which they discuss data representation in immersive situated visualizations.

Surroundings inclusion. This dimension maps to the concept of *environment* in the ProxSituating visualization framework [267].

Visualization richness. This dimension maps to the *information richness* dimension in Willett et al.’s [321] visualization-as-superpower framework, and the notion of *information bandwidth* discussed in Ens and al.’s [125] Grand Challenges paper. It is also close to the *semantic relevance* dimension in the taxonomy of annotation in outdoor augmented reality [324], and the concept of *value* in Lin et al.’s [210] taxonomy of labeling techniques.

Representation/visualization refresh rate. These dimensions relate to the notion of *temporal indirection* in the embedded representation [322] and the ProxSituating [267] frameworks, and to the notion of *temporal relevance* in the visualization-as-superpower framework [321].

Visualization integration. This dimension relates to the distinction between situated and embedded visualization in the embedded representation framework [322]: embedded visualizations are higher on the integration dimension than non-embedded situated visualizations.

Dimensions values. The choice to use three values (low, medium, high) in most of our dimensions was inspired by the design space of anthropographics by Morais and colleagues [229].

5.6.3 Descriptive and Generative Power

Our conceptual framework (terminology, typology of techniques, and design space) can help describe and compare situated visualization systems in terms of their support for non-visible referents – including the extent of this support and how this support is achieved. We successfully applied our design space to characterize the few existing situated visualization systems we know of that have some support for non-visible referents (see table in supplementary material section A.4). The General and Representation categories of our design space (see Table 5.1) can also be used to characterize the mixed-reality systems we surveyed in section 5.4, which support non-visible objects of interest but generally include no data visualization. We, however envision that our conceptual framework will be most useful in the hands

of designers of new situated visualizations to help them choose what types of non-visible referents they want to support and how. This ability to guide the design of new situated visualizations confers to our framework some generative power, which we illustrated through use cases in subsection 5.6.1. Our framework is probably less suited for designing novel techniques for handling hidden referents, although it provides an overview of existing techniques that can help identify gaps. Future work could involve validation of the framework’s descriptive power, for example, by asking designers to prototype situated visualizations with and without our framework.

5.6.4 Techniques not Discussed

In Section 5.4, we defined a scope for our survey and our design space, which helped us keep our framework simple but inevitably led us to exclude potentially interesting techniques. We discuss them here by revisiting the three main criteria for inclusion we listed in Section 5.4.

1. Visual only. We only considered using visual techniques to handle non-visible objects of interest. Although we are not aware of non-visual techniques in the context of situated visualization, Prouzeau et al. [254] proposed the use of vibrotactile feedback to manage occlusion and overplotting in immersive 3D scatterplots. Their technique encodes occluded information (e.g., the local density of points) with a VR controller’s vibration strength. Their study suggested that using haptic feedback is beneficial compared to a classic visual technique (in that case, a cutting plane) in cases where the scene is already visually complex. Nonetheless, the paper also showed that the haptic channel has a limited bandwidth compared to the visual channel. The use of non-visual channels, such as haptic and audio, deserves to be investigated to convey limited amounts of information about non-visible referents in situated visualizations, especially when the scenes are visually cluttered. Future work in this area would need to investigate how to link this non-visual information to the physical referents.

2. Object-aware. We assumed the existence of an object of interest that the system knows about, and we therefore excluded all object-agnostic approaches. Nevertheless, object-agnostic approaches could still help deal with non-visible referents in situated visualizations. For example, an x-ray vision technique that is object-agnostic and renders anything that is occluded could still help users see potential physical referents or their physical surroundings, even though it may not be able to link visualizations to their physical referents. Similarly, a situated visualization system could display feeds of video cameras placed in hidden areas where physical referents are likely to be situated, such as areas under surveillance [42], places where maintenance tasks are being carried out [325], or the parts of a body that are undergoing surgery [175]. Among object-agnostic techniques, some techniques exist that widen the user’s field of view, sometimes providing a 360° field of view [36]. Some of these techniques combine wide-angle views with regular views [130, 228]. All such techniques can help users see physical referents, although again, they cannot help them find referents or relate them with their corresponding visualizations.

3. No movement required. We only considered systems allowing users to see information about non-visible referents from where they stand and therefore excluded systems requiring them to move. Such systems could nevertheless facilitate the use of situated visualizations where physical referents are not always in view. For example, tour planning techniques can show users how to get to a specific location [35, 90, 96, 123], and could potentially guide them to a non-visible physical referent; once the user is there, the referent stops being non-visible. The out-of-view pointing techniques we discussed in subsection 5.4.7 sometimes serve the same function, but we included them nonetheless because the arrow does give immediate information about the physical referent (its location), albeit very limited information. Attention-guiding techniques are similar to tour planning, except they do not involve locomotion – only posture changes [288]. Another approach is directly moving the user to the non-visible object of interest. This is possible in virtual worlds using teleportation techniques (e.g., [169]), but such motions are virtual. In situated visualization systems, actually teleporting users is not possible, but autonomous vehicles could, in principle, be used to physically bring users closer to physical referents of interest.

5.7 Conclusion

We proposed a conceptual framework to help researchers and user interface designers think about how to handle non-visible physical referents in situated visualizations. This framework consists of 1) a terminology, 2) a typology of techniques, 3) a design space. Our framework reflects the diversity of mixed-reality techniques that can be used to address various situations in which physical referents are not visible. It also highlights areas that remain unexplored and opportunities for future research. However, our framework is far from capturing all aspects relevant to the design of situated visualizations that can handle hidden referents, nor does it cover all possible designs. Many considerations go into the design of such interactive systems, and a single design space cannot possibly capture them all. However, we believe our conceptual framework is a useful starting point for discussion and reflection. In particular, we hope that our terms and concepts will make it easier for researchers and practitioners to think and discuss about the issue of non-visible physical referents in situated visualizations and that they will be extended to cover broader cases. On a general level, we hope that this work will open interesting discussions and inspire more research about this neglected topic and, more broadly, the question of how to consider the complexity of physical-world constraints when designing situated visualizations.

This chapter closes the presentation of the different axes I worked on during this PhD. In the next chapter, I will summarize my contributions and clarify several perspectives for future research.

Chapter 6

Conclusion and Perspectives

In this PhD thesis, I explored immersive situated visualizations. Through different axes, I chose to consider immersive visualizations, in conjunction with the opportunities and constraints of the physical world, as the starting point for my reflections. In particular, I envisioned exploring the usage and effects of concrete visualizations, proposing a way to use them to complement traditional ways of visualizing data, and investigating how the current state of immersive technology presents opportunities for implementing situated visualizations in real-world settings, with a specific emphasis on non-visible objects. I will first summarize these works and recall clearly my contributions before presenting stemming opportunities and perspectives for future research.

6.1 Summary

In chapter 2, I provided a global background to the reader about immersive situated visualizations. I introduced various challenges faced by the visualization and MR-HCI community, which puts our exploration choices into context.

Within chapter 3, I proposed a new type of visualization for eco-feedback, which we qualify as concrete, inspired by the physical world, and intended to be displayed within the physical world. After giving specific related work about concrete visualizations, immersive concrete visualizations, and existing systems and studies from the eco-feedback research community, I:

- Identified and named this family of eco-feedback visualization techniques (ARWAVs) that have received very little attention so far and presented scenarios illustrating the different situations in which they can be useful.

- Described prototypes that implement such visualization techniques and that were tested in an initial feedback session, both with colleagues and with a larger public during a tech exhibition event;
- Introduced a terminology and a set of design considerations for ARWAVs.
- Reported a user study suggesting that ARWAVs are emotionally more engaging than less immersive display modalities (3D on screen) and simpler information representations (numerals).

After assessing some value for immersive concrete visualizations, in chapter 4, I presented explorations on animated transitions to use concrete visualizations as complementary visualizations to traditional representations. After going through some background around unit visualizations and animated transitions and presenting scenarios for which these animated transitions would be useful, I:

- Presented a design space with 4 axes (Marks, Embedding, View, and Scale) to frame our approach for designing these transitions;
- Presented a controlled experiment in VR, yielding quantitative and qualitative evidence that the [ME-VS] condition (which displays changes in Marks and Embedding first, followed by View and Scale transformations) outperforms the [MEVS] condition (a transition showing all transformations simultaneously over the same total duration);
- Proposed design guidelines for immersive animated transitions between concrete and abstract representations.

In chapter 5, I tackle another challenge: managing physical world constraints when implementing situated visualizations in the particular case of a temporarily or permanently non-visible referent. After presenting illustrative scenarios and background around occlusion management and the issue of non-visible referents, I:

- Laid a terminology characterizing different types of non-visible objects.
- Presented a typology of MR techniques handling non-visible referents.
- Proposed a Design Space characterizing those techniques that apply to situated visualizations.

I believe these contributions fill gaps in research around immersive situated visualization and hope they inspire the visualization and HCI community. In the next section, I clarify some perspectives of research that I wish to grasp or let other researchers take hold of.

6.2 Future work and Perspectives

6.2.1 Exploratory Data Analysis with Concrete Visualizations

In chapter 4, we have mostly focused on conveying waste quantities when the data consists of a single quantity to be conveyed. However, waste data is dynamic by nature and can have rich attributes associated with it, such as the nature of the waste, where it was produced, and by whom. This can also be exemplified by the book example presented in the section 4.1 of chapter 4: What types of books have been borrowed? How long did their borrower keep them? This opens up a vast array of possibilities for interactive data exploration.

Interaction techniques used in traditional computer-based visualizations could be mapped to AR gestures involving remote or direct touch. For example, a concrete visualization could support zooming through a pinch gesture and let users go back and forth between a 1:1 view and a miniature view (physical to reduced scale on Figure 4.5). Gestures could also be used to change the time scale of a concrete representation, e.g., to see the amount of trash produced or books borrowed in a week, a month, or a year; To facet the data by sub-categories, e.g., see the trash amount recycled or not, or the amount produced per person; Or to make comparisons, assess trends and distributions (e.g., see twelve piles of trash bags or books next to each other, one for each month – also see the device in Figure 3.2c, which supports comparison between two periods). These embodied transitions could exist between concrete and abstract representations and support animated transitions presented chapter 4 (see an early prototype of an embodied animated transition between a bar chart and a volume of waste at <https://youtu.be/XTwYBnhC8xQ>). Such interactions could help users understand patterns in their consumption habits and perhaps help them find ways to optimize their behavior. Considering augmented-reality visualizations, concrete visualizations such as ARWAVs could support interactions that make smart use of the physical environment and go beyond what is possible on desktop or virtual-reality displays. For example, an ARWAV system like in Figure 3.4a could first show an overview of total water consumption in the house and let the user break down this data by consumption source (e.g., bathroom shower, bathroom sink, toilet, kitchen sink, etc.). This transition could be shown with an animation of a large volume of water splitting into smaller volumes, each moving next to its respective source. Such a feature could help users identify activities for which changes in behavior are likely to be the most impactful, and transitions between concrete and abstract visualizations presented chapter 4 are just entry points to this endeavor.

We have also mostly explored the use of concrete visualizations such as ARWAVs made of collections of simple and identically-looking objects, but object appearance (color, shape) can be used to encode richer information. For example, a pile of plastic cups could be composed of crushed cups representing cups that have been used and disposed of, green cups representing the part that will be recycled, and new cups

for the cups that have not been thrown away. Finally, the prototypes we have presented only implement basic audio effects, but work in immersive analytics has already demonstrated the use of haptics [254] and sound [233] to encode richer data and enhance data exploration.

6.2.2 Immersive Data Storytelling with Concrete Visualizations

Data-driven stories “*seek to inform and persuade audiences through the use of data visualizations and engaging narratives*” [332]. VR and AR immersiveness has been identified as a valuable addition to 2D-based data-driven stories. Zhou et al. [332] highlight opportunities for transitions between different displays, creating a hybrid space including screens and immersive displays. In particular, they point out the possibility of leveraging other types of representations thanks to immersive displays (e.g., visualizing topographic water levels *concretely* while reading about it on a screen).

The animated transitions detailed in chapter 4 could enhance data-driven storytelling by leveraging a unique visualization medium. Users could immerse themselves in a data-driven narrative within a VR environment, manipulating abstract and concrete representations. These narratives could also be experienced in the real world through AR, potentially at the location pertinent to the visualized data, resulting in what I would term *immersive situated data-driven stories*. For example, while following a narrative about French food habits in a kitchen setting, a user could interact with virtual representations of statistics regarding the health of French citizens or the state of intensive farming in the country. This experience could involve transitions between abstract and concrete representations (e.g., using actual piles of fruits, cereals, and vegetables to illustrate each proportion or employing concrete representations of filled chicken cages to depict average chicken density on farms).

Zhou et al.’s concept of employing various modalities also presents an intriguing perspective to expand upon the work discussed in chapter 4. Our transitions could initiate with on-screen bar charts and culminate with concrete representations in the physical space. This extends to other mediums, such as scientific posters or books, which typically feature numerous abstract (scientific) visualizations. These could be augmented like a MagicBook [69] (a physical book with 3D animated virtual object pop-ups). For example, a student in the field of life and earth sciences could read about the distribution of various types of volcanic rock in a particular area through a graph and then access a concrete representation through their smartphone, revealing the layers of the Earth and the actual distribution of these rocks. Notably, they could transition between the book and the 3D representation to correlate numeric and spatial information.

Of course, these ideas would need various user experiments, notably to assess the user’s performance in understanding multiple transitions occurring at the same time (i.e., animated transitions between different representations and between different displays) or understanding and following a transition while keeping

up with a story.

6.2.3 ARWAVs for a Sustainable Behaviour Change

6.2.3.1 Ecological cost of ARWAVs

A seeming paradox with ARWAVs is that they are meant to promote eco-friendly behavior but are costly to the environment due to their reliance on AR hardware. However, it is possible that at least some forms of ARWAVs are so effective at changing behavior that their benefits outweigh the environmental costs of the hardware necessary to run them. The net ecological benefit of an ARWAV likely depends on a range of factors, including the type of AR display technology it uses. For example, most people in industrialized countries already own AR-ready smartphones, so running ARWAVs like the one in 3.4a is possible at almost no extra environmental cost. In contrast, AR-HMD devices and spatial AR installations are uncommon and require many hardware components. Nevertheless, a single device can be shared across people and can reach hundreds of people in exhibitions and other public events (see, e.g., Scenario 3). In addition, the enhanced sense of immersion brought by AR-HMDs and spatial AR installations may make ARWAVs more effective than smartphones, which could justify their use. Doing cost-benefit analyses like these will be crucial but require empirical knowledge from user studies. Some ARWAV designs (software and hardware) may have a net ecological cost while others have a net ecological benefit, so not studying ARWAVs just because of the environmental cost of AR hardware could cause us to miss important opportunities.

6.2.3.2 ARWAVs for Behaviour Change

As just explained, assessing the benefits of such AR visualizations seems urgent. Our work on ARWAVs initiated discussions among our team and participated in the idea of proposing an ANR project called *Be · Aware*¹, that joins the efforts of our team (Potioc working on HCI-Visu), the Lessac team from Dijon specialized in Behavioural Economics² and the CIRED based in Paris leading research in Environmental Sciences³.

Using controlled experiments in behavioral economics and real environmental data, this interdisciplinary research aims at evaluating behavior change induced by interactive visualizations with game-theoretical models (see Figure 6.1). More precisely, one objective is to help reduce the *cognitive distance* between people and the actual environmental consequences of their actions, using interactive visualizations. Environmental issues such as climate change may express themselves more intensively in a far

¹<https://beaware.inria.fr/>

² <https://lessac.bsb-education.com/index.php?page=presentation>

³ <https://www.centre-cired.fr/>

temporal or geographic distance (from an occidental point of view), and that influences individuals' attitudes and behaviors. For instance, people tend to feel more indifferent to events situated far in space and time [330], which affects their pro-activity to carry pro-environmental behaviors. Using visualization strategies to reduce cognitive distances encompasses embedding ARWAVs in a user's direct surroundings or displaying on-screen images such as Figure 3.14. Currently, the Lessac team is evaluating ARWAVs prototypes for this endeavor.




Figure 6.1 Be · Aware project process.

This approach already gathered interest within the HCI community, which led several members of Be · Aware to participate in a CHI'23 Workshop [183]. Aside academics, several territorial organizations contacted us following an article in *The Conversation France*⁴, a popular science journal addressing current topics directly from the academic community. Local authorities lack a diverse range of resources to effectively raise awareness among their residents about matters such as waste recycling or energy consumption. We met with the Bordeaux city council to explore concrete means of collaboration. This is a kind of collaboration I'm willing to explore in the future, will that be in research or another environment.

6.2.4 Engage Multiple and Diverse Display Mediums when Designing Immersive Situated Visualizations

Head-mounted AR is the dominant technology envisioned for immersive situated visualizations, but many opportunities remain poorly explored [81]. As extensively discussed in this thesis, situated visualizations are inherently tied to the physical world. Using fully digitalized overlaying is one approach to designing immersive situated visualizations. Another approach is to build situated physicalizations as we discussed section 3.7 (i.e., *“physical artifact whose geometry or material properties encode data”* [184]), which present favors user engagement and collaboration [248]. In chapter 3, we took inspiration from several physical accumulations whose natural existence poses specific hygienic limitations. However, their artifi-

⁴<https://theconversation.com/546-kilos-de-dechets-dans-votre-salon>



(a) Balloons to convey CO2 emission, from <https://interface.fh-potsdam.de/dataobjects/>, credit: Boris Müller and Fabian Morón Zirfas.



(b) Spherical wooden sculpture to convey CO2 emission from <https://semiconductorfilms.com/art/cosmos/>, credit: Ruth Jarman and Joe Gerhardt.

Figure 6.2 Two physicalizations to convey CO2 emissions. via <http://dataphys.org/list/>

cial and physical recreation could be valuable for concrete representations, such as waste accumulations. These objects could be manipulated and placed directly in the user's immediate environment, potentially enhancing user engagement. It would be possible to physically encode additional information related to the visualized data. This type of artifact has already been used, such as in Figure 6.2a and Figure 6.2b in which carbon dioxide emissions are represented with either balloons or with a 2-meter spherical wooden sculpture in the woods.

Taking the example of garbage bags, green bags could signify recyclable materials, while neutral-colored bags could represent non-recyclables, akin to what we propose in subsection 6.2.1. Other attributes, subtly encoded or partially concealed, could convey specific information to the user through manipulation, such as filling some bags with plastic and others with glass to demonstrate the proportion of each type of waste material. While it is conceivable to achieve this holographically, thereby circumventing practical implementation, replicability or updating hurdles, the trade-offs compared to the alternative I have proposed need to be clarified and warrant further research. Our collaborators at Lessac (see subsection 6.2.3) have employed this artifact to present the Be·aware project by populating a scene with garbage bags. Although this required substantial preparation, it was the most straightforward approach to enable all viewers to visualize waste accumulations on a large scale, complementing the videos displayed on the screen. Alternatively, other scenarios are better suited for using physical objects, particularly when embedding representations in environments like the sky or the sea (as in [248]), which are less conducive to AR visualization. Additionally, we can consider shape-changing physical objects in the same way as shape-changing UI [192] to get closer to the malleability of the conceivable forms of virtual objects. This raises questions such as: in the context of a collaborative task involving both AR and physical conditions, which would promote greater collaboration? We have demonstrated that waste stacking elicits stronger

emotional responses when shown in AR rather than on a screen, but what about physical accumulations?

To take this a step further, physicalizations offer a promising avenue to explore, but it would also be valuable to investigate possibilities involving both physical and virtual objects. This would manifest as a hybrid space where some elements are physical (that may be more engaging) while others are virtual (that may allow interactivity and flexibility in the representation).

6.2.5 More on the Study of Situatedness

Situated visualizations appear promising for several reasons already discussed. However, there still needs to be studies affirming their effectiveness. For instance, comparative studies to evaluate situated versus non-situated or situated versus embedded visualization conditions. Also, it would be valuable to understand better the impacts of temporal and spatial indirections (see subsection 2.2.1) on data comprehension and decision-making in context. Consider the example from the study conducted in chapter 3. The AR condition represents waste-related data for the year preceding the current year and where the waste was generated. Investigating the implications of temporal data indirection could be achieved by using the same metric (or a different one evaluating understanding, fatigue, or arousal, for instance) with data explicitly dated from the previous month or a decade ago (while keeping as many variables constant, such as considering a population equally affected by this data). Measuring the implications of spatial indirection would involve comparing the embedding of the representation of this data in the cafeteria, in the Inria parking lot, or the participant's personal space. When we add immersive technologies to the mix, questions arise regarding the influence of the physicality of contextual information (i.e., of the referent's representation). Finally, one could evaluate whether the same experience in a VR replica of the cafeteria frequented by the participants would yield similar results.

6.2.6 Human Perception of the physical world with Immersive Situated Visualizations

Genay et al. [143] demonstrated the potential for enhanced understanding of our abilities in the physical world at a distance through the display of one's avatar in Augmented Reality. These findings inspire the possibility of conducting similar investigations using concrete representations to understand the physical world better. The study by Genay et al. aimed to better comprehend certain distances between two points on participants' bodies and in physical space. An AR application from Ikea⁵ allows a user to place a 3D piece of furniture (which can be considered a concrete representation of data related to dimensions, materials, and color of the furniture) in their physical living room via their phone, partially enabling this kind of exercise, as the dimensions of the furniture are proportionate to physical space.

⁵<https://play.google.com/store/apps/details?id=com.ingka.ikea.app&hl=au>

We could contemplate conducting similar studies to assess if certain concrete representations facilitate a better understanding of the physical world. For instance, estimating specific quantities relative to pre-defined containers can be challenging (as previously addressed in[94]). Some of our ARWAV prototypes lend themselves well to this purpose. Consider the example of the water volume in Figure 3.4a, visualized through a smartphone. Suppose I display my weekly water consumption; would this assist me in estimating the number of water tanks I have used? Considering the trash bag prototype, would representing my waste consumption for any amount of time near my bins enable me to estimate when they are full in order to estimate the frequency of waste collection my household requires more easily?

6.2.7 Immersive Situated Art

More and more AR-using influencers are flourishing on social networks (like PiperZy⁶). They create paintings, sculptures, clothes, jewelry, and accessories intended to be virtually augmented. Although these are not visualizations, the 3D elements added to these objects are situated. However, user perception of such objects is not academically studied. Particularly in the case of clothing and accessories often commissioned by luxury brands and worn at social events, these elements are part of the social expression of individuals. This is another step towards a world where AR becomes a social tool. This raises interesting questions, such as: Does AR contribute to the perceived value of an object? How do these augmented objects affect social relations between individuals? How does it change the artistic experience of users? Can such augmentations contribute to purchasing decisions?



Figure 6.3 PiperZy's work on AR-art and digital fashion.

⁶<https://piperzy.com/>

6.2.8 Reflections on ethical stakeholders associated with Immersive Situated Visualizations

In 1977, Roman Gary wrote in *Charge d'âmes*, “*You always have to know the limits of what’s possible. Not to stop, but to undertake the impossible under the best possible conditions.*”⁷.

This citation underscores the meticulous precautions that must be taken when considering the perspectives we present in this chapter. While they offer exciting and potentially empowering opportunities for humans, enabling them to better comprehend data for informed decision-making, numerous uncertainties loom over the implications of using AR as a non-expert tool. Indeed, the effects of AR on human cognition, accessibility, and the social and environmental repercussions it carries are subjects of concern. Although seamlessly integrating virtual elements into the environment holds great promise (which I hope I tinily helped to nurture and identify in this manuscript), it also contributes to individuals’ challenges in distinguishing between reality and virtuality, which could have dramatic consequences in the future. These discussions are already underway within the scientific community, which is reassuring; nevertheless, it is imperative to advance these elements of research further and initiate extensive communication beyond the academic world, ensuring the safety of users who will benefit from the industrial outcomes of these investigations.

Acknowledgements

This research was supported by the Agence Nationale de la Recherche (ANR) (grant number ANR-19-CE33-0012), NSERC, and by the LaBRI mobility program.

⁷Translated from french.

Acknowledgements

I extend my sincere appreciation to my colleagues and loved ones. It is with great humility and gratitude that I find myself privileged to write this manuscript.

The past three years have been a remarkable journey, and I attribute a considerable part of the joy to the exceptional work environment that has enveloped me. Therefore, a heartfelt thank you to my esteemed advisors, Martin, Pierre and Arnaud. Beyond your academic guidance, it is the harmonious interplay of your distinct personalities that has provided a motivating and cohesive balance, pivotal to the success of this thesis.

Acknowledgment is also due to the entire Potioc/Bivvac team, as our shared discussions have often highlighted the unique qualities that each member brings. This collaborative environment, simultaneously dynamic and supportive, has created a space where one feels both challenged and nurtured. I will keep in mind this great time we had all together in Hamburg where we shared a lot of StoryTimes.

Thank you to Michael McGuffin for receiving me as an intern in Montreal where I had an amazing time enjoying the city and working under his great guidance.

A special expression of gratitude goes to my friends (Pamplight, Gang Go, les Foculs, Oceane and Ryan and al.). Your support have been a source of strength and inspiration.

Lastly, I extend my deepest thanks to my family, my aunt and cousins. To my parent: as the years pass, I increasingly recognize the shared values and experiences that form the foundation of who I am today. I am proud to emulate your admirable qualities. Special thank you to Ruben who is truly the best baby brother of all time. May you keep inspiring and impressing me as much as you do since you were born.

This thesis is not just a personal accomplishment but a collective success, and I am sincerely grateful to each one of you for being an integral part of it.

Bibliography

- [1] [Online]. Available: https://www.matec-conferences.org/articles/mateconf/pdf/2017/11/mateconf_etice2017_01098.pdf
- [2] “Desire path,” https://en.wikipedia.org/wiki/Desire_path, accessed: 2023.
- [3] “Filtering Shared Social Data in AR | Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems.” [Online]. Available: <https://dl.acm.org/doi/abs/10.1145/3170427.3188609>
- [4] “A Guide to VRML 2.0 - 2. VRML Browsers.” [Online]. Available: <http://www.agocg.ac.uk/train/vrml2rep/part1/guide2.htm>
- [5] “Location based experience design for mobile augmented reality,” in Proc. EICS, author = Fedosov, Anton and Misslinger, Stefan, year = 2014,.
- [6] “MapChecking,” <https://www.mapchecking.com/>.
- [7] “Microsoft HoloLens | Technologie de réalité mixte pour les entreprises.” [Online]. Available: <https://www.microsoft.com/fr-fr/hololens>
- [8] “Public visualization displays of citizen data_ Design, impact and implications | Elsevier Enhanced Reader.”
- [9] “Readings in Information Visualization: Using Vision to Think - Mackinlay Card - Google Livres.” [Online]. Available: https://books.google.fr/books?hl=fr&lr=&id=wdh2gqWfQmgC&oi=fnd&pg=PR13&ots=ooCK8wtINu&sig=-IU3vxBFn9iyhBnwOX1-ptDxKgU&redir_esc=y#v=onepage&q&f=false
- [10] “Snellen fraction.” [Online]. Available: <https://www.oxfordreference.com/view/10.1093/oi/authority.20110803100514387>
- [11] “Stava,” <https://www.strava.com/>, accessed: 2023.
- [12] “Tangible CityScape.” [Online]. Available: <https://tangible.media.mit.edu/project/tangible-cityscape>
- [13] “Ted talk by tom hulme,” https://www.ted.com/talks/tom_hulme_what_can_we_learn_from_shortcuts, accessed: 2023.
- [14] “Towards an Understanding of Situated AR Visualization for Basketball Free-Throw Training | Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.” [Online]. Available: <https://dl.acm.org/doi/abs/10.1145/3411764.3445649>
- [15] “U.S. Army to use HoloLens technology in high-tech headsets for soldiers - Transform.” [Online]. Available: <https://news.microsoft.com/transform/u-s-army-to-use-hololens-technology-in-high-tech-headsets-for-soldiers/>
- [16] “Virtual and augmented reality technologies for emergency management in the built environments: A state-of-the-art review | Elsevier Enhanced Reader.”
- [17] “Virtual Becomes Reality: A Stanford VR Experience sur Steam.” [Online]. Available: https://store.steampowered.com/app/1658530/Virtual_Becomes_Reality_A_Stanford_VR_Experience/

BIBLIOGRAPHY

- [18] “What’s the Situation with Situated Visualization? A Survey and Perspectives on Situatedness.” [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03319648/document>
- [19] “World-in-Miniature Interaction for Complex Virtual Environments: Media & Communications Book Chapter | IGI Global.” [Online]. Available: <https://www.igi-global.com/chapter/content/64047>
- [20] <https://www.douyin.com/video/7026283424738463010>, 2021.
- [21] “All the metals we mined,” <https://elements.visualcapitalist.com/all-the-metals-we-mined-in-one-visualization/>, 2021.
- [22] “Occlusion Handling in Augmented Reality: Past, Present and Future,” *IEEE TVCG*, 2021.
- [23] <https://www.douyin.com/video/7088254656287198478>, 2022.
- [24] <https://www.douyin.com/video/7083800279065382152>, 2022.
- [25] “À Marseille, la répétition des grèves de la collecte des ordures exaspère certains habitants,” Jan. 2022, section: Infos. [Online]. Available: <https://www.francebleu.fr/infos/economie-social/a-marseille-la-repetition-des-greves-de-la-collecte-des-ordures-exasperent-certains-habitants-1643046035>
- [26] ADEME, “Quelles obligations réglementaires ?” <https://expertises.ademe.fr/entreprises-monde-agricole/reduire-impacts/reduire-cout-dechets/obligations-reglementaires>, [Online; accessed 30-November-2022].
- [27] S. J. G. Ahn, J. N. Bailenson, and D. Park, “Short- and long-term effects of embodied experiences in immersive virtual environments on environmental locus of control and behavior,” *Computers in Human Behavior*, vol. 39, pp. 235–245, Oct. 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0747563214003999>
- [28] A. Ajanki, M. Billinghurst, H. Gamper, T. Järvenpää, M. Kandemir, S. Kaski, M. Koskela, M. Kurimo, J. Laaksonen, K. Puolamäki, T. Ruokolainen, and T. Tossavainen, “An augmented reality interface to contextual information,” *Virtual Reality*, vol. 15, no. 2-3, pp. 161–173, 2011.
- [29] G. Alce, M. Roszko, H. Edlund, S. Olsson, J. Svedberg, and M. Wallergård, “[POSTER] AR as a User Interface for The Internet of Things—Comparing Three Interaction Models,” in *ISMAR-Adjunct*, 2017.
- [30] F. Alpizar, F. Carlsson, G. Lanza, B. Carney, R. C. Daniels, M. Jaime, T. Ho, Z. Nie, C. Salazar, B. Tibesigwa et al., “A framework for selecting and designing policies to reduce marine plastic pollution in developing countries,” *Environmental Science & Policy*, vol. 109, pp. 25–35, 2020.
- [31] F. Amini, N. Henry Riche, B. Lee, C. Hurter, and P. Irani, “Understanding data videos: Looking at narrative visualization through the cinematography lens,” in *ACM CHI*, 2015, pp. 1459–1468.
- [32] F. Amini, N. H. Riche, B. Lee, J. Leboe-McGowan, and P. Irani, “Hooked on data videos: Assessing the effect of animation and pictographs on viewer engagement,” in *AVI*, 2018. [Online]. Available: <https://doi.org/10.1145/3206505.3206552>
- [33] V. Amrhein, S. Greenland, and B. McShane, “Scientists rise up against statistical significance,” *Nature*, vol. 567, pp. 305–307, 2019.
- [34] C. Andujar, F. Argelaguet, and R. Trueba, “Hand-Based Disocclusion for the World-in-Miniature Metaphor,” *Presence: Teleoperators and Virtual Environments*, vol. 19, no. 6, 2010.
- [35] C. Andujar, P. Vázquez, and M. Fairén, “Way-Finder: guided tours through complex walkthrough models,” *Computer Graphics Forum*, vol. 23, no. 3, 2004.
- [36] J. Ardouin, A. Lécuyer, M. Marchal, C. Riant, and E. Marchand, “FlyVIZ: a novel display device to provide humans with 360° vision by coupling catadioptric camera with hmd,” in *Proc. of VRST*, 2012.
- [37] P. Arpaia, E. De Benedetto, C. A. Dodaro, L. Duraccio, and G. Servillo, “Metrology-Based Design of a Wearable Augmented Reality System for Monitoring Patient’s Vitals in Real Time,” *IEEE Sensors Journal*, vol. 21, no. 9, 2021.

BIBLIOGRAPHY

- [38] C. Arth and D. Schmalstieg, "Challenges of large-scale augmented reality on smartphones," Oct. 2011, *iSMAR 2011 : ISMAR 2011, ISMAR 2011* ; Conference date: 26-10-2011 Through 29-10-2011.
- [39] A. Assor, A. Prouzeau, P. Dragicevic, and M. Hachet, "Exploring augmented reality waste data representations for eco feedback," in *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. *CHI EA '23*. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: <https://doi.org/10.1145/3544549.3583905>
- [40] —, "Augmented-reality waste accumulation visualizations," *ACM J. Comput. Sustain. Soc.*, jan 2024, just Accepted. [Online]. Available: <https://doi.org/10.1145/3636970>
- [41] A. Assor, A. Prouzeau, M. Hachet, and P. Dragicevic, "Handling non-visible referents in situated visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 1, pp. 1336–1346, 2024.
- [42] B. Avery, W. Piekarski, and B. H. Thomas, "Visualizing Occluded Physical Objects in Unfamiliar Outdoor Augmented Reality Environments," in *Proc. of ISMAR, 2007*.
- [43] B. Avery, C. Sandor, and B. H. Thomas, "Improving Spatial Perception for Augmented Reality X-Ray Vision," in *2009 IEEE Virtual Reality Conference*, Mar. 2009, pp. 79–82, ISSN: 2375-5334.
- [44] R. T. Azuma, "A survey of augmented reality," *Presence: teleoperators & virtual environments*, vol. 6, no. 4, pp. 355–385, 1997.
- [45] B. Bach, N. Henry-Riche, T. Dwyer, T. Madhyastha, J.-D. Fekete, and T. Grabowski, "Small multipiles: Piling time to explore temporal patterns in dynamic networks," *IEEE CG&A*, vol. 34, no. 3, pp. 31–40, 2015.
- [46] B. Bach, E. Pietriga, and J.-D. Fekete, "Graphdiaries: Animated transitions and temporal navigation for dynamic networks," *IEEE TVCG*, vol. 20, no. 5, pp. 740–754, 2013.
- [47] B. Bach, R. Sicat, J. Beyer, M. Cordeil, and H. Pfister, "The hologram in my hand: How effective is interactive exploration of 3d visualizations in immersive tangible augmented reality?" *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 457–467, 2017.
- [48] J. O. Bailey, J. N. Bailenson, J. Flora, K. C. Armel, D. Voelker, and B. Reeves, "The impact of vivid messages on reducing energy consumption related to hot water use," *Environment and Behavior*, vol. 47, no. 5, pp. 570–592, 2015. [Online]. Available: <https://doi.org/10.1177/0013916514551604>
- [49] R. Bane and T. Hollerer, "Interactive tools for virtual x-ray vision in mobile augmented reality," in *Proc. of ISMAR, 2004*.
- [50] Q. Bao, E. Burnell, A. M. Hughes, and M. C. Yang, "Investigating User Emotional Responses to Eco-Feedback Designs," *Journal of Mechanical Design*, vol. 141, no. 2, Dec. 2018. [Online]. Available: <https://doi.org/10.1115/1.4042007>
- [51] Q. Bao, M. M. Shaukat, A. Elantary, and M. C. Yang, "Eco-Feedback Designs: A Balance Between the Quantitative and the Emotional." *American Society of Mechanical Engineers Digital Collection*, Dec. 2016. [Online]. Available: <https://asmedigitalcollection.asme.org/IDETC-CIE/proceedings/IDETC-CIE2016/50190/V007T06A022/258466>
- [52] W. Barfield, D. Zeltzer, T. Sheridan, and M. Slater, "Presence and performance within virtual environments," *Virtual environments and advanced interface design*, pp. 473–513, 1995.
- [53] S. J. Barnes, "Out of sight, out of mind: Plastic waste exports, psychological distance and consumer plastic purchasing," *Global Environmental Change*, vol. 58, p. 101943, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0959378018312081>
- [54] L. Bartram, "Design Challenges and Opportunities for Eco-Feedback in the Home," *IEEE Computer Graphics and Applications*, vol. 35, no. 4, pp. 52–62, Jul. 2015, conference Name: *IEEE Computer Graphics and Applications*.
- [55] L. Bartram, J. Rodgers, and K. Muise, "Chasing the negawatt: visualization for sustainable living," *IEEE Comput Graph Appl*, vol. 30, no. 3, pp. 8–14, Jun. 2010.
- [56] J. Bastiras and B. H. Thomas, "Combining Virtual Reality and Narrative Visualisation to Persuade," in *2017 International Symposium on Big Data Visual Analytics (BDVA)*, Nov. 2017, pp. 1–8.

BIBLIOGRAPHY

- [57] P. Baudisch and R. Rosenholtz, “Halo: a Technique for Visualizing Off-Screen Locations,” NEW HORIZONS, no. 5, p. 8, 2003.
- [58] P. Baudisch, D. Tan, M. Collomb, D. Robbins, K. Hinckley, M. Agrawala, S. Zhao, and G. Ramos, “Phosphor: explaining transitions in the user interface using afterglow effects,” in ACM UIST, 2006, pp. 169–178.
- [59] R. M. Baños, C. Botella, M. Alcañiz, V. Liaño, B. Guerrero, and B. Rey, “Immersion and Emotion: Their Impact on the Sense of Presence,” Feb. 2005, archive Location: 2 Madison Avenue Larchmont, NY 10538 USA Publisher: Mary Ann Liebert, Inc. 2 Madison Avenue Larchmont, NY 10538 USA. [Online]. Available: <https://www.liebertpub.com/doi/epdf/10.1089/cpb.2004.7.734>
- [60] M. Beaudouin-Lafon, “Designing interaction, not interfaces,” in Proceedings of the Working Conference on Advanced Visual Interfaces, ser. AVI '04. New York, NY, USA: Association for Computing Machinery, 2004, p. 15–22. [Online]. Available: <https://doi.org/10.1145/989863.989865>
- [61] R. K. Becker, Ana, “Is the NASDAQ in Another Bubble? A virtual reality tour of the NASDAQ.” [Online]. Available: <http://graphics.wsj.com/3d-nasdaq/>
- [62] B. Bell, S. Feiner, and T. Höllerer, “View management for virtual and augmented reality,” in Proc. of UIST, 2001.
- [63] B. Bell, T. Höllerer, and S. Feiner, “An annotated situation-awareness aid for augmented reality,” in Proc. of UIST, 2002.
- [64] J. Belo, M. Lystbæk, A. Feit, K. Pfeuffer, P. Kán, A. Oulasvirta, and K. Grønbaek, “Auit - the adaptive user interfaces toolkit for designing xr applications,” 10 2022.
- [65] A. Bezerianos, P. Dragicevic, and R. Balakrishnan, “Mnemonic rendering: an image-based approach for exposing hidden changes in dynamic displays,” in ACM UIST, 2006, pp. 159–168.
- [66] C. Bichlmeier, S. M. Heining, M. Feuerstein, and N. Navab, “The Virtual Mirror: A New Interaction Paradigm for Augmented Reality Environments,” IEEE Transactions on Medical Imaging, vol. 28, no. 9, 2009.
- [67] C. Bichlmeier, S. M. Heining, M. Rustae, and N. Navab, “Virtually Extended Surgical Drilling Device: Virtual Mirror for Navigated Spine Surgery,” in Proc. of MICCAI, 2007.
- [68] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose, “Toolglass and Magic Lenses: The See-Through Interface,” p. 8.
- [69] M. Billingham, S. Cambell, I. Poupyrev, K. T. H. Kato, W. Chinthammit, and D. Hendrickson, “Magic book: Exploring transitions incollaborative arinterfaces,” in Proc. of SIGGRAPH, vol. 87. Citeseer, 2000.
- [70] M. Billingham, H. Kato, and I. Poupyrev, “Tangible augmented reality,” in ACM SIGGRAPH ASIA 2008 courses on - SIGGRAPH Asia '08. Singapore: ACM Press, 2008, pp. 1–10. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1508044.1508051>
- [71] O. Bimber and R. Raskar, Spatial augmented reality: merging real and virtual worlds. CRC press, 2005.
- [72] R. Blanco, Z. Salazar, and T. Isenberg, “Exploring carnivorous plant habitats based on images from social media,” in IEEE VIS 2019: IEEE Conference on Visualization, 2019.
- [73] A. Bluff and A. Johnston, “Don’t Panic: Recursive Interactions in a Miniature Metaworld,” in Proc. of International Conference on Virtual-Reality Continuum and its Applications in Industry, 2019.
- [74] C. Boomsma, J. Goodhew, S. Goodhew, and S. Pahl, “Improving the visibility of energy use in home heating in England: Thermal images and the role of visual tailoring,” Energy Research & Social Science, vol. 14, pp. 111–121, Apr. 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214629616300056>
- [75] F. Bork, U. Eck, and N. Navab, “Birds vs. Fish: Visualizing Out-of-View Objects in Augmented Reality using 3D Minimaps,” in Adjunct Proc. of ISMAR, 2019.
- [76] F. Bork, C. Schnelzer, U. Eck, and N. Navab, “Towards Efficient Visual Guidance in Limited Field-of-View Head-Mounted Displays,” IEEE TVCG, vol. 24, no. 11, 2018.

BIBLIOGRAPHY

- [77] J. A. Boudell and B. A. Middleton, "Plot locator: An app for locating plots in the field," Applications in plant sciences, vol. 7, no. 12, p. e11311, 2019.
- [78] D. A. Bowman and R. P. McMahan, "Virtual reality: how much immersion is enough?" Computer, vol. 40, no. 7, pp. 36–43, 2007.
- [79] J. Boy and J.-D. Fekete, "The co2 pollution map: Lessons learned from designing a visualization that bridges the gap between visual communication and information visualization," in IEEE Conference on Information Visualization [Poster paper], 2014.
- [80] J. Boy, A. V. Pandey, J. Emerson, M. Satterthwaite, O. Nov, and E. Bertini, "Showing people behind data: Does anthropomorphizing visualizations elicit more empathy for human rights data?" in Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, ser. CHI '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 5462–5474. [Online]. Available: <https://doi.org/10.1145/3025453.3025512>
- [81] N. Bressa, H. Korsgaard, A. Tabard, S. Houben, and J. Vermeulen, "What's the situation with situated visualization? a survey and perspectives on situatedness," IEEE TVCG, vol. 28, no. 1, 2022.
- [82] J. Brodsky, "From subatomic to cosmic: Visually communicating extreme scale in science communication," 2014.
- [83] G. Caggianese, V. Colonnese, and L. Gallo, "Situated Visualization in Augmented Reality: Exploring Information Seeking Strategies," in Proc. of SITIS, 2019.
- [84] S. Card, "INFORMATION VISUALIZATION," in The Human-Computer Interaction Handbook, 2nd ed. CRC Press, 2007, num Pages: 36.
- [85] J. Carmigniani, B. Furht, M. Anisetti, P. Ceravolo, E. Damiani, and M. Ivkovic, "Augmented reality technologies, systems and applications," Multimedia tools and applications, vol. 51, pp. 341–377, 2011.
- [86] J. P. Carneiro, A. Aryal, and B. Becerik-Gerber, "Influencing occupant's choices by using spatiotemporal information visualization in Immersive Virtual Environments," Building and Environment, vol. 150, pp. 330–338, Mar. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0360132319300320>
- [87] J. P. Carneiro, M. P. Varnosfaderani, V. Balali, and A. Heydarian, "Comprehensible and interactive visualizations of spatial building data in augmented reality," in Computing in Civil Engineering 2019: Visualization, Information Modeling, and Simulation. American Society of Civil Engineers Reston, VA, 2019, pp. 79–86.
- [88] G. Carrus, P. Passafaro, and M. Bonnes, "Emotions, habits and rational choices in ecological behaviours: The case of recycling and use of public transportation," Journal of Environmental Psychology, vol. 28, no. 1, pp. 51–62, Mar. 2008. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0272494407000746>
- [89] T. Chandler, M. Cordeil, T. Czuderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, K. Marriott, F. Schreiber et al., "Immersive analytics," in 2015 Big Data Visual Analytics (BDVA). IEEE, 2015, pp. 1–8.
- [90] I. Chaturvedi, F. H. Bijarbooneh, T. Braud, and P. Hui, "Peripheral vision: a new killer app for smart glasses," in Proc. of IUI, 2019.
- [91] Z. Chen, Y. Su, Y. Wang, Q. Wang, H. Qu, and Y. Wu, "MARVisT: Authoring glyph-based visualization in mobile augmented reality," IEEE TVCG, vol. 26, no. 8, pp. 2645–2658, 2019.
- [92] F. Chevalier, P. Dragicevic, and S. Franconeri, "The not-so-staggering effect of staggered animated transitions on visual tracking," IEEE TVCG, vol. 20, no. 12, pp. 2241–2250, 2014.
- [93] F. Chevalier, N. H. Riche, C. Plaisant, A. Chalbi, and C. Hurter, "Animations 25 years later: New roles and opportunities," in AVI, 2016, pp. 280–287.
- [94] F. Chevalier, R. Vuillemot, and G. Gali, "Using Concrete Scales: A Practical Framework for Effective Visual Depiction of Complex Measures," IEEE Transactions on Visualization and Computer Graphics, vol. 19, no. 12, pp. 2426–2435, Dec. 2013, conference Name: IEEE Transactions on Visualization and Computer Graphics.

BIBLIOGRAPHY

- [95] L. Chittaro and S. Burigat, “3D location-pointing as a navigation aid in Virtual Environments,” in Proc. of AVI, 2004.
- [96] L. Chittaro, R. Ranon, and L. Ieronutti, “Guiding visitors of Web3D worlds through automatically generated tours,” in Proc. of WEB 3D, 2003.
- [97] L. Chittaro and I. Scagnetto, “Is semitransparency useful for navigating virtual environments?” in Proc. of VRST, 2001.
- [98] M. M. Chun, “Contextual cueing of visual attention,” Trends in Cognitive Sciences, vol. 4, no. 5, pp. 170–178, May 2000. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1364661300014765>
- [99] P.-A. Cinquin, P. Giraudeau, T. Lainé, and L. Cousin, “Erlen : Une interface tangible pour la réduction de la consommation énergétique sur le poste de travail,” in 30eme conférence francophone sur l’interaction homme-machine, ser. Demos, AFIHM, Ed., Brest, France, Oct. 2018, p. 2p. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01900054>
- [100] A. Cockburn, C. Gutwin, and A. Dix, “HARK No More: On the Preregistration of CHI Experiments,” in Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, ser. CHI ’18. New York, NY, USA: Association for Computing Machinery, Apr. 2018, pp. 1–12. [Online]. Available: <https://doi.org/10.1145/3173574.3173715>
- [101] A. Cockburn, A. Karlson, and B. B. Bederson, “A review of overview+detail, zooming, and focus+context interfaces,” ACM Comput. Surv., vol. 41, no. 1, pp. 2:1–2:31, Jan. 2009. [Online]. Available: <https://doi.org/10.1145/1456650.1456652>
- [102] D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiropoulos, and D. F. Keefe, “Slice WIM: a multi-surface, multi-touch interface for overview+detail exploration of volume datasets in virtual reality,” in Proc. of I3D, 2011.
- [103] C. J. S. Collins and G. R. Barnes, “The occluded onset pursuit paradigm: prolonging anticipatory smooth pursuit in the absence of visual feedback,” Exp Brain Res, vol. 175, no. 1, 2006.
- [104] M. Conlen, J. Heer, H. Mushkin, and S. Davidoff, “Cinematic techniques in narrative visualization,” 2023.
- [105] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer, “IATK: An immersive analytics toolkit,” in IEEE VR, 2019, pp. 200–209.
- [106] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott, “Imaxes: Immersive axes as embodied affordances for interactive multivariate data visualisation,” in Proc. of UIST, 2017.
- [107] G. Cumming, Understanding the New Statistics: Effect Sizes, Confidence Intervals, and Meta-Analysis. New York: Routledge, Jun. 2011.
- [108] —, “The new statistics: Why and how,” Psychological Science, vol. 25, no. 1, pp. 7–29, 2014, pMID: 24220629. [Online]. Available: <https://doi.org/10.1177/0956797613504966>
- [109] K. Danyluk, B. Ens, B. Jenny, and W. Willett, “A design space exploration of worlds in miniature,” in Proc. of CHI, 2021.
- [110] K. Danyluk, T. Ulusoy, W. Wei, and W. Willett, “Touch and beyond: Comparing physical and virtual reality visualizations,” IEEE TVCG, vol. 28, no. 4, pp. 1930–1940, 2020.
- [111] C. Dede, “Immersive Interfaces for Engagement and Learning,” Science, vol. 323, no. 5910, pp. 66–69, Jan. 2009, publisher: American Association for the Advancement of Science. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.1167311>
- [112] J. Dickinson, R. Crain, S. Yalowitz, and T. Cherry, “How Framing Climate Change Influences Citizen Scientists’ Intentions to Do Something About It,” The Journal of Environmental Education, vol. 44, pp. 145–158, Jan. 2013.
- [113] P. Dourish, Where the action is: the foundations of embodied interaction. MIT press, 2001.
- [114] P. Dragicevic, “Fair statistical communication in HCI,” in Modern Statistical Methods for HCI.

BIBLIOGRAPHY

- [115] —, “Statistical dances: Why no statistical analysis is reliable and what to do about it,” June 2017, <https://youtu.be/UKX9iN0p5%5fA>.
- [116] P. Dragicevic, A. Bezerianos, W. Javed, N. Elmqvist, and J.-D. Fekete, “Temporal distortion for animated transitions,” in *ACM CHI*, 2011, pp. 2009–2018.
- [117] P. Dragicevic, Y. Jansen, and A. Vande Moere, “Data physicalization,” *Handbook of Human Computer Interaction*, pp. 1–51, 2020.
- [118] T. Drey, “Data Visualisation using Augmented Reality with a Focus set on Head Mounted Displays and Collaborative Tasks,” 2018.
- [119] S. Drucker and R. Fernandez, “A unifying framework for animated and interactive unit visualizations,” *Microsoft Research*, 2015.
- [120] K. Einsfeld, A. Ebert, and J. Wölle, “Modified virtual reality for intuitive semantic information visualization,” in *Proc. Int. Conf. Information Visualisation (IV)*, 2008, pp. 515–520.
- [121] N. Elmqvist and P. Tsigas, “A Taxonomy of 3D Occlusion Management for Visualization,” *IEEE TVCG*, vol. 14, no. 5, 2008.
- [122] N. ElSayed, B. Thomas, K. Marriott, J. Piantadosi, and R. Smith, “Situated analytics,” in *2015 Big Data Visual Analytics (BDVA)*, 2015, pp. 1–8.
- [123] T. Elvins, D. Nadeau, R. Schul, and D. Kirsh, “Worldlets: 3D Thumbnails for Wayfinding in Virtual Environments,” *Journal of the ACM*, pp. 21–30, Aug. 1997.
- [124] M. R. Endsley and D. J. Garland, *Situation Awareness Analysis and Measurement*. CRC Press, Jul. 2000, google-Books-ID: WrJGDsjJakcC.
- [125] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne, T. Dwyer, J. Grubert, J. H. Haga, N. Kirshenbaum, D. Kobayashi, T. Lin, M. Olaosebikan, F. Pointecker, D. Saffo, N. Saquib, D. Schmalstieg, D. A. Szafir, M. Whitlock, and Y. Yang, “Grand challenges in immersive analytics,” in *ACM CHI*, 2021. [Online]. Available: <https://doi.org/10.1145/3411764.3446866>
- [126] B. Ens, S. Goodwin, A. Prouzeau, F. Anderson, F. Wang, S. Gratzl, Z. Lucarelli, B. Moyle, J. Smiley, and T. Dwyer, “Uplift: A Tangible and Immersive Tabletop System for Casual Collaborative Visual Analytics,” *IEEE TVCG*, 2020.
- [127] B. Ens, S. Goodwin, A. Prouzeau, F. Anderson, F. Y. Wang, S. Gratzl, Z. Lucarelli, B. Moyle, J. Smiley, and T. Dwyer, “Uplift: A Tangible and Immersive Tabletop System for Casual Collaborative Visual Analytics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1193–1203, Feb. 2021, conference Name: *IEEE Transactions on Visualization and Computer Graphics*.
- [128] O. Erat, W. A. Isop, D. Kalkofen, and D. Schmalstieg, “Drone-Augmented Human Vision: Exocentric Control for Drones Exploring Hidden Areas,” *IEEE TVCG*, vol. 24, no. 4, 2018.
- [129] M. T. Eren and S. Balcisoy, “Evaluation of X-ray visualization techniques for vertical depth judgments in underground exploration,” *The Visual Computer*, vol. 34, no. 3, 2018.
- [130] K. Fan, J. Huber, S. Nanayakkara, and M. Inami, “Spidervision: Extending the human field of view for augmented awareness,” in *Proc. of AH*, 2014.
- [131] S. Feiner and B. MacIntyre, “A Touring Hachine: Prototgping 3D Hobite Augmented Reatitg Sgsystems for Exptoring the Urban Environment,” *Personal Technologies*.
- [132] A. H. Fischer and A. S. R. Manstead, *The relation between gender and emotion in different cultures*, ser. *Studies in Emotion and Social Interaction*. Cambridge University Press, 2000, p. 71–94.
- [133] R. J. Fisher, “Social desirability bias and the validity of indirect questioning,” *Journal of consumer research*, vol. 20, no. 2, pp. 303–315, 1993.
- [134] J. J. Flannery, “The relative effectiveness of some common graduated point symbols in the presentation of quantitative data,” *Cartographica: The International Journal for Geographic Information and Geovisualization*, vol. 8, no. 2, pp. 96–109, 1971. [Online]. Available: <https://doi.org/10.3138/J647-1776-745H-3667>

BIBLIOGRAPHY

- [135] D. C. Foyle, A. D. Andre, and B. L. Hooley, "Situation Awareness in an Augmented Reality Cockpit: Design, Viewpoints and Cognitive Glue," in Proc of HCI Int., 2005.
- [136] A. D. Fredericks, Z. Fan, and S. I. Woolley, "Visualising the invisible: Augmented reality and virtual reality as persuasive technologies for energy feedback," in 2019 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCOM/IOP/SCI). IEEE, 2019, pp. 1209–1212.
- [137] J. Frey, R. Gervais, S. Fleck, F. Lotte, and M. Hachet, "Teegi: tangible EEG interface," in Proc. of UIST, 2014.
- [138] T. Friedrich, A. Prouzeau, and M. J. McGuffin, "The effect of increased body motion in virtual reality on a placement-retrieval task," in ACM VRST, 2021, pp. 1–5.
- [139] J. Froehlich, L. Findlater, and J. Landay, "The design of eco-feedback technology," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ser. CHI '10. New York, NY, USA: Association for Computing Machinery, Apr. 2010, pp. 1999–2008. [Online]. Available: <https://doi.org/10.1145/1753326.1753629>
- [140] J. Froehlich, L. Findlater, M. Ostergren, S. Ramanathan, J. Peterson, I. Wragg, E. Larson, F. Fu, M. Bai, S. Patel, and J. A. Landay, "The design and evaluation of prototype eco-feedback displays for fixture-level water usage data," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ser. CHI '12. New York, NY, USA: Association for Computing Machinery, May 2012, pp. 2367–2376. [Online]. Available: <https://doi.org/10.1145/2207676.2208397>
- [141] H. Fuchs, A. State, and J.-C. Bazin, "Immersive 3D Telepresence," Computer, vol. 47, no. 7, 2014.
- [142] S. Fukatsu, Y. Kitamura, T. Masaki, and F. Kishino, "Intuitive control of "bird's eye" overview images for navigation in an enormous virtual environment," in Proceedings of the ACM symposium on Virtual reality software and technology 1998 - VRST '98. Taipei, Taiwan: ACM Press, 1998, pp. 67–76. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=293701.293710>
- [143] A. Genay, A. Lécuyer, and M. Hachet, "What can i do there? controlling ar self-avatars to better perceive affordances of the real world," in 2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), 2022, pp. 450–459.
- [144] S. Ghosh, G. Bailly, R. Despouys, E. Lecolinet, and R. Sharrock, "SuperVision: Spatial Control of Connected Objects in a Smart Home," in Proc. of CHI EA, 2015.
- [145] M. Golparvar-Fard, F. Peña-Mora, and S. Savarese, "Monitoring changes of 3D building elements from unordered photo collections," in 2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops), Nov. 2011, pp. 249–256.
- [146] C. Gonzalez, "Does animation in user interfaces improve decision making?" in ACM CHI, 1996, pp. 27–34.
- [147] F. Gortana, F. von Tenspolde, D. Guhlmann, and M. Dörk, "Off the grid: Visualizing a numismatic collection as dynamic piles and streams," Open Library of Humanities, vol. 4, no. 2, 2018, <https://uclab.fh-potsdam.de/coins/>.
- [148] R. Grasset, T. Langlotz, D. Kalkofen, M. Tatzgern, and D. Schmalstieg, "Image-driven view management for augmented reality browsers," in Proc. of ISMAR, 2012.
- [149] U. Gruenefeld, A. E. Ali, S. Boll, and W. Heuten, "Beyond Halo and Wedge: visualizing out-of-view objects on head-mounted virtual and augmented reality devices," in Proc. of MobileHCI. ACM, Sep. 2018.
- [150] U. Gruenefeld, A. E. Ali, W. Heuten, and S. Boll, "Visualizing out-of-view objects in head-mounted augmented reality," in Proc. of MobileHCI, 2017.
- [151] U. Gruenefeld, D. Ennenga, A. E. Ali, W. Heuten, and S. Boll, "EyeSee360: designing a visualization technique for out-of-view objects in head-mounted augmented reality," in Proc. of SUI, 2017.
- [152] U. Gruenefeld, D. Lange, L. Hammer, S. Boll, and W. Heuten, "FlyingARrow: Pointing Towards Out-of-View Objects on Augmented Reality Devices," in Proc. of PerDis, 2018.

BIBLIOGRAPHY

- [153] U. Gruenefeld, L. Prädell, and W. Heuten, “Locating nearby physical objects in augmented reality,” in Proc. of MUM, 2019.
- [154] U. Gruenefeld, T. C. Stratmann, A. E. Ali, S. Boll, and W. Heuten, “RadialLight: exploring radial peripheral LEDs for directional cues in head-mounted displays,” in Proc. of MobileHCI, 2018.
- [155] U. Gruenefeld, T. C. Stratmann, W. Heuten, and S. Boll, “PeriMR: a prototyping tool for head-mounted peripheral light displays in mixed reality,” in Proc. of MobileHCI, 2017.
- [156] U. Gruenefeld, T. C. Stratmann, J. Jung, H. Lee, J. Choi, A. Nanda, and W. Heuten, “Guiding Smombies: Augmenting Peripheral Vision with Low-Cost Glasses to Shift the Attention of Smartphone Users,” in Proc. of ISMAR-Adjunct, 2018.
- [157] U. Gruenefeld, T. C. Stratmann, L. Prädell, and W. Heuten, “MonoculAR: a radial light display to point towards out-of-view objects on augmented reality devices,” in Proc. of MobileHCI-Adjunct, 2018.
- [158] R. Guarese, J. Becker, H. Fensterseifer, M. Walter, C. Freitas, L. Nedel, and A. Maciel, “Augmented situated visualization for spatial and context-aware decision-making,” in Proceedings of the International Conference on Advanced Visual Interfaces, 2020, pp. 1–5.
- [159] D. Guilmaine, C. Viau, and M. J. McGuffin, “Hierarchically animated transitions in visualizations of tree structures,” in AVI, 2012, pp. 514–521.
- [160] Y. Guo, N. Cao, L. Cai, Y. Wu, D. Weiskopf, D. Shi, and Q. Chen, “Datamator: An intelligent authoring tool for creating datamations via data query decomposition,” 2023.
- [161] S. Gustafson, P. Baudisch, C. Gutwin, and P. Irani, “Wedge: clutter-free visualization of off-screen locations,” in Proc. of CHI, 2008.
- [162] P. K. Han, W. M. Klein, B. Killam, T. Lehman, H. Massett, and A. N. Freedman, “Representing randomness in the communication of individualized cancer risk estimates: Effects on cancer risk perceptions, worry, and subjective uncertainty about risk,” Patient Education and Counseling, vol. 86, no. 1, pp. 106–113, 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0738399111000747>
- [163] S. Haroz, R. Kosara, and S. L. Franconeri, “ISOTYPE visualization: Working memory, performance, and engagement with pictographs,” in ACM CHI, 2015, pp. 1191–1200.
- [164] R. L. Harris, Information graphics: A comprehensive illustrated reference. Oxford University Press, USA, 1999.
- [165] G. Hastings, M. Stead, and J. Webb, “Fear appeals in social marketing: Strategic and ethical reasons for concern,” Psychology & Marketing, vol. 21, no. 11, pp. 961–986, 2004, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mar.20043>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mar.20043>
- [166] H. A. He, S. Greenberg, and E. M. Huang, “One size does not fit all: applying the transtheoretical model to energy feedback technology design,” in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ser. CHI ’10. New York, NY, USA: Association for Computing Machinery, Apr. 2010, pp. 927–936. [Online]. Available: <https://doi.org/10.1145/1753326.1753464>
- [167] J. Heer and G. Robertson, “Animated transitions in statistical data graphics,” IEEE TVCG, vol. 13, no. 6, pp. 1240–1247, 2007.
- [168] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell, “Passive real-world interface props for neurosurgical visualization,” in Proceedings of CHI, 1994.
- [169] K. Hirose, T. Ogawa, K. Kiyokawa, and H. Takemura, “Interactive Reconfiguration Techniques of Reference Frame Hierarchy in the Multi-viewport Interface,” in Proc. of 3DUI, 2006.
- [170] D. Holstius, J. Kembel, A. Hurst, P.-H. Wan, and J. Forlizzi, “Infotropism: living and robotic plants as interactive displays,” in Proc. DIS, 2004.
- [171] D. Honee, W. Hurst, and A. J. Luttikhoud, “Harnessing augmented reality for increasing the awareness of food waste amongst Dutch consumers,” Augmented Human Research, vol. 7, no. 1, p. 2, 2022.

BIBLIOGRAPHY

- [172] C. Hull, S. Knudsen, S. Carpendale, and W. Willett, “Simultaneous worlds: Supporting fluid exploration of multiple data sets via physical models,” in Proc. of GI, 2022.
- [173] J. Hullman, “Why authors don’t visualize uncertainty,” IEEE Transactions on Visualization and Computer Graphics, vol. 26, no. 1, pp. 130–139, 2020.
- [174] S. Huron, R. Vuillemot, and J.-D. Fekete, “Visual sedimentation,” IEEE TVCG, 2013.
- [175] A. Ilie, K.-L. Low, G. Welch, A. Lastra, H. Fuchs, and B. Cairns, “Combining Head-Mounted and Projector-Based Displays for Surgical Training,” Presence: Teleoperators and Virtual Environments, vol. 13, no. 2, 2004.
- [176] P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen, “Collaborative visualization: Definition, challenges, and research agenda,” Information Visualization, vol. 10, no. 4, pp. 310–326, 2011.
- [177] P. Isenberg, B. Lee, H. Qu, and M. Cordeil, “Immersive Visual Data Stories,” in Immersive Analytics, ser. Lecture Notes in Computer Science, K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, Eds. Cham: Springer International Publishing, 2018, pp. 165–184. [Online]. Available: https://doi.org/10.1007/978-3-030-01388-2_6
- [178] Y. Ishiguro and J. Rekimoto, “Peripheral vision annotation: noninterference information presentation method for mobile augmented reality,” in Proc. of AH, Tokyo, Japan, 2011.
- [179] Y. Itoh, J. Orlosky, K. Kiyokawa, and G. Klinker, “Laplacian vision: Augmenting motion prediction via optical see-through head-mounted displays,” in Proc. Augmented Human International Conference, 2016, pp. 1–8.
- [180] A. Ivanov, K. Danyluk, C. Jacob, and W. Willett, “A Walk Among the Data,” IEEE Computer Graphics and Applications, vol. 39, no. 3, pp. 19–28, May 2019, conference Name: IEEE Computer Graphics and Applications.
- [181] R. J. Jacob, A. Girouard, L. M. Hirshfield, M. S. Horn, O. Shaer, E. T. Solovey, and J. Zigelbaum, “Reality-based interaction: A framework for post-wimp interfaces,” in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ser. CHI ’08. New York, NY, USA: Association for Computing Machinery, 2008, p. 201–210. [Online]. Available: <https://doi.org/10.1145/1357054.1357089>
- [182] Y. Jansen, “Physical and tangible information visualization,” Theses, Université Paris Sud - Paris XI, Mar. 2014. [Online]. Available: <https://theses.hal.science/tel-00981521>
- [183] Y. Jansen, A. Assor, P. Dragicevic, A. Ferron, A. Prouzeau, M. Hachet, T. Le Gallic, I. Ajdukovic, E. Montagner, E. Spiegelman et al., “An interdisciplinary approach to designing and evaluating data-driven interactive experiences for sustainable decision-making,” in CHI’23 Workshop on HCI for Climate Change, 2023.
- [184] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, and K. Hornbæk, “Opportunities and challenges for data physicalization,” in Proc. of CHI, 2015.
- [185] R. Jianu, C. Demiralp, and D. Laidlaw, “Exploring 3D DTI fiber tracts with linked 2D representations,” IEEE TVCG, vol. 15, no. 6, pp. 1449–1456, 2009.
- [186] H. Jo, S. Hwang, H. Park, and J.-h. Ryu, “Aroundplot: Focus+context interface for off-screen objects in 3D environments,” Computers & Graphics, vol. 35, no. 4, 2011.
- [187] D. Kalkofen, M. Tatzgern, and D. Schmalstieg, “Explosion Diagrams in Augmented Reality,” in 2009 IEEE Virtual Reality Conference, Mar. 2009, pp. 71–78, iSSN: 2375-5334.
- [188] T. Kantonen, C. Woodward, and N. Katz, “Mixed reality in virtual world teleconferencing,” in Proc of IEEE VR, 2010.
- [189] R. K. Katschmann, B. Araki, and D. Rus, “Safe Local Navigation for Visually Impaired Users With a Time-of-Flight and Haptic Feedback Device,” IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 26, no. 3, pp. 583–593, Mar. 2018, conference Name: IEEE Transactions on Neural Systems and Rehabilitation Engineering.

BIBLIOGRAPHY

- [190] J. Keirstead, “Behavioural responses to photovoltaic systems in the uk domestic sector,” *Energy policy*, vol. 35, no. 8, pp. 4128–4141, 2007.
- [191] H. Kim, J. Kim, and S. Park, “A bird’s-eye view system using augmented reality,” in *Proceedings 32nd Annual Simulation Symposium*, 1999, pp. 126–131.
- [192] H. Kim, C. Coutrix, and A. Roudaut, “Knobslider: Design of a shape-changing ui for parameter control,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1–13. [Online]. Available: <https://doi.org/10.1145/3173574.3173913>
- [193] Y. J. Kim, D. Njite, and M. Hancer, “Anticipated emotion in consumers’ intentions to select eco-friendly restaurants: Augmenting the theory of planned behavior,” *International Journal of Hospitality Management*, vol. 34, pp. 255–262, Sep. 2013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0278431913000510>
- [194] K. N. Kirby and D. Gerlanc, “BootES: An R package for bootstrap confidence intervals on effect sizes,” *Behavior Research Methods*, vol. 45, no. 4, pp. 905–927, Dec. 2013. [Online]. Available: <https://doi.org/10.3758/s13428-013-0330-5>
- [195] S. Kitagawa, S. Hasegawa, N. Yamaguchi, K. Okada, and M. Inaba, “Miniature Tangible Cube: Concept and Design of Target-Object-Oriented User Interface for Dual-Arm Telemanipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6977–6984, Oct. 2021, conference Name: IEEE Robotics and Automation Letters.
- [196] B. Kondo and C. Collins, “DimpVis: Exploring time-varying information visualizations by direct manipulation,” *IEEE TVCG*, vol. 20, no. 12, pp. 2003–2012, 2014.
- [197] A. M. Kring and A. H. Gordon, “Sex differences in emotion: expression, experience, and physiology,” *Journal of personality and social psychology*, vol. 74, no. 3, p. 686, 1998.
- [198] A. Kristoffersson, S. Coradeschi, and A. Loutfi, “A Review of Mobile Robotic Telepresence,” *Advances in Human-Computer Interaction*, 2013.
- [199] M. Kritzler, M. Funk, F. Michahelles, and W. Rohde, “The virtual twin: controlling smart factories using a spatially-correct augmented reality representation,” in *Proc. of IoT*, 2017.
- [200] H. Laaki, Y. Miche, and K. Tammi, “Prototyping a Digital Twin for Real Time Remote Control Over Mobile Networks: Application of Remote Surgery,” *IEEE Access*, vol. 7, 2019, conference Name: IEEE Access.
- [201] B. Lee, D. Brown, B. Lee, C. Hurter, S. Drucker, and T. Dwyer, “Data Visceralization: Enabling Deeper Understanding of Data Using Virtual Reality,” *IEEE TVCG*, pp. 1–1, 2020, arXiv: 2009.00059. [Online]. Available: <http://arxiv.org/abs/2009.00059>
- [202] B. Lee, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer, “A design space for data visualisation transformations between 2D and 3D in mixed-reality environments,” in *ACM CHI*, 2022. [Online]. Available: <https://doi.org/10.1145/3491102.3501859>
- [203] B. Lee, A. Satyanarayan, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer, “Deimos: A grammar of dynamic embodied immersive visualisation morphs and transitions,” in *ACM CHI*, 2023, pp. 1–18.
- [204] D. K. Lee, “Alternatives to p value: confidence interval and effect size,” *Korean Journal of Anesthesiology*, vol. 69, no. 6, pp. 555–562, 2016.
- [205] H. Lee, S.-T. Noh, and W. Woo, “TunnelSlice: Freehand Subspace Acquisition Using an Egocentric Tunnel for Wearable Augmented Reality,” *IEEE Transactions on Human-Machine Systems*, vol. 47, no. 1, pp. 128–139, Feb. 2017, conference Name: IEEE Transactions on Human-Machine Systems.
- [206] J. Li, Q. Smithwick, and D. Chu, “Holobricks: modular coarse integral holographic displays,” *Light: Science & Applications*, vol. 11, no. 1, pp. 1–15, 2022.
- [207] N. Li, Z. Zhang, C. Liu, Z. Yang, Y. Fu, F. Tian, T. Han, and M. Fan, “vMirror: Enhancing the Interaction with Occluded or Distant Objects in VR with Virtual Mirrors,” in *Proc. of CHI*, 2021.

BIBLIOGRAPHY

- [208] W. Li, Y. Wang, H. Huang, W. Cui, H. Zhang, H. Qu, and D. Zhang, “AniVis: Generating animated transitions between statistical charts with a tree model,” 2021.
- [209] K. Liliya, H. Pohl, S. Boring, and K. Hornbæk, “Augmented Reality Views for Occluded Interaction,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Glasgow Scotland Uk: ACM, May 2019, pp. 1–12. [Online]. Available: <https://dl.acm.org/doi/10.1145/3290605.3300676>
- [210] T. Lin, Y. Yang, J. Beyer, and H. Pfister, “Labeling out-of-view objects in immersive analytics to support situated visual searching,” *IEEE TVCG*, vol. 29, 2021.
- [211] Y.-T. Lin, Y.-C. Liao, S.-Y. Teng, Y.-J. Chung, L. Chan, and B.-Y. Chen, “Outside-In: Visualizing Out-of-Sight Regions-of-Interest in a 360° Video Using Spatial Picture-in-Picture Previews,” in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, ser. *UIST '17*. New York, NY, USA: Association for Computing Machinery, Oct. 2017, pp. 255–265. [Online]. Available: <https://doi.org/10.1145/3126594.3126656>
- [212] R. Lindgren, M. Tscholl, S. Wang, and E. Johnson, “Enhancing learning and engagement through embodied interaction within a mixed reality simulation,” *Computers & Education*, vol. 95, pp. 174–187, Apr. 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S036013151630001X>
- [213] M. Livingston, J. Swan, J. Gabbard, T. Hollerer, D. Hix, S. Julier, Y. Baillet, and D. Brown, “Resolving multiple occluded layers in augmented reality,” in *Proc. of ISMAR*, 2003.
- [214] M. A. Livingston, A. Dey, C. Sandor, and B. H. Thomas, *Pursuit of “X-Ray Vision” for Augmented Reality*, 2013, ch. Human Factors in Augmented Reality Environments.
- [215] P. Lubos, G. Bruder, O. Ariza, and F. Steinicke, “Ambiculus: LED-based low-resolution peripheral display extension for immersive head-mounted displays,” in *Proc. of VRIC*, 2016.
- [216] A. Maimone, X. Yang, N. Dierk, A. State, M. Dou, and H. Fuchs, “General-purpose telepresence with head-worn optical see-through displays and projector-based lighting,” in *Proc. of IEEE VR*, 2013.
- [217] B. Marques, B. S. Santos, T. Araújo, N. C. Martins, J. B. Alves, and P. Dias, “Situated visualization in the decision process through augmented reality,” in *Proc. Int. Conf. Information Visualisation (IV)*, 2019, pp. 13–18.
- [218] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, *Immersive analytics*. Springer, 2018, vol. 11190.
- [219] V. Masson-Delmotte, P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P. R. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J. B. R. Matthews, Y. Chen, X. Zhou, M. I. Gomis, E. Lonnoy, T. Maycock, M. Tignor, and T. Waterfield, “Global warming of 1.5° c: an ipcc special report on the impacts of global warming of 1.5° c above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty,” *Intergovernmental Panel on Climate Change (IPCC)*, Tech. Rep., 2018.
- [220] E. P. McCormick, C. D. Wickens, R. Banks, and M. Yeh, “Frame of reference effects on scientific visualization subtasks,” *Human Factors*, vol. 40, no. 3, pp. 443–451, 1998.
- [221] M. J. McGuffin, G. Davison, and R. Balakrishnan, “Expand-ahead: a space-filling strategy for browsing trees,” in *IEEE Symp. Information Visualization (InfoVis)*, 2004, pp. 119–126.
- [222] J. Mercier-Ganady, F. Lotte, E. Loup-Escande, M. Marchal, and A. Lécuyer, “The Mind-Mirror: See your brain in action in your head using EEG and augmented reality,” in *IEEE VR*, 2014.
- [223] H. S. Meyerhoff, F. Papenmeier, and M. Huff, “Studying visual attention using the multiple object tracking paradigm: A tutorial review,” *Attention, Perception, & Psychophysics*, vol. 79, pp. 1255–1274, 2017.
- [224] P. Milgram and F. Kishino, “A Taxonomy of Mixed Reality Visual Displays,” *IEICE Trans. Information Systems*, 1994.

BIBLIOGRAPHY

- [225] K. J. Minnaugh, M. Suomalainen, I. Becerra, E. Lozano, R. Murrieta-Cid, and S. M. LaValle, “Defining Preferred and Natural Robot Motions in Immersive Telepresence from a First-Person Perspective,” [arXiv:2102.12719 \[cs\]](https://arxiv.org/abs/2102.12719), Feb. 2021, arXiv: 2102.12719. [Online]. Available: <http://arxiv.org/abs/2102.12719>
- [226] R. Minerva, G. M. Lee, and N. Crespi, “Digital Twin in the IoT Context: A Survey on Technical Features, Scenarios, and Architectural Models,” Proceedings of the IEEE, vol. 108, no. 10, pp. 1785–1824, Oct. 2020, conference Name: Proceedings of the IEEE.
- [227] D. Mirk and H. Hlavacs, “Virtual Tourism with Drones: Experiments and Lag Compensation,” in Proceedings of Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use. Florence Italy: ACM, May 2015, pp. 45–50. [Online]. Available: <https://dl.acm.org/doi/10.1145/2750675.2750681>
- [228] F. Mizuno, T. Hayasaka, and T. Yamaguchi, “Virtual chameleon- a system to provide different views to both eyes,” in World Congress on Medical Physics and Biomedical Engineering - Information and Communication in Medicine, Telemedicine and e-Health, no. 5, 2009.
- [229] L. Morais, Y. Jansen, N. Andrade, and P. Dragicevic, “Showing data about people: A design space of anthropographics,” IEEE TVCG, vol. 28, no. 3, pp. 1661–1679, 2020.
- [230] S. Mori, S. Ikeda, and H. Saito, “A survey of diminished reality: Techniques for visually concealing, eliminating, and seeing through real objects,” IPSN Transactions on Computer Vision and Applications, vol. 9, no. 1, 2017.
- [231] T. Munzner, Visualization analysis and design, ser. A.K. Peters visualization series. Boca Raton: CRC Press, Taylor & Francis Group, CRC Press is an imprint of the Taylor & Francis Group, an informa business, 2015.
- [232] S. Murata, S. Kagatsume, H. Taguchi, and K. Fujinami, “PerFridge: An Augmented Refrigerator That Detects and Presents Wasteful Usage for Eco-Persuasion,” in 2012 IEEE 15th International Conference on Computational Science and Engineering, Dec. 2012, pp. 367–374.
- [233] H. R. Nagel, E. Granum, S. Bovbjerg, and M. Vittrup, “Immersive Visual Data Mining: The 3DVDM Approach,” in Visual Data Mining: Theory, Techniques and Tools for Visual Analytics, ser. Lecture Notes in Computer Science, S. J. Simoff, M. H. Böhlen, and A. Mazeika, Eds. Berlin, Heidelberg: Springer, 2008, pp. 281–311. [Online]. Available: https://doi.org/10.1007/978-3-540-71080-6_18
- [234] T. Nakao, M. Nakatani, L. Chan, and K. Kunze, “Smart glasses with a peripheral vision display,” in Proc. of VRIC, 2016.
- [235] J. W. Nam, K. McCullough, J. Tveite, M. M. Espinosa, C. H. Perry, B. T. Wilson, and D. F. Keefe, “Worlds-in-Wedges: Combining Worlds-in-Miniature and Portals to Support Comparative Immersive Visualization of Forestry Data,” in 2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR), Mar. 2019, pp. 747–755, iSSN: 2642-5254.
- [236] A. Nassani, G. Lee, M. Billingham, T. Langlotz, and R. W. Lindeman, “Using visual and spatial cues to represent social contacts in AR,” in SIGGRAPH Asia 2017 Mobile Graphics & Interactive Applications on - SA '17. Bangkok, Thailand: ACM Press, 2017, pp. 1–6. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3132787.3139199>
- [237] M. Neurath, “Isotype,” Instructional science, pp. 127–150, 1974.
- [238] A. L. Nichols and J. K. Maner, “The good-subject effect: Investigating participant demand characteristics,” The Journal of general psychology, vol. 135, no. 2, pp. 151–166, 2008.
- [239] A. Nieman, “Concrete vs abstract visualisation: the real world as a canvas for data visualisation,” in Proc. ADS-VIS2011: Making visible the invisible: Art, design and science in data visualisation, 2012, pp. 49–56, <http://eprints.hud.ac.uk/id/eprint/12775/>.
- [240] D. Offenhuber, “Data by proxy—material traces as autographic visualizations,” IEEE transactions on visualization and computer graphics, vol. 26, no. 1, pp. 98–108, 2019.
- [241] A. Olwal and B. Kress, “1D eyewear: peripheral, hidden LEDs and near-eye holographic displays for unobtrusive augmentation,” in Proc. of ISWC, 2018.

BIBLIOGRAPHY

- [242] I. P. on Climate Change, Global warming of 1.5° C: an IPCC special report on the impacts of global warming of 1.5° C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. Intergovernmental Panel on Climate Change, 2018.
- [243] J. Orlosky, “AR as an Entry Point to Parallel Consciousness,” Mar. 2017.
- [244] A. V. Pandey, A. Manivannan, O. Nov, M. Satterthwaite, and E. Bertini, “The persuasive power of data visualization,” IEEE Transactions on Visualization and Computer Graphics, vol. 20, no. 12, pp. 2211–2220, 2014.
- [245] D. Park, S. M. Drucker, R. Fernandez, and N. Elmqvist, “Atom: A grammar for unit visualizations,” IEEE TVCG, vol. 24, no. 12, pp. 3032–3043, 2018.
- [246] E. Paulos and T. Jenkins, “Jetsam: Exposing our everyday discarded objects,” 2006.
- [247] R. Pausch, T. Burnette, D. Brockway, and M. E. Weiblen, “Navigation and locomotion in virtual worlds via flight into hand-held miniatures,” in Proc. of SIGGRAPH, 1995.
- [248] L. J. Perovich, S. A. Wylie, and R. Bongiovanni, “Chemicals in the creek: designing a situated data physicalization of open government data with the community,” IEEE Transactions on Visualization and Computer Graphics, vol. 27, no. 2, pp. 913–923, 2020.
- [249] P. Petkov, S. Goswami, F. Köbler, and H. Krcmar, “Personalised eco-feedback as a design technique for motivating energy saving behaviour at home,” in Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design, 2012, pp. 587–596.
- [250] J. Pierce, W. Odom, and E. Blevis, “Energy aware dwelling: a critical survey of interaction design for eco-visualizations,” in Proceedings of the 20th Australasian Conference on Computer-Human Interaction: Designing for Habitus and Habitat, 2008, pp. 1–8.
- [251] C. Plaisant, “The challenge of information visualization evaluation,” in Proceedings of the Working Conference on Advanced Visual Interfaces, ser. AVI '04. New York, NY, USA: Association for Computing Machinery, 2004, p. 109–116. [Online]. Available: <https://doi.org/10.1145/989863.989880>
- [252] C. Plaisant, J. Grosjean, and B. B. Bederson, “SpaceTree: Supporting exploration in large node link tree, design evolution and empirical evaluation,” in IEEE Symp. Information Visualization (InfoVis), 2002, pp. 57–64.
- [253] Presseliteapps, “Twitter 360 for the iPhone 3GS with Augmented Reality.” [Online]. Available: <https://www.youtube.com/watch?v=5w7EAz8-uwU>
- [254] A. Prouzeau, M. Cordeil, C. Robin, B. Ens, B. H. Thomas, and T. Dwyer, “Scaptics and Highlight-Planes: Immersive Interaction Techniques for Finding Occluded Features in 3D Scatterplots,” in Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, ser. CHI '19. New York, NY, USA: Association for Computing Machinery, May 2019, pp. 1–12. [Online]. Available: <https://doi.org/10.1145/3290605.3300555>
- [255] A. Prouzeau, A. Lhuillier, B. Ens, D. Weiskopf, and T. Dwyer, “Visual link routing in immersive visualisations,” in Proceedings of the 2019 ACM international conference on interactive surfaces and spaces, 2019, pp. 241–253.
- [256] A. Prouzeau, Y. Wang, B. Ens, W. Willett, and T. Dwyer, “Corsican Twin: Authoring In Situ Augmented Reality Visualisations in Virtual Reality,” in Proc. of AVI, 2020.
- [257] L. Qian, A. Plopski, N. Navab, and P. Kazanzides, “Restoring the Awareness in the Occluded Visual Field for Optical See-Through Head-Mounted Displays,” IEEE Transactions on Visualization and Computer Graphics, vol. 24, no. 11, pp. 2936–2946, Nov. 2018, conference Name: IEEE Transactions on Visualization and Computer Graphics.
- [258] C. Qiu, S. Zhou, Z. Liu, Q. Gao, and J. Tan, “Digital assembly technology based on augmented reality and digital twins: a review,” Virtual Reality & Intelligent Hardware, vol. 1, no. 6, 2019.
- [259] R. Reddy and D. Rompapas, VisuoTouch: Enabling Haptic Feedback in Augmented Reality through Visual Cues, Dec. 2020.

BIBLIOGRAPHY

- [260] H. Regenbrecht, M. Wagner, and G. Baratoff, “MagicMeeting: A Collaborative Tangible Augmented Reality System,” *Virtual Reality*, vol. 6, no. 3, pp. 151–166, Oct. 2002. [Online]. Available: <https://doi.org/10.1007/s100550200016>
- [261] D. Ren, T. Goldschwendt, Y. Chang, and T. Höllerer, “Evaluating wide-field-of-view augmented reality with mixed reality simulation,” in *Proc. of IEEE VR*, 2016.
- [262] D. Ren, B. Lee, and T. Höllerer, “XRCreator: interactive construction of immersive data-driven stories,” in *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology*, ser. VRST ’18. New York, NY, USA: Association for Computing Machinery, Nov. 2018, pp. 1–2. [Online]. Available: <https://doi.org/10.1145/3281505.3283400>
- [263] J. S. Roo, R. Gervais, J. Frey, and M. Hachet, “Inner garden: Connecting inner states to a mixed reality sandbox for mindfulness,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’17. New York, NY, USA: Association for Computing Machinery, 2017, p. 1459–1470. [Online]. Available: <https://doi.org/10.1145/3025453.3025743>
- [264] C. Sandor, A. Cunningham, A. Dey, and V.-V. Mattila, “An Augmented Reality X-Ray system based on visual saliency,” in *2010 IEEE International Symposium on Mixed and Augmented Reality*, Oct. 2010, pp. 27–36.
- [265] C. Sandor, A. Dey, A. Cunningham, S. Barbier, U. Eck, D. Urquhart, M. R. Marner, G. Jarvis, and S. Rhee, “Egocentric space-distorting visualizations for rapid environment exploration in mobile mixed reality,” in *2010 IEEE Virtual Reality Conference (VR)*, Mar. 2010, pp. 47–50, iSSN: 2375-5334.
- [266] B. Santos, T. Romão, A. E. Dias, P. Centieiro, and B. Teixeira, “Changing environmental behaviors through smartphone-based augmented experiences,” in *International Conference on Advances in Computer Entertainment Technology*. Springer, 2012, pp. 553–556.
- [267] K. Satriadi, A. Cunningham, R. Smith, T. Dwyer, A. Drogemuller, and B. Thomas, “Proxiated visualization: An extended model of situated visualization using proxies for physical referents,” in *Proc. of CHI*, 2023.
- [268] K. A. Satriadi, A. Cunningham, B. H. Thomas, A. Drogemuller, A. Odi, N. Patel, C. Aston, and R. T. Smith, “Augmented scale models: Presenting multivariate data around physical scale models in augmented reality,” in *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2022, pp. 54–63.
- [269] K. A. Satriadi, J. Smiley, B. Ens, M. Cordeil, T. Czauderna, B. Lee, Y. Yang, T. Dwyer, and B. Jenny, “Tangible globes for data visualisation in augmented reality,” in *Proc. of CHI*, 2022.
- [270] K. Sauvé, S. Bakker, and S. Houben, “Econundrum: Visualizing the climate impact of dietary choice through a shared data sculpture,” in *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, 2020, pp. 1287–1300.
- [271] S. Scarr and M. Hernandez, “Drowning in plastic – visualising the world’s addiction to plastic bottles,” Web article <https://graphics.reuters.com/ENVIRONMENT-PLASTIC/0100B275155/index.html>. Last visited 2022-04-05/, 2019.
- [272] T. Schinke, N. Henze, and S. Boll, “Visualization of off-screen objects in mobile augmented reality,” in *Proceedings of the 12th international MobileHCI ’10*. Lisbon, Portugal: ACM Press, 2010, p. 313. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1851600.1851655>
- [273] D. Schwajda, J. Friedl, F. Pointecker, H.-C. Jetter, and C. Anthes, “Transforming graph data visualisations from 2D displays into augmented reality 3D space: A quantitative study,” *Frontiers in Virtual Reality*, vol. 4, 2023.
- [274] D. W. Seo, H. Kim, J. S. Kim, and J. Y. Lee, “Hybrid reality-based user experience and evaluation of a context-aware smart home,” *Computers in Industry*, vol. 76, 2016.
- [275] B. Shelton and N. Hedley, “Exploring a Cognitive Basis for Learning Spatial Relationships with Augmented Reality,” *ITLS Faculty Publications*, vol. 1, Jan. 2003.
- [276] W. R. Sherman and A. B. Craig, “Understanding Virtual Reality—Interface, Application, and Design,” *Presence: Teleoperators & Virtual Environments*, vol. 12, no. 4, 2003.

BIBLIOGRAPHY

- [277] R. Sicat, J. Li, J. Choi, M. Cordeil, W.-K. Jeong, B. Bach, and H. Pfister, “DXR: A toolkit for building immersive data visualizations,” *IEEE TVCG*, vol. 25, no. 1, pp. 715–725, 2018.
- [278] T. Siu and V. Herskovic, “SidebARs: improving awareness of off-screen elements in mobile augmented reality,” in *Proc. of ChileCHI*, Temuco, Chile, 2013.
- [279] R. Skarbez, N. F. Polys, J. T. Ogle, C. North, and D. A. Bowman, “Immersive analytics: Theory and research agenda,” *Frontiers in Robotics and AI*, vol. 6, 2019.
- [280] T. Skog, “Activity wallpaper: ambient visualization of activity information,” Jan. 2004, pp. 325–328.
- [281] P. Soltani and A. H. P. Morice, “Augmented reality tools for sports education and training,” *Computers & Education*, vol. 155, p. 103923, Oct. 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0360131520301226>
- [282] J. Sorger, P. Mindek, P. Rautek, E. Gröller, G. Johnson, and I. Viola, “Metamorphers: Storytelling templates for illustrative animated transitions in molecular visualization,” in *Proc. Spring Conference on Computer Graphics*, 2017.
- [283] M. Sousa, D. Mendes, S. Paulo, N. Matela, J. Jorge, and D. S. Lopes, “VRRRRRoom: Virtual Reality for Radiologists in the Reading Room,” in *Proc. of CHI*, Denver Colorado USA, 2017.
- [284] S. Y. Ssin, J. A. Walsh, R. T. Smith, A. Cunningham, and B. H. Thomas, “GeoGate: Correlating Geo-Temporal Datasets Using an Augmented Reality Space-Time Cube and Tangible Interactions,” in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, Mar. 2019, pp. 210–219, iSSN: 2642-5254.
- [285] N. Stark, “Visualizing impact through augmented reality – using concrete scales to combat waste,” Web page <https://nicstark.com/visualizing-impact-with-augmented-reality/>. Last visited 2022-04-05/, 2020.
- [286] R. Stoakley, M. J. Conway, and R. Pausch, “Virtual reality on a WIM: interactive worlds in miniature,” in *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '95*. Denver, Colorado, United States: ACM Press, 1995, pp. 265–272. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=223904.223938>
- [287] Y. A. Strengers, “Designing eco-feedback systems for everyday life,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '11. New York, NY, USA: Association for Computing Machinery, May 2011, pp. 2135–2144. [Online]. Available: <https://doi.org/10.1145/1978942.1979252>
- [288] M. Sukan, C. Elvezio, O. Oda, S. Feiner, and B. Tversky, “ParaFrustum: visualization techniques for guiding a user to a constrained set of viewing positions and orientations,” in *Proc. of UIST*, 2014.
- [289] S. Tachi, N. Kawakami, H. Nii, K. Watanabe, and K. Minamizawa, “TELEsarPHONE: Mutual Telexistence Master-Slave Communication System Based on Retroreflective Projection Technology,” *SICE Journal of Control, Measurement, and System Integration*, vol. 1, no. 5, pp. 335–344, Sep. 2008. [Online]. Available: <https://www.tandfonline.com/doi/full/10.9746/jcmsi.1.335>
- [290] S. Tachi, K. Komoriya, K. Sawada, T. Nishiyama, T. Itoko, M. Kobayashi, and K. Inoue, “Telexistence cockpit for humanoid robot control,” *Advanced Robotics*, vol. 17, no. 3, 2003.
- [291] S. Tachi, K. Watanabe, K. Takeshita, K. Minamizawa, T. Yoshida, and K. Sato, “Mutual telexistence surrogate system: ESAR4 - telexistence in real environments using autostereoscopic immersive display -,” in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sep. 2011, pp. 157–162, iSSN: 2153-0866.
- [292] Y. Tajima and T. Nagatani, “Scaling behavior of crowd flow outside a hall,” *Physica A: Statistical Mechanics and its Applications*, vol. 292, no. 1-4, pp. 545–554, 2001.
- [293] F. Tao, M. Zhang, and A. Y. C. Nee, “Chapter 11 - Digital Twin and Virtual Reality and Augmented Reality/Mixed Reality,” in *Digital Twin Driven Smart Manufacturing*, 2019.

BIBLIOGRAPHY

- [294] M. Tatzgern, R. Grasset, E. Veas, D. Kalkofen, H. Seichter, and D. Schmalstieg, “Exploring real world points of interest: Design and evaluation of object-centric exploration techniques for augmented reality,” *Pervasive and Mobile Computing*, vol. 18, 2015.
- [295] B. Thomas, G. Welch, P. Dragicevic, N. Elmqvist, P. Irani, Y. Jansen, D. Schmalstieg, A. Tabard, N. Elsayed, R. Smith, and W. Willett, “Situating Analytics,” in *Immersive Analytics*, ser. Lecture Notes in Computer Science. Springer, Oct. 2018, vol. 11190, pp. 185–220. [Online]. Available: <https://inria.hal.science/hal-01947243>
- [296] B. H. Thomas, G. F. Welch, P. Dragicevic, N. Elmqvist, P. Irani, Y. Jansen, D. Schmalstieg, A. Tabard, N. A. M. ElSayed, R. T. Smith, and W. Willett, “Situating Analytics,” in *Immersive Analytics*, 2018, vol. 11190.
- [297] B. Thoravi Kumaravel, F. Anderson, G. Fitzmaurice, B. Hartmann, and T. Grossman, “Loki: Facilitating Remote Instruction of Physical Tasks Using Bi-Directional Mixed-Reality Telepresence,” in *Proc. of UIST*, 2019.
- [298] timetraveler augmented, “Timetraveler Augmented The Berlin Wall App - English Voice Over,” Sep. 2014. [Online]. Available: <https://www.youtube.com/watch?v=HlbA7G8V3jA>
- [299] X. Tong, C. Li, and H.-W. Shen, “GlyphLens: View-Dependent Occlusion Management in the Interactive Glyph Visualization,” *IEEE TVCG*, vol. 23, no. 1, 2017.
- [300] A. Treuille, S. Cooper, and Z. Popović, “Continuum crowds,” *ACM Trans. Graph.*, vol. 25, no. 3, p. 1160–1168, jul 2006. [Online]. Available: <https://doi.org/10.1145/1141911.1142008>
- [301] R. Trueba, C. Andujar, and F. Argelaguet, “Complexity and Occlusion Management for the World-in-Miniature Metaphor,” in *Smart Graphics*, ser. Lecture Notes in Computer Science, A. Butz, B. Fisher, M. Christie, A. Krüger, P. Olivier, and R. Therón, Eds. Berlin, Heidelberg: Springer, 2009, pp. 155–166.
- [302] S. Truman and S. von Mammen, “An Integrated Design of World-in-Miniature Navigation in Virtual Reality,” in *Proc. of FDG*, 2020, no. 69.
- [303] E. R. Tufte, “The visual display of quantitative information,” *The Journal for Healthcare Quality (JHQ)*, vol. 7, no. 3, p. 15, 1985.
- [304] F. Tyndiuk, “Référentiels Spatiaux des Tâches d’Interaction et Caractéristiques de l’Utilisateur influençant la Performance en Réalité Virtuelle,” p. 235.
- [305] H. van der Valk, H. Haße, F. Möller, M. Arbter, J.-L. Henning, and B. Otto, *A Taxonomy of Digital Twins*, Aug. 2020.
- [306] E. Veas, R. Grasset, E. Kruijff, and D. Schmalstieg, “Extended Overview Techniques for Outdoor Augmented Reality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 4, pp. 565–572, Apr. 2012, conference Name: *IEEE Transactions on Visualization and Computer Graphics*.
- [307] J. Vermeulen, F. Kawsar, A. L. Simeone, G. Kortuem, K. Luyten, and K. Coninx, “Informing the design of situated glyphs for a care facility,” in *Proc. of VL/HCC*, 2012.
- [308] A. Visuri, N. van Berkel, C. Luo, J. Goncalves, D. Ferreira, and V. Kostakos, “Challenges of quantified-self: encouraging self-reported data logging during recurrent smartphone usage,” in *Proceedings of the 31st International BCS Human Computer Interaction Conference (HCI 2017)* 31, 2017, pp. 1–7.
- [309] F. Volonté, F. Pugin, P. Bucher, M. Sugimoto, O. Ratib, and P. Morel, “Augmented reality and image overlay navigation with OsiriX in laparoscopic and robotic surgery: not only a matter of fashion,” *Journal of Hepato-Biliary-Pancreatic Sciences*, vol. 18, no. 4, 2011.
- [310] J. Wagner, W. Stuerzlinger, and L. Nedel, “The effect of exploration mode and frame of reference in immersive analytics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 9, pp. 3252–3264, 2022.
- [311] L. Wang, J. Chen, Q. Ma, and V. Popescu, “Disocclusion Headlight for Selection Assistance in VR,” in *Proc of IEEE VR*, 2021.

BIBLIOGRAPHY

- [312] X. Wang, P. E. D. Love, M. J. Kim, and W. Wang, "Mutual awareness in collaborative design: An Augmented Reality integrated telepresence system," *Computers in Industry*, vol. 65, no. 2, 2014.
- [313] D. Watson, L. A. Clark, and A. Tellegen, "Development and validation of brief measures of positive and negative affect: The PANAS scales," *Journal of Personality and Social Psychology*, vol. 54, no. 6, pp. 1063–1070, 1988, place: US Publisher: American Psychological Association.
- [314] H. Wei-Che, C.-M. Tseng, and K. Shih-Chung, "Using Exaggerated Feedback in a Virtual Reality Environment to Enhance Behavior Intention of Water-Conservation," *Journal of Educational Technology & Society*, vol. 21, no. 4, pp. 187–203, 2018, num Pages: 187-203 Place: Palmerston North, Canada Publisher: International Forum of Educational Technology & Society Section: Special Issue Articles. [Online]. Available: <https://www.proquest.com/docview/2147876359/abstract/2861DD1A79B54748PQ/1>
- [315] G. Westheimer, "Scaling of Visual Acuity Measurements," *Archives of Ophthalmology*, vol. 97, no. 2, pp. 327–330, Feb. 1979. [Online]. Available: <https://doi.org/10.1001/archopht.1979.01020010173020>
- [316] S. White, S. Feiner, and J. Kopylec, "Virtual Vouchers: Prototyping a Mobile Augmented Reality User Interface for Botanical Species Identification," in *Proc. of 3DUI*, 2006.
- [317] S. White and S. Feiner, "SiteLens: situated visualization techniques for urban site visits," in *Proc. of CHI*, 2009.
- [318] M. Whitlock, D. A. Szafir, and K. Gruchalla, "Hydrogenar: Interactive data-driven presentation of dispenser reliability," in *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2020, pp. 704–712.
- [319] C. D. Wickens and T. T. Prevett, "Exploring the dimensions of egocentricity in aircraft navigation displays." *Journal of Experimental Psychology: Applied*, vol. 1, no. 2, p. 110, 1995.
- [320] W. Willett, "Cetonia: A dynamic swarm at your fingertips." [Online]. Available: <http://www.wjwillett.net/content/cetonia/>
- [321] W. Willett, B. A. Aseniero, S. Carpendale, P. Dragicevic, Y. Jansen, L. Oehlberg, and P. Isenberg, "Superpowers as inspiration for visualization," *IEEE TVCG*, vol. 2021, 2021.
- [322] W. Willett, Y. Jansen, and P. Dragicevic, "Embedded Data Representations," *IEEE Trans. Visual. Comput. Graphics*, vol. 23, no. 1, 2017.
- [323] M. Wilson, M. Wilson, and M. Wilson, "How GM Is Saving Cash Using Legos As A Data Viz Tool," Apr. 2012. [Online]. Available: <https://www.fastcompany.com/1669468/how-gm-is-saving-cash-using-legos-as-a-data-viz-tool>
- [324] J. Wither, S. DiVerdi, and T. Höllerer, "Annotation in outdoor augmented reality," *Computers & Graphics*, vol. 33, no. 6, 2009.
- [325] M.-L. Wu and V. Popescu, "Efficient VR and AR Navigation Through Multiperspective Occlusion Management," *IEEE TVCG*, vol. 24, no. 12, 2018.
- [326] Y. Yang, T. Dwyer, K. Marriott, B. Jenny, and S. Goodwin, "Tilt map: Interactive transitions between choropleth map, prism map and bar chart in immersive environments," *IEEE TVCG*, vol. 27, no. 12, pp. 4507–4519, 2020.
- [327] E. Yigitbas, K. Karakaya, I. Jovanovikj, and G. Engels, "Enhancing Human-in-the-Loop Adaptive Systems through Digital Twins and VR Interfaces," *arXiv:2103.10804 [cs]*, Mar. 2021, arXiv: 2103.10804. [Online]. Available: <http://arxiv.org/abs/2103.10804>
- [328] D. Yu, H.-N. Liang, K. Fan, H. Zhang, C. Fleming, and K. Papangelis, "Design and Evaluation of Visualization Techniques of Off-Screen and Occluded Targets in Virtual Reality Environments," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 9, pp. 2762–2774, Sep. 2020, conference Name: IEEE Transactions on Visualization and Computer Graphics.
- [329] J. Yu, E. Williams, and M. Ju, "Analysis of material and energy consumption of mobile phones in china," *Energy Policy*, vol. 38, no. 8, pp. 4135–4141, 2010. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0301421510002041>

BIBLIOGRAPHY

- [330] D. A. Yudkin, R. Pick, E. Y. Hur, N. Liberman, and Y. Trope, “Psychological distance promotes exploration in search of a global maximum,” Personality and Social Psychology Bulletin, vol. 45, no. 6, pp. 893–906, 2019.
- [331] J. Zhao and A. V. Moere, “Embodiment in data sculpture: a model of the physical visualization of information,” in Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts, 2008, pp. 343–350.
- [332] X. Zhou, Y. Yang, F. Ortega, A. U. Batmaz, and B. Lee, “Data-driven storytelling in hybrid immersive display environments,” 2023.
- [333] S. Zollmann, D. Kalkofen, C. Hoppe, S. Kluckner, H. Bischof, and G. Reitmayr, “Interactive 4D overview and detail visualization in augmented reality,” in Proc. of ISMAR, 2012.
- [334] S. Zollmann, C. Poglitsch, and J. Ventura, “VISGIS: Dynamic situated visualization for geographic information systems,” in Proc. of IVCNZ, 2016.

Appendices

A.1 Glossary

I define here frequently used words in the thesis. These definitions are fully or partially inspired by the literature put in reference. Italic signifies an unmodified quote.

- **Representation:** Refers to how something is presented. Visual representations rely on graphical elements to convey ideas, phenomena or data.
- **Visualization:** A visualization is a form of representation that focuses on presenting information in an understandable and clear way. These aim at helping people to complete tasks more effectively [231].
- **Situated data representation:** *Data representation located close to the data's **physical referent**, meaning close to the space, object, or person the data is about* [322].
- **Embedded data representation:** A situated data representation that involves multiple referents, enabling *multiple physical presentations that each independently display data related to their respective physical referent*. [322]
- **Visceral representation:** A representation that restores the basic understanding of units and measures that are often abstracted away [201].
- **Abstract visualization:** Conventional data visualization, where data values are mapped to elementary visual variables such as position, size, or color [231].
- **Concrete visualization:** Visualizations conveying data using representations that resemble objects of the physical world.
- **AR:** Augmented Reality, technology that enables the integration of 3D virtual objects into a 3D real environment in real time [44].

- **VR:** *VR simulates a virtual environment that immerses users to the extent that they have the feeling of “being there” [78]. A virtual environment is a software representation of real (or imagined) agents, objects, and processes; and a human–computer interface for displaying and interacting with these models [52].*
- **MR:** A Mixed Reality environment is an environment in which the real and the virtual world are displayed through the same display [224].
- **HMD:** Head-Mounted Displays, displays that are worn on the head and typically cover the user’s eyes.
- **Embodied interaction:** *“Interactions with software systems emphasizing skilled, engaged practice rather than disembodied rationality”. [113]*
- **Reality-based interaction:** Interaction with digital content using intuitive gestures, voice commands, or even direct manipulation of virtual objects within their physical environment. [181]
- **ARWAVs:** Augmented Reality Waste Accumulation Visualizations. Those are holographic, literal, concrete visualizations of waste data embedded in the user’s surroundings.

A.2 Supplementary material chapter 3

B Experiment questionnaire

APPENDIX A. APPENDICES

Inria

Ordering :
Participant Id:

The employees of the Inria research laboratory produce waste within the centre. This waste includes household waste (partly produced during meals, here in the centre's cafeteria), cardboard, recyclable waste, paper and glass. This volume of waste is collected by General Services when the waste collection and transportation company employed by the centre comes by.

1. How often do you use the center's cafeteria ? *(Check the box)*

Never	Less than one time per month	Several times per month	Between one and two time a week	Three times per week or more

2. In your opinion, how many 50L waste bags do [anonymized] employees produce on average in one week (5 days) in 2021? *(Write the number)*

The same data about the waste produced by the research center in one week will be shown in three different ways:

- In 3D in augmented reality
- In text on a screen
- In 3D on a screen.

After each one, you will have to complete the associated part of the questionnaire. The questions concern the emotions felt at the time of the visualization. Throughout the

Figure A.1 Experiment questionnaire from chapter 3, Part 1 (pre-questionnaire). Participants filled out this part before the experiment.

APPENDIX A. APPENDICES

Ordering :
Condition:

Indicate the intensity of the emotions you felt while watching the visualisation.

The intensity scale is used in a similar way for "positive" and "negative" feelings. For example, if you feel « extremely strong" you would tick box 5 on the corresponding line. Similarly, if you feel « extremely guilty" you would tick 5 of the corresponding line.

		Very slightly or not at all	A little	Moderately	Quite a bit	Extremely
PANAS 1	Interested	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 2	Distressed	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 3	Excited	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 4	Upset	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 5	Strong	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 6	Guilty	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 7	Scared	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 8	Hostile	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 9	Enthusiastic	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 10	Proud	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 11	Irritable	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 12	Alert	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 13	Ashamed	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 14	Inspired	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 15	Nervous	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 16	Determined	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 17	Attentive	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 18	Jittery	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 19	Active	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5
PANAS 20	Afraid	<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5

Comments :

Figure A.2 Experiment questionnaire from chapter 3, Part 2 (PANAS questionnaire). Participants filled this part after the experiment. They had this document in three copies and had to fill it for each condition.

APPENDIX A. APPENDICES

experience you can go back to your previous answers. We encourage you to compare the intensity of your emotions between the different conditions.

During the experiment you will be asked to visualize a number of trash bags; you will not be asked to count them, the exact number will be communicated to you at the appropriate time.

3. Answer the following questions by ranking the visualisations from 1 to 3.

- To you, which visualisation evoked the most intense emotion(s)?

(Write in the corresponding box 1, 2 or 3 to place them in the first, second or third position).

Text on screen	
3D on screen	
3D in Augmented Reality	

- Which visualisation best enabled you to understand this data?

(Write in the corresponding box 1, 2 or 3 to place them in the first, second or third position).

Text on screen	
3D on screen	
3D in Augmented Reality	

- If you can, justify the answers to the previous questions.

Figure A.3 Experiment questionnaire from chapter 3, Part 3 (post-questionnaire). Participants filled out this part after the experiment and the PANAS questionnaires.

B Videos

- CHI Interactivity Submission Video: <https://youtu.be/BHcUDGuSHh4>
- Descriptive video of our work on ARWAVs: https://osf.io/7jgse?view_only=cd8b973f9df54d7590f8421be6b

B Pre-Registration

https://osf.io/v78ay/?view_only=193d16fbde7641c896018024130c9063

A.3 Supplementary material chapter 4

C Experiment Questionnaire

APPENDIX A. APPENDICES

Pre-questionnaire

Participant ID Number: _____ Sex : _____ Age : _____

Which of your hands is dominant? (Circle one answer)

Left Right Ambidextrous

Which hand do you use with the mouse? (Circle one answer)

Left Right Ambidextrous

Are you color blind? Circle : No Yes What type : _____

Describe your program of study, or your profession, or your area of expertise.
(For example: bachelor's student in electrical engineering)

Do you sometimes use 3D software (games, modeling software, etc.)? Which ones?

Have you ever used a virtual reality or augmented reality headset? For what activity?

Interpupillary distance: _____

Figure A.4 Pre-questionnaire, filled by participants before performing the tracking task.

APPENDIX A. APPENDICES

Post-Questionnaire

In the below table, evaluate each of the conditions according to each of the 4 criteria by circling a number between 1 and 7.

	What mental effort was required?	What physical effort was required?	How did you find the speed of the transition?	How much frustration did you feel?
	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 too slow too fast	1 2 3 4 5 6 7 little a lot
	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 too slow too fast	1 2 3 4 5 6 7 little a lot
	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 too slow too fast	1 2 3 4 5 6 7 little a lot
	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 little a lot	1 2 3 4 5 6 7 too slow too fast	1 2 3 4 5 6 7 little a lot

Which kind of transition did you like the **most**? Circle one.

- [MES] (transition in 1 stage)
- [M, ES] (transition in 2 stages)
- [EM, S] (transition in 2 stages)
- [M, E, S] (transition in 3 stages)

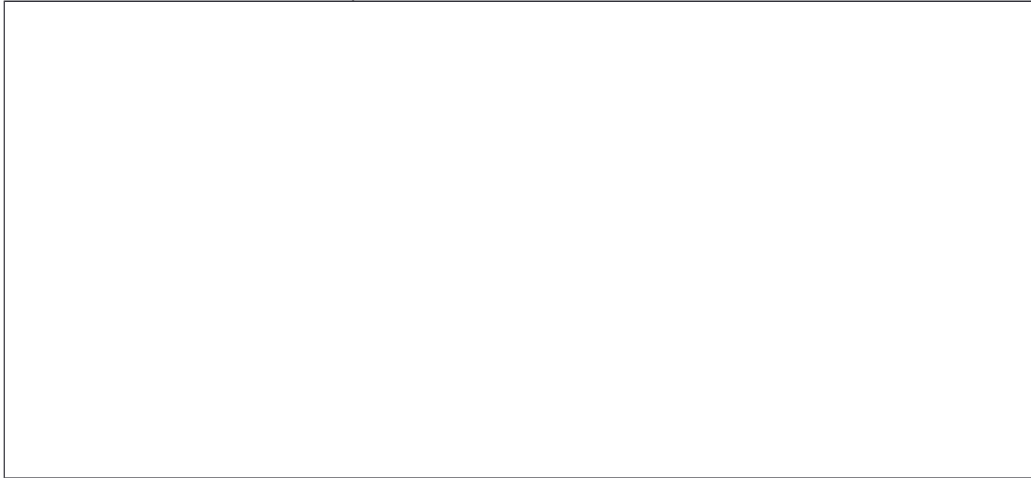
Which kind of transition did you like the **least**? Circle one.

- [MES] (transition in 1 stage)
- [M, ES] (transition in 2 stages)
- [EM, S] (transition in 2 stages)
- [M, E, S] (transition in 3 stages)

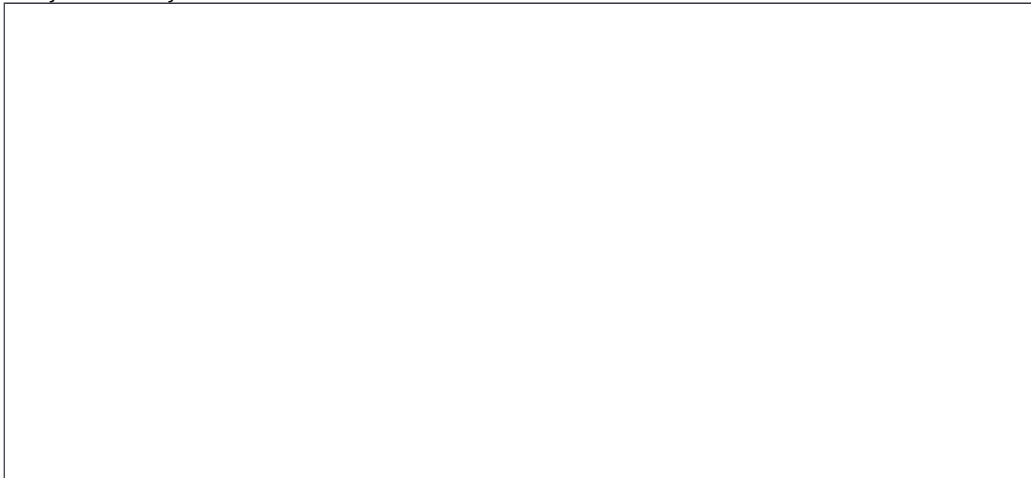
Figure A.5 Post-questionnaire, filled by participants after performing the tracking task, Part 1/2.

General questions

How could the user interface be improved ?

A large, empty rectangular box with a thin black border, intended for participants to write their suggestions for improving the user interface.

Do you have any other comments ?

A large, empty rectangular box with a thin black border, intended for participants to provide any other comments.

Thank you for your time and comments !

Figure A.6 Post-questionnaire, filled by participants after performing the tracking task, Part 2/2.

C Videos

- Density of people prototype: <https://youtube.com/shorts/xAAuaoUaj04>
- Embodied transition prototype: <https://youtu.be/XTwYBnhC8xQ>
- Animated Transitions between Concrete and Abstract Visualizations: <https://youtu.be/eE4CFsuK65U>

C Pre-Registration

https://osf.io/8mu73?view_only=f5ed74fd9e2346228214f5dec973b208

A.4 Supplementary material chapter 5

D Complementary documents

Correspondance list for Figure A.7, Figure A.8, Figure A.9.

1. [49]
2. [222]
3. [137]
4. [199]
5. [284]
6. [235]
7. [126]
8. [118]
9. [268]
10. [172]
11. [269]

APPENDIX A. APPENDICES

Techniques	Non-visibility support	Interactivity	Interaction Realism	Actionability
1	1.Occluded 2.Out of view, occluded, illegible	Medium	Medium	Low
2	Out-of-view	Low	Low	Low
3	Out-Of-view	Medium	Medium	Low
4	Out of view, occluded, illegible	Medium	Medium	High
5	Illegible	High	Low	Low
6	Out of view, occluded, illegible	Medium	Medium	Low
7	Out of view, occluded, illegible	High	Medium	Low
8	Out of view, occluded, illegible	Low	Low	Low
9	Out of view, occluded, illegible	High	Low	Low
10	Out of view, occluded, illegible	High	Medium	Low
11	Illegible	High	Medium	Low

Figure A.7 Techniques handling non-visible referents including visualizations. General Category. See the corresponding papers on the previous page.

Techniques	Representation Fidelity	Representation Refresh Rate	Surroundings Inclusion	Representation Medium	Representation Integration
1	Medium	Low	High	Digital	Medium
2	Medium	Low	Medium	Digital	Medium
3	Low	Low	Low	Physical	High
4	Medium	Low	Low	1.Digital/2.Mixed	1.Low/2.High
5	Medium	Low	Medium	Digital	High
6	Medium	Low	Low (VR)	Digital	Low
7	Medium	Low	High	Physical	High
8	Medium	Low	Medium	Digital	Medium
9	Medium	Low	Low	Physical	High
10	Medium	Low	High	Physical	High
11	Medium	Low	Low	Mixed	High

Figure A.8 Techniques handling non-visible referents including visualizations. Representation Category. See the corresponding papers on the previous page.

Techniques	Location Information	Visualization Richness	Visualization Refresh Rate	Visualization Medium	Visualization Integration
1	1.High 2. Low	Medium	Low	Digital	High
2	High	Medium	High	Digital	High
3	Medium	Medium	High	Digital	High
4	Low	High	High	Digital	1.Low/2.Medium
5	High	High	Low	Digital	Low
6	Low	Medium	Low	Digital	Low
7	Low	High	Low	Digital	Low to High
8	Low	Medium	Low	Digital	Medium
9	Low	High	Low	Digital	Low
10	Low	Medium	Low	Digital	High
11	Low	Medium	Low	Digital	High

Figure A.9 Techniques handling non-visible referents including visualizations. Visualization Category. See the corresponding papers on the previous page.

APPENDIX A. APPENDICES

DOI/uri	Year	Title	Family
https://doi.org/10.1145/19.1686...191821	1994	Passive real-world interface props for neurosurgical visualization	physical replica
https://doi.acm.org/doi/10.1145/223904.223938	1995	Virtual Reality on a WIM: interactive world in miniature	WIM
https://doi.acm.org/doi/pdf/10.1145/218380.218495	1995	Navigation and Locomotion in Virtual Worlds via Flight into Hand-Held Miniatures	WIM
https://link.springer.com/content/pdf/10.1007/BF01682023.pdf	1997	A Touring Machine: Prototyping 3D Hobbit Augmented Reality Systems for Exploring the Urban Environment	labeling
https://www.pitcan.com/papers/ujs199.pdf	1999	Real-world interaction using the FieldMouse	Xray
https://doi.acm.org/doi/pdf/10.1145/502348.502363	2001	View management for virtual and augmented reality	labeling
https://doi.acm.org/doi/pdf/10.1145/571985.572017	2002	An Annotated Situation-Awareness Aid for Augmented Reality	WIM
https://ieeexplore.ieee.org/document/1240754	2003	Optical camouflage using retro-reflective projection technology	Xray
https://doi.org/10.1163/15685903764018468	2003	Telepresence cockpit for humanoid robot control	telepresence
http://portal.acm.org/citation.cfm?doid=989863.989910	2003	Resolving multiple occluded layers in augmented reality	Xray
http://portal.acm.org/citation.cfm?doid=989863.989910	2004	3D location-pointing as a navigation aid in Virtual Environments	out-of-view pointing
http://portal.acm.org/citation.cfm?doid=989863.989910	2004	Interactive Tools for Virtual X-Ray Vision in Mobile Augmented Reality	Xray
https://doi.acm.org/doi/10.1145/1056808.1056883	2005	Attention-based design of augmented reality interfaces	Xray
https://doi.org/10.1163/15685903764018468	2007	Virtually Extended Surgical Drilling Device: Virtual Mirror for Navigated Spine Surgery	virtual mirror
https://doi.acm.org/doi/10.1145/1056808.1056883	2008	TELESARPHONE: Mutual Telepresence Master-Slave Communication System Based on Retroreflective Projection Technology	telepresence
https://www.landonline.com/doi/full/10.9746/jcmst.1.335	2008	A Taxonomy of 3D Occlusion Management for Visualization	general
https://doi.acm.org/doi/10.1109/VR.2009.4811001	2009	Explosion Diagrams in Augmented Reality	Xray
https://www.cs.cmu.edu/~pbrum/seethrough/papers/barnur	2009	Dynamic See-through: Synthesizing Hidden Views of Moving Objects	labeling
https://www.sciencedirect.com/science/article/pii/S009784930	2009	Annotation in outdoor augmented reality	labeling
https://doi.org/10.1109/VR.2009.4811002	2009	The Virtual Mirror: A New Interaction Paradigm for Augmented Reality Environments	virtual mirror
https://doi.org/10.1109/VR.2009.4811002	2009	Complexity and Occlusion Management for the World-in-Miniature Metaphor	WIM
https://ieeexplore.ieee.org/abstract/document/5444792?casa_	2009	Improving Spatial Perception for Augmented Reality X-Ray Vision	Xray
https://ieeexplore.ieee.org/abstract/document/5649547?casa_	2010	Mixed reality in virtual world teleconferencing	Telepresence
https://doi.org/10.1162/jres.2010.1162	2010	An Augmented Reality X-Ray system based on visual saliency	labeling
https://doi.org/10.1162/jres.2010.1162	2010	An augmented reality interface to contextual information	out-of-view pointing
https://doi.org/10.1162/jres.2010.1162	2010	Visualization of off-screen objects in mobile augmented reality	virtual twin
https://doi.org/10.1162/jres.2010.1162	2010	Mixed reality in virtual world teleconferencing	WIM
https://doi.org/10.1162/jres.2010.1162	2010	Egocentric space-distorting visualizations for rapid environment exploration in mobile mixed reality	Xray
https://doi.org/10.1162/jres.2010.1162	2010	Hand-Based Disocclusion for the World-in-Miniature Metaphor	Xray
https://doi.org/10.1162/jres.2010.1162	2010	Document Search Support by Making Physical Documents Transparent in Projection-Based Mixed Reality	Xray
https://doi.org/10.1162/jres.2010.1162	2011	Visualization Techniques for Augmented Reality	Xray
https://doi.org/10.1162/jres.2010.1162	2011	Roundplot: Focus+context interface for off-screen objects in 3D environments	out-of-view pointing
https://doi.org/10.1162/jres.2010.1162	2011	Peripheral vision annotation: noninterference information presentation method for mobile augmented reality	out-of-view pointing
https://doi.org/10.1162/jres.2010.1162	2011	Mutual telepresence surrogate system: TELESAR4 - telepresence in real environments using autostereoscopic immersive display -	telepresence
https://doi.org/10.1162/jres.2010.1162	2011	Slice WIM: a multi-surface, multi-touch interface for overview/detail exploration of volume datasets in virtual reality	WIM
https://doi.org/10.1162/jres.2010.1162	2011	Augmented reality and image overlay navigation with OsiriX in laparoscopic and robotic surgery: not only a matter of fashion	Xray
https://doi.org/10.1162/jres.2010.1162	2012	Image-driven view management for augmented reality browsers	labeling
https://doi.org/10.1162/jres.2010.1162	2012	World-in-Miniature Interaction for Complex Virtual Environments	WIM
https://doi.org/10.1162/jres.2010.1162	2013	Supporting interoperability and presence awareness in collaborative mixed reality environments	Telepresence
https://doi.org/10.1162/jres.2010.1162	2013	Sidebars: Improving awareness of off-screen elements in mobile augmented reality	out-of-view pointing
https://doi.org/10.1162/jres.2010.1162	2013	A Review of Mobile Robotic Telepresence	telepresence

Figure A.10 List of selected papers.

D Video

VIS presentation: <https://youtu.be/YZBh0GpBqSw>