



HAL
open science

New approach of 3D for monumental heritage

Corentin Cou

► **To cite this version:**

Corentin Cou. New approach of 3D for monumental heritage. Graphics [cs.GR]. Université de Bordeaux, 2023. English. NNT : 2023BORD0382 . tel-04364363v2

HAL Id: tel-04364363

<https://hal.science/tel-04364363v2>

Submitted on 12 Jan 2024

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

ÉCOLE DOCTORALE
MATHÉMATIQUE ET INFORMATIQUE
SPÉCIALITÉ : INFORMATIQUE

Présentée par
Corentin COU

**Nouvelle vision de la 3D pour le
patrimoine monumental**

Sous la direction de : **Xavier GRANIER**

Soutenue le 8 décembre 2023

Membre du jury :

Mathieu HEBERT Professeur	Institut d'Optique	Président
Daniel MENEVEAUX Professeur	Université de Poitiers	Rapporteur
Adolfo MUNOZ Professeur	Universidad de Zaragoza	Rapporteur
Xavier GRANIER Professeur	Institut d'Optique	Directeur
Romain PACANOWSKI Chargé de Recherche	Inria	Invité
Mercedes VOLAIT Directrice de Recherche	CNRS	Invitée

Nouvelle vision de la 3D pour le patrimoine monumental

Résumé : Le sujet de la thèse consiste à développer une nouvelle vision de la 3D pour le patrimoine monumental. Il s'agit à la fois de répondre aux questionnements en histoire des arts par l'incorporation d'une information visuelle enrichie, et aux nécessités de validation scientifique. Pour cela, elle s'appuie sur la structuration spatiale qu'offre la 3D pour les restitutions, pour être un support d'agrégation des données historiques mobilisées.

La base de l'expérimentation proposée est un grand décor néo-islamique étudié par Mercedes Volait dans *Maisons de France au Caire : le remploi de grands décors mamelouks et ottomans dans une architecture moderne*. Son intérêt est d'illustrer un dispositif décoratif complexe soulignant le réemploi islamique dans le Caire moderne du XIXe siècle. Réalisé au Caire (Égypte) entre 1875-1879 pour servir à l'aristocrate Gaston de Saint-Maurice (1831-1905), cette résidence n'a pas survécu dans sa forme initiale, il a été partiellement déplacé et remonté en 1937 dans une nouvelle localisation au sein de la capitale égyptienne.

La recherche doctorale proposée dans cette thèse est de nature interdisciplinaire et se situe à la croisée des domaines scientifiques suivants : histoire de l'art de l'Islam, archéologie historique, visualisation 3D et archéométrie. Elle consiste plus précisément à développer des méthodes d'acquisition de données visuelles variées sur site, de façon non invasive, et appliqué au patrimoine monumental. Ces méthodes se concentrent sur trois axes : la capture d'information de forme, de couleur, et des propriétés de réflexion des matériaux. Enfin, l'aspect visuel est acquis sur place au Caire.

Mots-clés : Acquisition, 3D, Patrimoine, Apparence, Système de référencement, Informatique graphique

New approach of 3D for monumental heritage

Abstract: The subject of the thesis consists in developing a new vision of 3D for monumental heritage. It is both to respond to questions in the art history by incorporating enriched visual information, and to the needs for scientific validation. For this, it relies on the spatial structuring offered by 3D for restitutions, to be a medium for aggregating the historical data used.

The basis of the proposed experiment is a large neo-Islamic decor studied by Mercedes Volait in *Maisons de France au Caire : le remploi de grands décors mamelouks et ottomans dans une architecture moderne*. Its interest is to illustrate a complex decorative device showing Islamic re-use in modern 19th century Cairo. Made in Cairo (Egypt) between 1875-1879 to serve the aristocrat Gaston de Saint-Maurice (1831-1905), this residence has not survived in its initial form, it was partially moved and reassembled in 1937 in a new location within the Egyptian capital.

The doctoral research proposed in this thesis is of an interdisciplinary nature and lies at the crossroads of the following scientific fields: history of Islamic art, historical archaeology, 3D visualization and archaeometry. It consists more precisely in developing methods for acquiring various visual data on site, in a non-invasive way, and applied to monumental heritage. These methods focus on three axes: the capture of information of shape, color, and reflection properties of materials. Finally, the visual aspect is acquired on site in Cairo.

Keywords: Acquisition, 3D, Humanities, Appearance, Referencing Systems, Computer Graphics

Unités de recherche

InVisu USR 3103, CNRS, Institut National d'Histoire de l'Art, Paris, France.

LP2N UMR 5298, Institut d'Optique Graduate School, Talence, France.

Inria Bordeaux Sud-Ouest, Talence, France.

Remerciements

Je tiens d'abord à remercier Mathieu Hébert, Daniel Menneveux et Adolfo Munoz pour avoir accepté d'être membres du jury de cette thèse. Je les remercie pour leur temps, leur bienveillance, ainsi que les discussions très intéressantes et constructives qui en ont découlées.

Je suis de plus très reconnaissant envers mes deux co-encadrants de thèse : Romain Pacanowski et Gaël Guennebaud. Votre investissement, chacun sur vos parties de ma thèse, m'a permis d'avancer sereinement et de passer un doctorat qui restera un très bon souvenir. Mais je tiens aussi et surtout à souligner votre rigueur scientifique qui m'a fait progresser et me dépasser en tant que jeune chercheur. De plus, merci Gaël pour m'avoir donné l'opportunité de réaliser ce stage de fin d'étude avec toi, c'est ce qui m'a fait découvrir, et surtout aimer, la recherche et m'a poussé à continuer dans cette voie.

Je tiens d'autant plus à remercier mon directeur de thèse, Xavier Granier, envers qui je suis extrêmement reconnaissant. Je te remercie du fond du cœur de m'avoir fait confiance sur ce si beau projet qui te tenait tant à cœur, puis pour toute notre relation doctorant/encadrant qui dépasse le simple cadre du labo. Tu auras été un super encadrant, que ce soit pour la partie scientifique, mais surtout pour toutes les difficultés passées à côté, et je ne peux que te tirer mon chapeau pour cela.

Je remercie aussi les personnes que j'ai pu rencontrer au cours de ces 4 années de thèse, que ce soit Hortense, Aurélie et Ayed à IRAMAT, Pascal et Vincent à Archéovision, ou Mercedes à InVisu. Mais aussi au sein de l'équipe Manao, où je remercie les membres permanents Pascal, Pierre et Pat, mais aussi les non permanents qui sont devenus plus que des collègues au cours de ma thèse : David, Camille, Mégane, Simon, Gary, Melvin, Pana, Pierre(s), Jean, Marjorie. Ce sont tous ces moments ensemble qui donnaient envie de venir le matin au boulot, et c'est grâce à vous, donc merci.

Parmi les personnes que j'ai côtoyées à Manao, je tiens particulièrement à dire merci à Morgane, qui est devenue bien plus qu'une collègue. Si ton soutien quotidien est devenu au fur et à mesure un boost (et notamment au moment de la rédaction et de la fin de thèse), il a été, et demeure, surtout une source de force lorsque ma santé refait des siennes, et je ne saurais te remercier assez pour cela.

De la même façon, je tiens à remercier mes amis les plus proches qui m'auront suivi pendant cette drôle d'aventure. Je sais que vous vous reconnaîtrez, mais si je dois essayer de citer quelques noms : les PB, dont Arnaud qui m'a jamais lâché depuis cette mauvaise journée du 31 décembre 2019 ; mes anciens colloqs de Russie ; les Mathildes, Fantine, Flavien, Paul, Gaby qui sont un soutien sans faille depuis le lycée ; Anto, Louis, Max, Etienne, Théo, et tous les copains du basket qui sont là chaque jour au moins par message ; Nano, JB, Paul, Etienne et tout le groupe de Bordeaux... Je ne peux tous vous citer ici mais je pense à vous, et je vous suis terriblement reconnaissant d'avoir été là, de l'aide en hospitalisation aux simples messages de soutien.

Enfin, je tiens bien sûr à chaudement remercier ma famille. Mes grands-parents, d'abord. Puis Arthur à qui j'espère t'avoir donné le goût des sciences en plus du basket désormais. Maëla, j'ai la chance d'avoir une petite sœur aussi géniale, avec qui je m'entends si bien. Puis mes parents. Maman, car tu as toujours été là et je pourrais te remercier assez. Je pense que plus ça va, plus je te ressemble. Papa, j'espère t'avoir rendu fier, car c'est toi qui m'a donné cette amour de la science et ce sérieux dans les études.

Je tiens juste à rajouter un merci pour finir aux équipes médicales du CHU de Brest oncologie, de l'Institut Bergonié, ainsi que les kinés, les APA et le personnel soignant m'ayant entouré. Si je suis là à écrire ces remerciements, c'est grâce à vous.

Ce travail n'est ainsi pas seulement le résultat de mes efforts seuls, mais a été réalisé grâce au soutien moral et scientifique des personnes ci-dessus sans qui je ne serais pas arrivé ici.

Ce doctorat n'aurait pas pu avoir lieu sans les financements du CNRS, via le projet Smart 3D et le projet tremplin μ ChaOS, que je tiens donc à remercier.

Résumé en français

Du monde réel à une représentation informatique : l'acquisition du patrimoine architectural

Avec l'émergence d'images photoréalistes, l'optique numérique et l'informatique graphique ont cherché à représenter les objets qui nous entourent de la manière la plus fidèle possible. Plusieurs domaines de recherche ont émergé pour l'améliorer : la simulation du transport de la lumière, le traitement de la géométrie, ou les propriétés de réflexion des matériaux. Nous avons atteint un point où certaines scènes entièrement générées par ordinateur peuvent imiter la réalité et tromper nos yeux.

Cependant, la question s'est posée de savoir pourquoi nous devrions nous limiter à recopier informatiquement notre monde alors que nous avons la capacité de le capturer directement avec des dispositifs d'acquisition. Des études visent alors à envisager de nouvelles méthodes d'acquisition, d'abord du monde réel au monde virtuel, puis des ordinateurs aux observateurs. Pas à pas, elles se sont spécialisées dans la capture de ce qui rend les objets qui nous entourent si reconnaissables.

Dans ce contexte, nous avons vu l'émergence de dispositifs et de logiciels pour la reconstruction de la forme 3D. Nous avons pu ainsi synthétiser les objets du vrai monde sous la forme de maillages, composés de points 3D (appelés sommets) reliés par des triangles (appelés faces). Ensuite, nous avons associé des couleurs aux objets lors de leur capture sous la forme de triplets rouge-vert-bleu (RVB) associés à chaque sommet ou chaque face.

Cependant, une forme et des couleurs ne suffisent pas à capturer et reproduire notre environnement parfaitement. Une boîte noire en bois ou en métal ne présente effectivement pas du tout le même aspect malgré une forme et une couleur commune. En réalité, il est également nécessaire de prendre en compte les propriétés optiques du matériau pour représenter avec précision la réflexion de la lumière. L'apparence d'un matériau peut être résumée par la *Bi-directional Reflectance Distribution Function* (BRDF). Cette fonction représente la proportion de lumière réfléchiée par un matériau pour une observation et une direction lumineuse données. Par conséquent, les systèmes d'acquisition les plus précis capturent maintenant un maillage associé à une BRDF, voire une BRDF variant spatialement (SVBRDF), au lieu d'une simple couleur RVB. Chaque sommet est lié à un SVBRDF avec des paramètres particuliers.

L'avantage de l'acquisition d'un modèle 3D est qu'elle permet d'ajouter encore plus d'informations. À l'ère des données, les acquisitions deviennent de plus en plus diversifiées (spectres de réflectance, données médicales, annotations). Ces données peuvent ensuite être associées aux sommets du maillage du modèle. De nouvelles visualisations sont développées et l'objet 3D du monde réel devient plus attrayant pour une large gamme d'applications.

C'est notamment le cas du patrimoine culturel. En effet, c'est un domaine où la numérisation apporte de nombreuses réponses. Tout d'abord, la question de la sauvegarde du patrimoine humain est l'une des questions clés de l'histoire récente : depuis le début du XIXe siècle, l'émergence des musées a montré la préoccupation de sauvegarder les biens culturels. La numérisation d'objets anciens et fragiles permet de conserver une copie virtuelle à un moment donné, par exemple. De plus, la numérisation permet l'exploitation de modèles informatiques enrichis en

données. La recherche en synergie entre historiens et physiciens ou informaticiens génère donc de nouveaux progrès significatifs autour de ces deux axes. D'une part, le besoin de conservation nécessite un processus d'acquisition de plus en plus réaliste, tandis que d'autre part, la plupart des applications nécessitent des méthodes de visualisation plus simples et plus intuitives.

Dans ce contexte, si le patrimoine fait partie intégrante de la gamme d'objets qui peuvent être transportés en laboratoire, l'architecture aborde également progressivement des questions liées à la numérisation et à l'acquisition. De nouvelles contraintes apparaissent alors, telles que des environnements d'acquisition non contrôlés ou des tailles beaucoup plus imposantes. De plus, maintenir une représentation correcte des données d'acquisition aussi réaliste que possible soulève de nombreux autres défis.

L'application à l'Hôtel Saint-Maurice, un exemple du réemploi

C'est dans ce contexte que cette thèse s'inscrit au sein du projet *Smart 3D* mené conjointement par l'INHA, l'Inria, Archéoscience Bordeaux et l'IOGS, et bénéficie également du soutien de l'IFAO par le biais de sa contribution à travers le programme "La fabrique du Caire moderne". L'objectif du projet *Smart 3D* est d'établir une nouvelle perspective sur la représentation tridimensionnelle dans le contexte du patrimoine monumental (c'est à dire de l'étude architecturale). Cette approche vise à répondre à la fois aux questions relatives à l'histoire de l'art en incluant des informations visuelles améliorées, et aux exigences de validation scientifique en servant de plateforme pour la consolidation des données historiques utilisées dans les restaurations. L'objet d'application est un hôtel particulier du XIXe siècle situé au Caire appelé Hôtel particulier Gaston de Saint-Maurice, présentant un sujet d'étude important sur le réemploi islamique. À l'intérieur de cette résidence, les carreaux de céramique d'Iznik sont un exemple typique de la réutilisation caractéristique du réemploi.

Le réemploi islamique à l'hôtel Saint-Maurice. Bien que principalement associée à l'Antiquité et au Moyen Âge, la réutilisation architecturale représente une pratique largement diffusée en Europe et au-delà, tout au long du XIXe siècle. Dans des capitales du Moyen-Orient telles que Le Caire et Damas, des matériaux récupérés sur des sites délabrés ou en cours de rénovation ont été réutilisés pour le revêtement de sol et les murs (marbre), les portes et les meubles (menuiserie), le lambris (céramique) et les plafonds (plâtre peint). Ces arrangements architecturaux sont actuellement identifiés comme appartenant au **réemploi islamique**.

La réemploi islamique a été promue au Caire par des expatriés français dans les années 1870, un exemple majeur étant l'**Hôtel Saint-Maurice**. Construit en 1875-79, la structure associait des éléments récupérés et des répliques d'ornements historiques des monuments du Caire. Le bâtiment a été démonté en 1937 mais partiellement reconstruit à un autre emplacement du Caire. Bien que l'hôtel n'ait pas survécu dans sa forme initiale, des vestiges dispersés existent, ainsi qu'une documentation étendue (dessins, plans, photographies, textes) sur le projet initial.

Les céramiques d'Iznik de Saint-Maurice. Les céramiques émaillées constituent une part importante du patrimoine culturel visible dans les archives archéologiques. Elles sont relativement simples à produire et ont été découvertes dans le monde entier, du bassin méditerranéen jusqu'à des régions aussi éloignées que la Chine et le Japon, depuis des millénaires. En raison de leur utilisation répandue et variée (que ce soit la vaisselle et d'autres objets utilitaires, ou pour l'architecture d'intérieure et d'extérieure) et de leur résistance à la dégradation, elles sont parmi les objets culturels matériels les plus couramment trouvés lors des fouilles archéologiques.

Une céramique particulière est étudiée dans notre contexte : les **céramiques d'Iznik**. Ces céramiques ont été produites entre le XVe et le XVIIe siècle pour divers usages, de la vaisselle aux carreaux. Originaires de la ville d'Iznik, dans l'ouest de l'Anatolie, elles se sont dispersées dans tout l'Empire ottoman, le monde arabe et la région méditerranéenne.

Or, les carreaux examinés dans la résidence Saint-Maurice présentent 4 motifs caractéristiques de la production d'Iznik que l'on retrouve sur de nombreux bâtiments historiques du Caire actuel. Ces derniers soulèvent alors des questions ouvertes sur leurs origines et leur réutilisation, alors que très peu de documentation existe pour y répondre.

Le défi ici est de **chercher la possibilité de mettre en évidence des groupes**, ou d'identifier d'une manière ou d'une autre des tendances claires au sein des 4 groupes de céramiques partageant des propriétés similaires.

Cette thèse vise donc à proposer des solutions axées sur la production d'une plus grande densité d'informations visuelles et le développement d'un lien caractéristique entre les visuels et le monde physique dans le cadre d'observations cohérentes. Ces observations sont le résultat de méthodes d'acquisition portables et non invasives. À cette fin, nous présentons d'abord le contexte et les enjeux historiques autour de la résidence Saint-Maurice, puis nous développons le travail effectué autour de trois axes :

1. L'acquisition de la géométrie au niveau microscopique (Chapitre 1).
2. L'acquisition et le traitement des spectres de réflectance (Chapitre 2).
3. L'acquisition de la SVBRDF pour le patrimoine monumental (Chapitre 3).

L'acquisition de géométrie à l'échelle microscopique

Afin de caractériser les carreaux de céramique d'Iznik, nous listons trois éléments distincts. Il est ainsi possible d'étudier les motifs dessinés sur les carreaux, les pigments utilisés, et l'épaisseur de la glaçure transparente à la surface. Dans notre cas, nous nous concentrons sur les pigments et la glaçure. En effet, les motifs étant réalisés à la main, ils peuvent être changés plus facilement d'un carreau à l'autre, tandis que les autres procédés résultent davantage de méthodes de fabrication.

Dans un premier temps, nous nous intéressons à la mesure de l'épaisseur de la glaçure des céramiques. Pour cela, nous étudions les différentes techniques permettant l'acquisition de la géométrie d'un objet, à une résolution de quelques micromètres. L'état de l'art montre que les techniques basées sur un appareil photo sont très rapidement limitées à une résolution de plus de 50-100 μm en raison de la profondeur de champ réduite. Afin d'utiliser la diminution de la profondeur de champ, nous avons donc développé et amélioré une méthode basée sur le Depth from Focus pour capturer des profils 3D. Cette méthode est une méthode monoscopique, pour laquelle nous réalisons une pile d'images en variant la profondeur du plan focal. Pour chaque pixel, nous déterminons quel est l'image la plus nette, et donc quelle est la profondeur réelle du pixel. A la fin, nous obtenons une carte de profondeur, et donc une géométrie. Nous avons amélioré l'état de l'art grâce à un système de fenêtre glissante linéaire, que nous évaluons sur des exemples synthétiques.

Dans la suite, nous appliquons ce système pour mesurer la glaçure des céramiques. En effet, la glaçure est transparente, mais elle possède des imperfections à la surface, ce qui permet d'obtenir des cartes de profondeur où la mesure peut se faire successivement au-dessus et en dessous de la glaçure. En observant ces deux niveaux, et en mesurant la distance les séparant, nous pouvons ainsi déterminer l'épaisseur de la glaçure. Après une validation de la méthode en laboratoire sur un grand nombre de céramiques d'origines diverses, nous l'appliquons sur nos céramiques d'Iznik au Caire. Les résultats montrent que les variations au sein d'un même carreau ne permettent pas de caractériser les céramiques de Saint-Maurice une à une, mais la comparaison avec d'autres carreaux d'Iznik de Tunis permet de considérer que nous avons un procédé de fabrication commun avec les autres céramiques d'Iznik étudiées.

L'acquisition et le traitement d'image hyperspectrale dans l'étude des pigments

La caractérisation des carreaux étant impossible avec la simple mesure de l'épaisseur de leur glaçure, nous nous intéressons alors à l'étude des pigments utilisés. Pour cela, nous avons fait l'acquisition des carreaux par imagerie hyperspectrale (IHS). L'IHS permet de générer des images comme en photographie couleur classique, mais un spectre remplace le triplet RVB pour chaque pixel.

Dans ce cadre, l'émergence des caméras hyperspectrales a permis d'acquérir des millions de spectres sur des échantillons. Cela génère le besoin d'utiliser des méthodes de traitement et de visualisation des données car l'observation manuelle n'est plus possible. Cependant, lorsque les données deviennent complexes avec des variations de recettes, d'intensité ou de mélange au sein d'un même colorant ou pigment, les méthodes courantes de segmentation ne fonctionnent plus très bien (groupes selon l'intensité et non les décalages visibles sur les spectres par exemple). En effet, des problèmes surviennent en raison de l'énorme quantité de données qui rendent très sensibles les variations au sein d'un même pigment ou colorant.

Afin de réduire la dimensionnalité de nos données acquises, nous développons donc une méthode basée sur la comparaison entre les spectres étudiés et un jeu de spectres caractéristiques à l'aide de coefficients de corrélation de Pearson. Cette méthode améliore la robustesse du traitement des images hyperspectrales et réduit la quantité de données en générant des outils pour une étude de similarité entre les spectres étudiés et une base de données. La première étape consiste à créer une base de données de spectres caractéristiques utilisés pour les corrélations. Ensuite, certains prétraitements sont appliqués à l'image hyperspectrale étudiée (comme le filtrage spatial pour le débruitage). La dernière étape, le point principal de notre méthode, consiste à calculer un coefficient de corrélation de Pearson entre les spectres étudiés et chacun des spectres clés de la base de données. Ces nouvelles valeurs obtenues peuvent être utilisées pour des méthodes courantes de segmentation et de visualisation.

Nos processus ont été appliqués à différents cas. Après l'avoir testé sur des nuanciers de colorants simples avec une base de données interne pour validation, nous l'avons appliqué à une tapisserie d'Aubusson (XVIII^e siècle) avec un nuancier de colorants externe pour obtenir une ACP dans laquelle les clusters sont beaucoup plus nets qu'une ACP classique sur données brutes. Cela a donc permis une identification plus fine des colorants. À titre d'illustration, nous identifions les teintures rouges sur une tapisserie comme étant de la garance et non de la cochenille.

Enfin, la méthode a été appliquée à des groupes de pigments présents sur les carreaux de céramique d'Iznik de l'hôtel Saint-Maurice et a permis de réduire drastiquement la quantité de données en ne conservant que les informations pertinentes. Les visualisations permettent de montrer les similitudes ou les différences entre les carreaux : certains motifs très communs dans le Caire moderne présentent de grosses variations tandis que les carreaux dont les motifs sont plus rares semblent présenter des pigments identiques.

Ces cas illustrent l'amélioration de la robustesse tout en réduisant la quantité de données sur les critères de similarité de colorants ou de pigments.

L'acquisition de l'apparence sur site à l'échelle monumentale

Enfin, nous avons cherché à améliorer le modèle 3D issu de la restitution de l'hôtel Saint-Maurice. Pour cela, nous avons souhaité passer d'un modèle composé d'un mesh 3D associé à des couleurs RGB, à un mesh 3D associé à des propriétés de réflectance, c'est à dire à une apparence complète du palais.

Or, l'acquisition de la SVBRDF d'objets 3D est une tâche très complexe en raison de la dimensionnalité élevée de la mesure (c'est-à-dire les directions de la lumière et de la vue, mais également les variations spatiales). Dans ce contexte, les méthodes en laboratoire parviennent

à obtenir des données très précises grâce à des étalonnages rigoureux dans un environnement contrôlé. Cependant, les calibrations de caméras et de lumière effectués in situ sont beaucoup plus complexes. Les travaux antérieurs sont limités en distribution angulaire ou nécessitent des a priori forts sur la SVBRDF ou la forme de l'objet capturé. De plus, les contraintes in situ semblent montrer qu'une approche basée sur des modèles mathématiques peut être mieux adaptée qu'une approche basée sur les données mesurées. Cela explique pourquoi la création d'une configuration portable pour des environnements non contrôlés est un véritable défi.

Nous présentons ainsi une nouvelle méthode d'acquisition portable et adaptable à diverses tailles d'objets, de façon à pouvoir réaliser des mesures pour améliorer la restitution de l'hôtel particulier. Cette méthode consiste à acquérir et reconstruire la forme, la SVBRDF et la carte des normales à partir de deux appareils photos et d'un spot lumineux. De plus, nous proposons une nouvelle méthode de calibration rapide de la lumière, déployable sur site, en utilisant deux boules miroirs. Notre méthode a l'avantage d'être très évolutive et adaptable in situ. Elle repose sur un positionnement rapide pour chaque couple lumière/vue, qui permet une grande liberté sur la répartition angulaire, sans a priori fort.

Nous validons notre méthode numériquement sur des scènes synthétiques et nous présentons également des résultats sur des objets réels de réflectance et de forme complexes en laboratoire. Enfin, nous l'appliquons sur site sur des parties du plafonds et des mosaïques de l'hôtel Saint-Maurice.

Cette thèse présente donc trois aspects différents dans l'étude applicative de l'hôtel Saint-Maurice. Si ces trois chapitres peuvent sembler assez éloignés, ils se recroisent en réalité car ils permettent de faire une étude globale de l'acquisition numérique d'un objet : en effet, nous avons étudié la forme de l'objet, puis sa couleur en deux parties, c'est-à-dire en étudiant son spectre réfléchi, et les propriétés angulaires de cette réflexion. Ces axes de recherche permettent d'avoir une meilleure compréhension de comment adapter les contraintes sur site de l'acquisition à ces trois points fondamentaux de la caractérisation visuelle et la capture d'un objet du monde réel.

My publications

- [1] Vincent Baillet, Pascal Mora, Corentin Cou, Sarah Tournon-Valiente, Mercedes Volait, Xavier Granier, Romain Pacanowski, and Gaël Guennebaud. 3D for Studying Reuse in 19th Century Cairo : the Case of Saint-Maurice Residence. In GCH 2021-Eurographics Workshop on Graphics and Cultural Heritage, 2021.
- [2] Corentin Cou, Ayed Ben Amara, and Xavier Granier. Non-invasive on-site method for thickness measurement of transparent ceramic glazes based on depth from focus. Archaeometry, 2022.
- [3] Corentin Cou and Gaël Guennebaud. Depth from Focus using Windowed Linear Least Squares Regressions. The Visual Computer, pages 1–10, 2023.
- [4] Claudia Sciuto, Federico Cantini, Rémy Chapoulie, Corentin Cou, Hortense de la Codre, Gabriele Gattiglia, Xavier Granier, Aurélie Mounier, Vincenzo Palleschi, Germana Sorrentino, and Simona Raneri. What Lies Beyond Sight ? Applications of Ultraportable Hyperspectral Imaging (VIS-NIR) for Archaeological Fieldwork. Journal of Field Archaeology, pages 1–14, October 2022.

Contents

Introduction	17
Context	19
Reuse in 19th-Century Cairo: the case of Saint-Maurice residence	19
The Iznik ceramics and the questions of their reuse	21
1 Depth from Focus for Shape acquisition	25
1.1 3D acquisition of small reliefs - state of the art	25
1.1.1 On-site 3D acquisition techniques	25
1.1.2 Exploiting the depth of field information	30
1.2 Enhancing the Depth from Focus accuracy	32
1.2.1 Overview of the pipeline	32
1.2.2 Peak estimation using Linearized Sliding-window	33
1.2.3 Post-correction using MLS	36
1.2.4 Experiments and Results	36
1.3 Application to monumental heritage: measuring ceramics glaze thickness	43
1.3.1 Non-invasive on-site method for thickness measurement	44
1.3.2 Samples	46
1.3.3 Results and discussion	49
1.3.4 Conclusion on the thickness measurement method	51
1.3.5 Results on the tiles from Saint-Maurice residence	53
2 Hyperspectral Imaging for Reflectance Spectra Similarity Studies	57
2.1 The interest of hyperspectral imaging for pigments characterization	57
2.1.1 Differences between Hyperspectral and RGB Imaging systems	57
2.1.2 Ultraportable systems for on-site measure	60
2.1.3 Notion of metamerism on colors	61
2.1.4 Hyperspectral imaging data processing	63
2.1.5 Applications of hyperspectral imaging on Cultural Heritage	65
2.2 A new dimensionality reduction based on Pearson correlation	70
2.2.1 Description of the pipeline	72
2.2.2 Reducing the data for a better relevance	74
2.3 Experimentation of the method on dyes on Aubusson tapestry	75
2.3.1 Context	75
2.3.2 Method	76
2.3.3 Visualization on the color charts	79
2.3.4 Dye identification	79
2.4 Application to groups of pigments on Iznik ceramics	82
2.4.1 Method	82
2.4.2 The question of tiles segmentation	87
2.4.3 Similarity study results and discussions from visualizations	90

3	Capturing the in-situ BRDF	93
3.1	BRDF of materials	93
3.1.1	BRDF definition	93
3.1.2	Models for simplification and utilization purposes	96
3.1.3	Adding spatial information	99
3.2	State of the art of BRDF acquisition methods	100
3.2.1	Acquisitions made in a controlled environment	100
3.2.2	Acquisitions under constraints	102
3.3	In-situ and highly portable acquisition method	106
3.3.1	Acquisition process	107
3.3.2	From image data to BRDF measurements	109
3.3.3	From BRDF measurements to BRDF model	122
3.4	Results and evaluation	126
3.4.1	Evaluation on synthetic scenes	126
3.4.2	Results on real objects	130
3.5	Application to Saint-Maurice residence	134
	Conclusion	137
	Future Work	139
	Annex	141
I	Results of Depth from focus for ceramics glaze thickness	141
II	Hyperspectral resolution resampling and impacts on our results	147
III	Pigments studies of the Iznik ceramics tiles	152
IV	Personal Blender/Python interface manual	158
V	List of symbols in Chapter 3	162
	Bibliography	165

Introduction

With the emergence of photorealistic images, Digital Optics and Computer Graphics have sought to represent objects around us as faithfully as possible. Several areas of research emerged for improvement: light transport simulation, geometry processing, reflection properties of materials. We have reached to a point where some fully computer-generated scenes can mimic reality and fool our eyes.

However, we also considered the question of why we should restrict ourselves to computer-based replication of our world when we have the capability to directly capture it with acquisition devices. Studies like [128, 193] then aim at considering new methods of acquisition, first from the real world to the virtual world, then from computers to observers. Step by step, they specialized in capturing what makes the objects around us so recognizable. We saw in this context the rise of devices and software for the reconstruction of 3D shape. Objects are thereby synthesized in the form of meshes, composed of 3D points (called vertices) connected by triangles (called faces). Then, colors are associated with the objects during the capture as red-green-blue (RGB) triplet associated to each vertex or each face.

But, one shape and colors is not enough to capture and reproduce our environment perfectly. A black box made of wood and metal effectively does not present the same appearance at all despite a common shape and color. At a matter of facts, it is also necessary to take into account the optical properties of the material to represent light reflection accurately. The appearance of a material can be summarized by the Bi-directional Reflectance Distribution Function (BRDF). This function represents the proportion of light reflected by a material for a given light observation and direction. Therefore, the most accurate acquisition systems now capture a mesh associated with a BRDF, or even a spatially varying BRDF (SVBRDF), instead of a simple RGB color. Each vertex is linked to an SVBRDF with particular parameters.

The advantage of acquiring a 3D model is that it allows to add even more information. In the era of data, acquisitions are becoming more and more diverse (reflectance spectra, medical data, annotations). These data can then be associated with vertices of the model mesh. New visualizations are developed and the 3D object from real world becomes more appealing for a large range of applications.

This is notably the case of Cultural Heritage. Indeed, this is an area where digitization brings many answers. First, the issue of saving human heritage is one of the key issues in recent History: since the beginning of the 19th century, the emergence of museums has shown the concern for safeguarding cultural property. The digitization of old, fragile objects makes it possible to keep a virtual copy at a given time for example. In addition, the digitalization enables the exploitation of data-enriched computer models. Research in synergy between historians and physicists or computer scientists therefore generates new significant advances around these two axes. On the one hand, the need for conservation requires an increasingly realistic acquisition process, whereas most applications require simpler and more intuitive visualization methods on the other one.

In this context, Cultural Heritage is first and foremost part of the scale of the objects that can be transported to the lab: Monumental Heritage is also slowly taking up issues related to digitization and acquisition. New constraints then appear such as uncontrolled acquisition environments, or much more imposing sizes. Upon to that, keeping correct rep. of acquisition

data as realistic as possible raises many other challenges.

It is in this context that this thesis is taking part of the *Smart 3D* project jointly led by INHA, Inria, Archéoscience Bordeaux and IOGS, which also benefits from the support of IFAO through its contribution to the research of the program “La fabrique du Caire moderne”. The objective of the *Smart 3D* project is to establish a novel perspective on three-dimensional representation in the context of monumental heritage. This approach aims to address both the inquiries pertaining to art history through the inclusion of enhanced visual information, and the requirements for scientific validation by serving as a platform for the consolidation of historical data used in restorations. The object of application is a 19th century private hotel located in Cairo called Saint-Maurice residence, presenting an important subject of study on Islamic revivalism. Within this residence, the ceramic tiles of Iznik are a typical example of the characteristic reuse of revivalism. It appears as a subject of study on its own.

This thesis seeks to propose solutions focusing on the production of a greater density of visual information and the development of a characteristic link between the visuals and the physical world within a framework of consistent observations. These observations are the result of portable and non-invasive acquisition methods. To this end, we first present the context and the historical issues around Saint-Maurice residence, then we develop the work carried out around 3 axes:

1. Shape acquisition at the microscale level (Chapter 1). In this chapter, I detail the precision improvement made on existing methods based on Depth from Focus algorithms and how I use it to measure the glaze thickness of old ceramics as part of a study of their physical characteristics.
2. Reflectance spectra acquisition and processing (Chapter 2). Here I present a novel method allowing better simplifications of data from hyperspectral acquisitions (instead of RGB color), together with its validations and applications.
3. SVBRDF acquisition for monumental heritage (Chapter 3). I present in this part the constraints linked to performing digital acquisitions in an uncontrolled environment, and the implementation of a process to measure SVBRDF models for a more realistic representation of objects in the context of monumental heritage.

Context

Reuse in 19th-century Cairo: Saint-Maurice residence



Figure 1: One of the challenge for the 3D restitution of Saint-Maurice residence (**upper-right**) is the collection and aggregation of multiple documents in multiple formats: colored photographs (**upper-center**, 1908 or 1926 by Jules Gervais-Courtellemont – National Geographic Society, Washington, D.C., US), photogrammetric 3D survey (doors at upper-right, fall 2019 – French Embassy, Cairo, Egypt), architectural blueprints (**lower-left**, Archives rapatriées des postes diplomatiques et consulaires, Nantes, France), watercolors of cross-section (**lower-center**, 1921 – location unknown), black and white photographs (**lower-right**, 1875-1879 by Beniamino Facchinelli – Bibliothèque de l'Institut national d'histoire de l'art, Paris, France).

Islamic revivalism at Saint-Maurice residence. Although mostly associated with Antiquity and the Middle Ages, architectural reuse represents a widely disseminated practice in Europe and beyond, all throughout the nineteenth century. In Middle Eastern capitals such as Cairo and Damascus, salvage purchased from dilapidated or under-renovation sites was repurposed as rich material for flooring and dados (marble), doors and furniture (woodwork), paneling (ceramics), and ceilings (painted plaster) [211]. These architectural arrangements are identified nowadays as belonging to **Islamic Revivalism** [69]. It was also common practice to sell installations of World's fairs for future reuse, once the events were over: an Egyptian pavilion from the 1867 Universal exhibition was reinstalled as an artistic studio and residence near Paris [210]. Part of the *Pavillon des Indes*, presented by Great Britain at the 1878 Universal

Exhibition, survives today in the vicinity of Paris (Courbevoie).

Islamic reuse was promoted in Cairo by French expatriates in the 1870s, a major instance being **Saint-Maurice residence**. Built in 1875-79, the structure associated salvage, with replicas of historic ornament from Cairo’s monuments and Revival pieces designed on purpose. The structure was dismantled in 1937 but partly reconstructed in another location in the same city. Although the building does not survive in its initial form, dispersed remains do exist, as well as a extensive documentation (drawings, plans, photographs, texts) on the initial project (c.f., Figure 1).

The restitution. Therefore, the aggregation of visual and material evidence permits to provide with color (and more) historic views, which are mainly monochrome, while offering the possibility to simulate a large decorative scheme that no longer exists in its original condition.

Such a restitution [86] is a challenge since it extends the notion of 3D puzzle (e.g., [166]) with documents in multiple formats, together with existing salvage scattered in distinct locations, either originals or copies, in modified form or not. It requires large surveys to collect and complete the required documentation. By putting together all available evidence, the 3D representation can help exploring some hypotheses regarding ornament reuse (see Figure 2).



Figure 2: *Two views of the 3D restitution of Saint-Maurice residence compared with historic photographs. **Left** – the main lounge decorated all the ornaments we had acquired. **Right** – an external view.*

The 3D restitution made by *Archeovision* [1] was difficult due to the available fragmentary documentation: while there was abundant historic iconography, there were no rigorous architectural plans nor detailed surveys of ornament. In brief, the shortage of detailed architectural data obliged to make full use of the plans, surveys, and photographs at their disposal.

Going beyond classic 3D. This colored 3D model leads to a visual restitution of the richness of the studied object. However, there is a need to go beyond classical 3D survey and closer to physical characteristics in order to formulate hypotheses on salvage reuse. Moreover, the metrologic aspect of existing technologies and models [89, 62] are not sufficient to answer the art historical questions raised by such type of architecture: a visual model is not enough for characterization of the underlying physical, chemical and optical properties. Indeed, as for metamerism in color, multiple physical compositions can lead to the same visual appearance. This is the second challenge of the study: how to acquire information that links more closely the visual aspect with the physical characteristics, in the limited field of study. Moreover, this acquisition has to be done on-site, in a non-invasive way. **This consideration is the central element that guided me during my doctorate, and that is presented in this manuscript.**

The Iznik ceramics of Saint-Maurice residence

Glazed ceramics represent a major part of the cultural heritage visible in the archaeological record. They are relatively simple to produce and have been found throughout the world from the Mediterranean basin, to as far as China and Japan for millennia. Due to its widespread and varied use (including dishes and other utilitarian items, both interior and exterior architecture, etc.) and its resistance to post-depositional degradation, it is one of the common material cultural items found during archaeological excavations. Initially applied to simple terracotta goods, the use of glazes first appeared in ancient Egypt. Once vitrified after firing, the glaze coating applied to ceramic surfaces serves to not only decorate the item, but also harden it, and make it water-proof [148, 202]. Though glazes were originally produced from ash, vitrification methods have evolved over time to also include the use of alkaline and lead glazes.

A particular ceramics is studied in our context: **Iznik ceramics**. These ceramics were produced between the 15th and 17th centuries for various uses, from tableware to tiles [16]. Coming from the city of Iznik in western Anatolia, they dispersed throughout the Ottoman Empire, the Arab world and the Mediterranean region [17].

From a technical point of view, they are composed successively from the bottom to the top of a siliceous substrate (Figure 4a), then of a layer of preparation support of the decoration also siliceous and of white color (usually called siliceous engobe, b). Afterwards a polychrome decoration based on metal oxides (generally red, black, blue, green, or even yellow pigments) is placed (c). Finally, we can find a layer of superficial glaze that covers the surface of the piece (d) [26].




Figure 3: *Schema of a section plan of the successive strata of a glazed ceramic. Depending on the type of pigment used, the polychrome decoration (c) can diffuse in the transparent glaze (e.g. blue or green pigments).*

The tiles surveyed in Saint-Maurice residence present 4 patterns characteristic of the Iznik production that can be found on many historic buildings in present-day Cairo [17] (see Figure 4). We named them C1, C2, P1 and P2 according to their ornamentation for better recognition. It then raises open questions about their origins and reuse while very few documentation exist for answering.

Three aspects of the ceramic tiles can be examined, namely their micro-scale shape, patterns, and pigments. Firstly, the analysis of the micro-scale geometry can reveal details on the granularity of the pigments or information on the glaze, while photogrammetry only generates a flat surface. Secondly, comparison of the patterns is an effective way to identify differences between different tiles. Although the tiles may appear similar at first glance, closer inspection can reveal significant variations in the pattern (look at the details in Figure 4). One way to study patterns is to segment an image of the tile based on its "color", from its RGB to spectral definition. Thirdly, the observation of pigments can facilitate both identification and similarity assessment. To achieve this, the reflectance spectrum of the pigment can be observed for identification purposes. As opposed to Gonizzi Barsanti *et al.* [20], the availability of on-site spectral imaging (e.g., Specim IQ) leads to the potential identification of the pigments used. To sum up, to identify these Iznik ceramic provenance, there is a need to acquire data beyond classic 3D method and RGB imaging.




Figure 4: Pictures of the 4 patterns of Iznik tiles visible in Saint-Maurice residence (**upper line, left to right**: groups C2, P1, C1, P2, c.f. Figure 5). Pictures of Iznik tiles presenting the same pattern in historic Cairo (**bottom left**), in Coptic Museum of Cairo (bottom - center), and in V&A Museum in London (**bottom right**)

The challenge here is to **seek the possibility of highlighting groups**, or identifying somehow clear tendencies within the 4 groups of ceramics sharing similar properties. Figure 5 shows the photos of these groups of tiles and their respective indexing. In this thesis, we will use the same indexing as chosen here.

Hyperspectral imaging of ceramics seems to be an interesting, fast and, non-invasive method to achieve such goal according to conclusions of Farjas *et al.* [66]. The hope is that it will be used to perform pattern segmentation as well as pigment recognition and semi-automatic classification. In addition, developing a 3D acquisition technique at the micro resolution for the glaze and micro details information is a very appealing method to enhance and validate our results.

Note: Acquisition campaign in Cairo

To achieve these purposes, I went to lead an acquisition campaign with Pascal Mora (*Archeovision*), Xavier Granier (*IOGS*), Mercedes Volait (*INHA*) and Matjaz Kacicnik (professional photograph) at the French Embassy of Cairo in November 2020 [19]. During 4 days, I carried out 3D macro scale acquisition (see Chapter 1) and hyperspectral imagery on the Iznik ceramics (see Chapter 2), and appearance measurements on a mosaic wall and a painted ceiling (see Chapter 3).

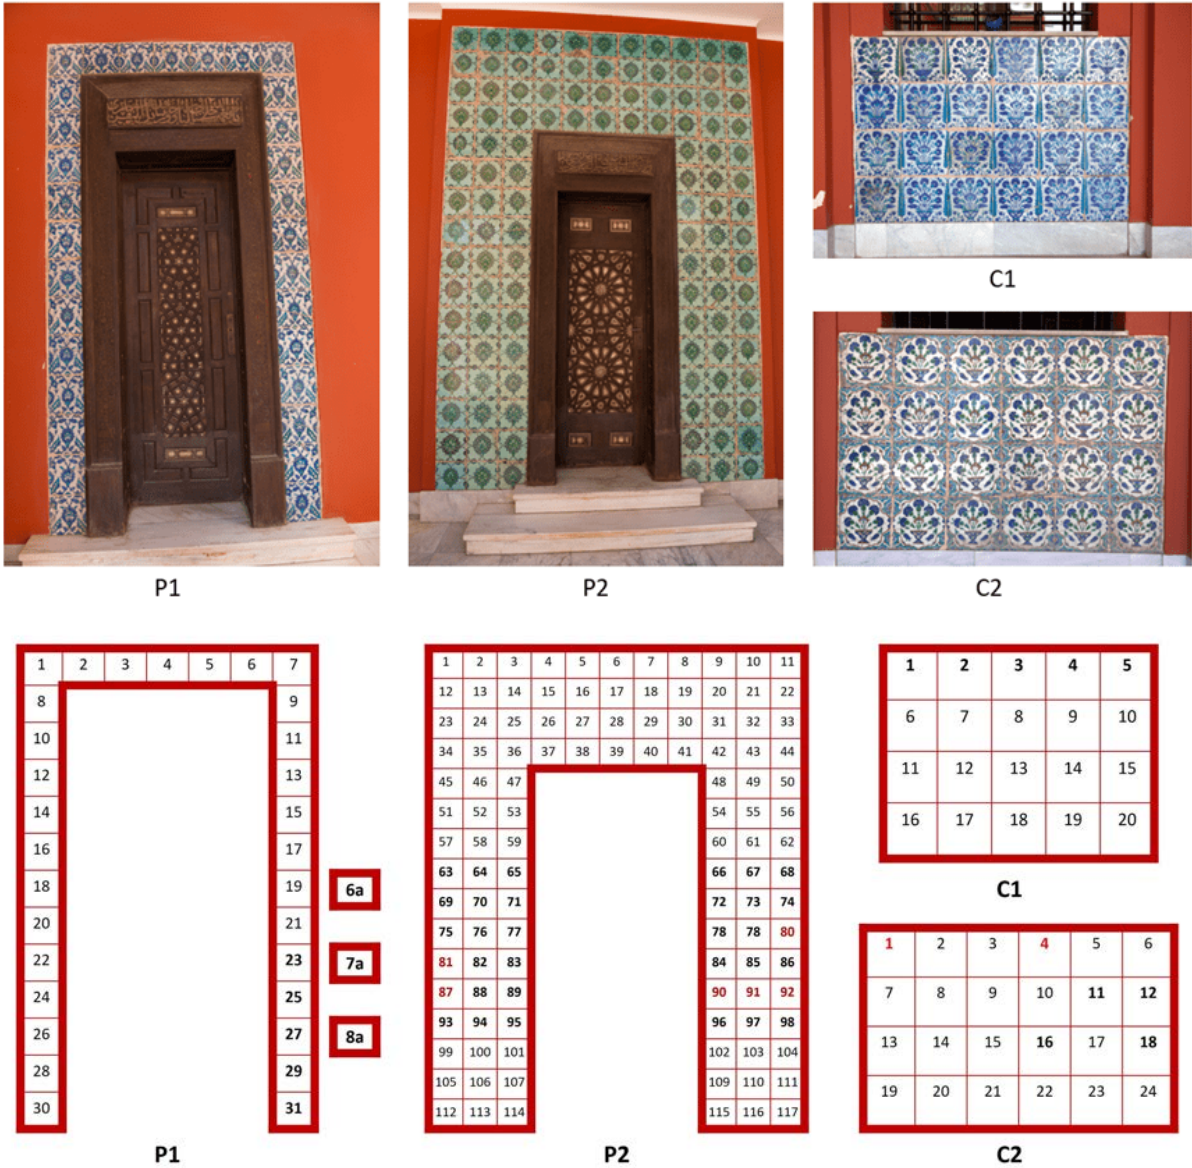


Figure 5: Top: Photographs of the Iznik ceramics tiles in the Saint-Maurice residence. Each of the group has a name P1, P2, C1, C2.

Bottom: Indexing of each group corresponding to the photographs. 3 tiles of P1 group are not visible on the photographs. Tiles measured with Depth from Focus are in red (see Chapter 1). Tiles photographed with hyperspectral imaging are in bold (see Chapter 2).

Chapter 1

Depth from Focus for Shape acquisition

As presented in the introduction, a computer-based reproduction of an object’s realistic appearance can be divided into two components: shape and appearance, including its color and the way it reflects light. The current chapter focuses on the first of these aspects: shape acquisition. Specifically, we present a method of Depth from Focus (DfF) that enables us to acquire shape information at small scales. Although a large number of 3D acquisition methods exist either for large-scale or microscopic objects, the literature lacks approaches capable of capturing details at the intermediate level. This gap is particularly significant when dealing with details that fall between the scale of microscopy and our own. To address this challenge, we detail the reasoning behind the use of a depth from focus strategy to acquire shape information. Moreover, we describe the improvements we made to the existing Depth from Focus methods for our specific application. Finally, we present an original extension of our approach allowing to measure the thicknesses of ancient ceramic glazes in a non-invasive manner. We specifically applied this new tool to the case of the Iznik ceramics case study from Saint-Maurice residence.

The results of the chapter have been validated in [49, 48].

1.1 3D acquisition of small reliefs - state of the art

1.1.1 On-site 3D acquisition techniques

Capturing the 3D structure of real world scenes is a long standing problem in the computer vision and computer graphics communities. The literature is extremely vast and the technological solutions are very diverse. In this section we focus on approaches potentially applicable to our case: portable high-resolution scanning (we want a resolution of 5-10 μm). Therefore we have deliberately ignored some technologies such as LiDAR remote sensing or time-of-flight cameras which are nevertheless widely used but clearly unsuited to our needs in terms of accuracy. The 3D acquisition techniques that we will see below are all based on the capture of 2D images (in the broad sense) obtained by varying either the parameters of the source, or the parameters of the sensor, or even both [177]. Shape information can then be reconstructed in different ways.

Before discussing about the different techniques of 3D acquisition, it remains important to underline the limitations caused by the optical systems. The major difficulty in macro and microscopy photography is that the depth of field relative to the field of view decreases drastically with magnification.

In order to maximize precision, it is necessary to acquire the sharpest images possible, and therefore to eliminate all sources of blur. Aside from motion blur which is irrelevant here, we have two main sources of blur: depth of field and diffraction. The width of the Airy disk can be approximated by: $2.44\lambda N$, with λ the wavelength and N the f-number. Taking into account the Bayer filter and the presence of a low-pass filter, it is commonly accepted that the diameter of the Airy disk is generally less than 3 pixels, which is acceptable for our need. So there is no problem here. Then, the depth of field (DoF) for an asymmetrical lens is defined by:

$$DoF = \frac{2f(\frac{1}{m} + \frac{1}{p})}{\frac{fm}{Nc} - \frac{Nc}{fm}} \quad (1.1)$$

with f the focal length, m the magnification, N the f-number, c the diameter of the circle of confusion (CoC), and p the ratio between the incoming and outgoing pupil (equal to 1 for a symmetrical lens). For sharp images at the pixel scale, the diameter c of the circle of confusion is directly constrained by the physical size of the pixels, and we can take it to be $2\times$ the size of a pixel. This confirms that the depth of field relative to the field of view is really related with magnification. For example, for a field of view of 1m, the depth of field is approximately 30cm, or 30% of the field of view. For an enlargement of 1:1, with a field of view of 24mm, the depth of field is 0.3mm, i.e., 1% only of the field of view.

This phenomenon is by far the main difficulty in 3D capture at the micron scale, regardless of the technology used.

3D laser scanning. 3D laser scanners are the most widely spread and known portable acquisition methods. This technology, greatly democratized over the period 1990 to 2010 [32], is based on a sensor (camera) and a structured source (laser) whose intrinsic and extrinsic properties (relative positions) are perfectly known via a calibration phase. The relative 3D position is obtained by simple triangulation, i.e. based on the pairing between two (or more) rays reaching the 3D physical point from different viewpoints established between a pre-calibrated transmitter-sensor couple. The main benefit of the laser source is that it has a coherent light source, which makes it possible to concentrate a beam on a very fine area. The beam can be modified into a line or a cross using some lenses, which allow to multiply the triangulated pixels by a single capture and greatly accelerate the acquisition of a portion of surface by scanning the beam in translation or rotation (see Figure 1.1). By using an ultraviolet source, they can also model a very large set of materials, including very specular metals. Regarding the resolution and precision that it is possible to obtain, this depends on the one hand on the fineness of the laser beam and the effective resolution (GSD) of the associated sensor(s).

However, they have several limitations. First, they only capture shape and not photometric (color) information. This is a major limitation because realigning two different sources to the precision of a few microns is a real challenge. Then, it can produce small occlusions between the laser and the sensors leading to measurement holes. Last but not the least, most recent “hand” scanners such as *CreaForm HandySCAN 700* are limited by their resolution at 50 μm theoretically, and around 100-200 μm in our tests. One of the main reasons for this is that the

“hand” scanners offering the highest resolutions are designed for fast scanning of objects ranging from 10cm to a few meters. This resolution can be greatly improved with laser profilometers such as *Keyence LJ-V7020K* or *Micro-Epsilon scanCONTROL 29xx-10/BL* which represent a resolution of few μm . But the drawback of this resolution is the acquisition time because they are designed for a pixel-to-pixel acquisition at microscopic scale. This would require a device that would allow the system to be translated slowly, but the acquisition times would be huge, and the risks of an error accumulation during processing would not be negligible.




Figure 1.1: Simple example of a 3D laser scan using a laser beam modified into a line. Images taken from [32]

Pattern Projection Scanning. A laser source permits to capture only a single line (or a few lines) of dots at once. An alternative approach consists in replacing the laser source by the projection of 2D patterns allowing the reconstruction of the depth of all the pixels of the sensor with a low number of shots. The methods differ mainly by the type of patterns used, which lead to different constraints on the sources and to very different algorithms.

We can mention first, structured light 3D scanners that project a light pattern onto the object and reconstruct depth information by analyzing deformations. This challenge is addressed by a joint use of different patterns and sophisticated algorithms, the literature being extremely vast on the subject [189]. The simplest techniques project binary patterns with decreasing frequency generating for each of the pixels of the camera a code from which the depth can be reconstructed by triangulation. Intuitively, low frequencies allow disambiguation, while high frequencies allow reconstruction of dense depth information. In order to improve the results, the joint use of non-binary and/or colored patterns makes it possible to further reduce the number of captures necessary while increasing the precision [190, 189]. This method has two main shortcomings. First, the quality and robustness of the reconstruction in structured light is conditioned by the reflectance properties of the objects. For example, areas that are too strongly specular (e.g., mirror) are generally very poorly captured because the projected pattern is unlikely to be reflected towards the camera [92]. Second, and most importantly, the accuracy is greatly limited by the resolution of the projector and requires a perfectly sharp projection on the surface.

The limits of structured light tend to be corrected with phase shift techniques (Fringe projection) [82]. The ancestor of this method is the projection of Moiré patterns [95] which was abandoned in favor of the projection of fringes thanks to the appearance of video projectors. It works by replacing the binary patterns with sinusoidal bands on the object successively with a phase shift of the sinusoidal signal. The images thus captured are analyzed to find the phase shift in each of them. In exchange for a very precise calibration, three images only are enough to solve this problem, and we can use a projector with a very low resolution while allowing sub-pixel precision (50 to 100 μm)[197].

Like 3D laser scanning, light projection methods have crippling flaws in our case. In addition

to the blackouts between the projector and the camera(s) leading to measurement holes, it is necessary to project high resolution images with sharp patterns on the surfaces to be captured. However, in our case at the macroscopic scale, the reduced depth of field makes it difficult to acquire areas with strong protuberances. This remark concerns both the camera(s) and the projection device. It becomes difficult to go below a resolution of the order of tens of μm .

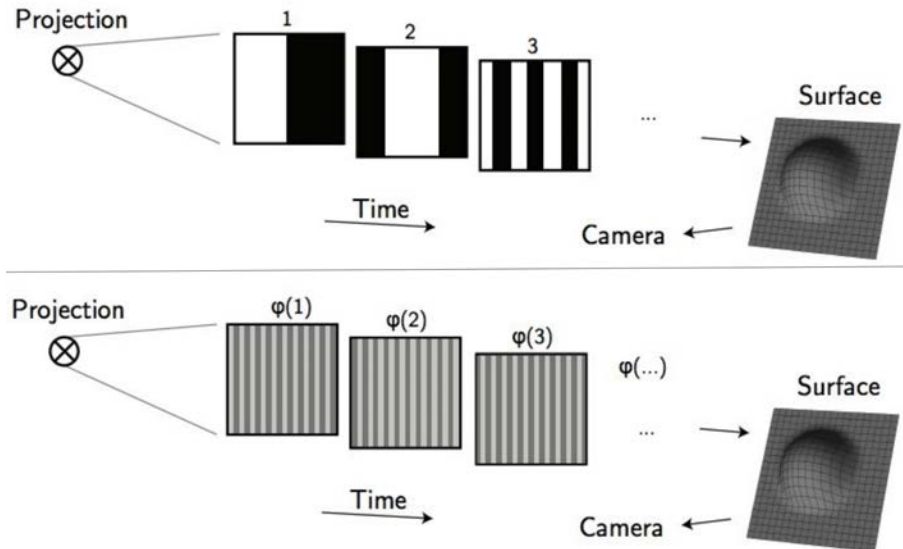


Figure 1.2: Illustration of the principle of digitization by pattern projection with binary patterns (*top*) and continuous patterns for fringe projection (*bottom*). Images taken from [220]

Shape from Shading. A last class of technique approaching the triangulation between a sensor and a light is put forward here. The subject is photographed from a single point of view but under different lighting conditions, typically from three or four different directions of light [217, 157]. By knowing the lighting conditions and the reflectance properties of the object, it is then possible to find in each of the pixels the slope (i.e., the normal) of the underlying surface directly from the different intensities of the pixel. Another algorithm like [111] then makes it possible to find the depth (within a constant) of the pixels from the slopes. Such a device is very simple to set up, and requires very little hardware, but the main limitation of this approach is that the reflectance of the object must be known and constant over the entire surface. Moreover, it requires to separate direct and indirect lighting for good results [156],



Figure 1.3: Illustration of the principle of shape-from-shading and application to a coin. Images taken from [94]

which is not very convenient in many cases. This restricts this technique to a few specific case studies such as the digitization of coins, for example [94] (see Figure 1.3).

Photogrammetry. Unlike the methods mentioned above, photogrammetry [138, 178] is a passive method that does not require special lighting. It is based on the triangulation not between a source and a camera, but between the positions of several cameras. By multiplying the sensors, sophisticated image processing algorithms make it possible to establish correspondences directly from the photometric data without relying on a particular source (Multi-View-Stereo - MVS) [191]. Other complementary algorithms even make it possible to dispense with the calibration step (Structure-From-Motion - SfM) [216]. Most of the recent photogrammetry methods are based on a principle of joint approach of both techniques. Firstly, SfM photogrammetry automatically determines the intrinsic and extrinsic parameters of a set of images captured by photographing an object from a myriad of view-points, using low-density point cloud of the object. Then, a dense point cloud is generated by the triangulation of rays casted from each of camera and a scene geometry is reconstructed incrementally. It is even possible to improve the obtained results by assuming that the object is perfectly diffused for the update of the geometry using estimated surface normals [165]. The more the number of views, that is the image set, increases, the more the precision on the measurement of the shape progresses. The results can reach the precision of 3D laser scanners for thousands of photos. Moreover, one of the greatest interests of photogrammetry remains the possibility of having the color of each point of the cloud by taking up that of the pixels of the images.



Figure 1.4: Image taken from *RealityCapture Software* by the website *3D Scan Expert*. Illustration of the 3D object acquired and the positions get by SfM for each camera.

However, this technique still has important limitations. First, passive lighting makes photogrammetry very constrained to material variations on the surface. Indeed, by matching points on different images, there is a big loss of robustness in the uniform and textureless areas. In addition, parts that are too strongly specular (e.g., mirror or metal) are generally very poorly captured because the very large color variations on the pixel make the triangulation obsolete.

Above all, it is also necessary here to process sharp images for triangulation. However, we are faced with a reduced depth of field in the case of close-range photogrammetry for macro

scanning. This limitation makes it difficult to obtain detail below 50-100 μm . It is nonetheless possible to overcome it by reconstructing sharp images with focus stacking (a technique which combines a set of images taken at different focus distances to generate an image with a larger depth of field [176]). The next step is to apply photogrammetry algorithms to the all-in-focus images to generate a 3D object [175]. The issue with the proposed approach is that it entails longer acquisition and processing times. Specifically, about fifty photos must be taken for each points-of-view, which necessitates the use of a fixed tripod. However, this presents a challenge in photogrammetry, as certain angles may be difficult to achieve without manual manipulation of the camera. Additionally, each set of pictures must be processed individually to generate all-in-focus images before further processing can be performed using the SfM-MVS approach.

Microscopy. If each of the proposed methods therefore works very well at the human scale, we understand that they all have the same problem related to the sensor: the very shallow depth of field prevents precise acquisition to the μm . On the other hand, there are acquisition solutions at much smaller scales, as we have seen with laser profilometers. Other systems thus exist by adapting microscopy optics to Fringe projection type methods [103]. But the problem nevertheless remains that with microscope optics, the area acquired is really small (in the case of [103], the acquisition zone is $8\times 6\text{mm}$ for an accuracy of 8 μm): we measure at a scale below our target that which we wish, and only on small objects transportable in the laboratory.

In summary, our objective is to perform measurements that bridge the gap between microscopic scale and traditional portable 3D acquisition methods. Unfortunately, techniques that rely on microscopic optics are too limited in their scope, while conventional techniques are limited by the reduced depth of field. This feature is a limitation in triangulation cases where sharp images are a requisite. But it is possible to use the information contained within the depth of field as an advantage to extract depth information.

1.1.2 Exploiting the depth of field information

A wide range of techniques have been proposed so far such as photogrammetry, laser scan or structured lights. But as explained in the previous part, when targeting micro-scale reliefs, such triangulation based optical systems become impracticable because the depth-of-field decreases quadratically with the field of view. For instance, when targeting a resolution of one pixel per micrometer with a classical digital camera, the typical depth of field will range between 5 and 50 micrometers only.

In this context, depth-from-focus methods become particularly appealing. Those techniques strive to recover 2.5D depth information from a focal stack by estimating, for each pixel, the depth of the focus plane. Most depth from focus methods start by applying a focus measure on the input focal stack, and then estimate the depth by seeking, for each pixel, a robust peak of this focus measure profile along the z direction. Most techniques manage to recover a discrete depth information with reasonably high accuracy, and the relative depth accuracy of the reconstruction is thus expected to increase as the field of view decrease. However, their performance at recovering continuous depth information remains unclear yet.

Depth from defocus. Exploiting depth of field information to reconstruct depth information was first proposed by Pentland [169]. Their *Depth from defocus* methods estimate the depth from the amount of defocus estimated from a pair of images taken at different focal lengths, and many variants of this approach have been proposed so far [120]. Using a *coded aperture*, Levin et al. [127] even showed how to get rid of the back and front ambiguity when working on a single image. However, since it is extremely difficult to estimate the amount of defocus independently of the texture frequency of the scene itself, those *Depth from defocus* methods

can only recover very coarse depth information that can only be used for simple tasks such as autofocus, foreground/background segmentation and the likes, but unlikely for 3D scanning.

Depth from focal stack. A focal stack is a sequence of images taken for different focal planes. It can be obtained by either moving the camera along a translation unit, or by changing the focal length of the optics. Usually, each point of the scene is in focus in one and only one slice such that an *all-in-focus* image can be reconstructed by peaking the sharpest pixel along each view direction [171]. A discrete depth map can be jointly obtained by assigning to each selected pixel the depth of its respective focal plane [88]. Depth from focal stack can then be used alone to determine depth, or in conjunction with Depth from Defocus and other methods [123].




Figure 1.5: A coarse focal stack of an outdoor scene and its surface mapped 3D depth obtained by depth from focal stack is shown from two different viewpoints. Image taken from [187].

Focus measure. In such *depth from focus* approaches, a central ingredient is the so called *focus measure*, whose goal is to estimate the level of sharpness of every pixel of the stack. Dozens of heuristics have been proposed, either based on differential operators, contrast estimator, or frequency analysis [195]. According to Pertuz et al [170]’s study, Laplacian-based operators turned out to be the best performing. This is confirmed by more recent works improving upon classical Laplacian operators such as the Ring Difference Filter (RDF) [196]. Other methods tried to learn optimal combinations of different measures [141, 187].

Depth estimation. Simply taking the extrema of the focus measure to recover depth information [88] is fast but subject to high noise and limited to discrete depth values. Higher accuracy and robustness is achieved through fitting methods using either local Gaussian [31] or global Laplacian [187] distributions. Some recent work have also experimented with neural networks for this task [99], but they face the hunger for data for the learning stage. Nonetheless, the raw output of such estimators still have to be post-processed to reduce the noise and complete incorrect parts, for instance using cost-aggregation techniques [196] or cost-volume methods based on a reliability map [187]. Javidnia et al. [108] globally optimize both the depth and

in-focus color of each pixel by minimizing the difference between the input reference stack and a procedurally generated stack from the current depth map. It therefore seems interesting to set up a method based on the exploitation of the depth of field, seeking to find an uncertainty less than the distance between two images of the stack.

1.2 Enhancing the Depth from Focus accuracy

In this chapter we propose a novel peak estimation method that strives to address this shortcoming. This chapter makes the following four key contributions:

1. *Sliding window.* To improve robustness to outliers while avoiding costly non-linear optimizations, we propose an original scheme that linearly scans the profile over a fixed size window, searching for the best peak within each window (Section 1.2.2).
2. *Linear least-Squares Laplace regression.* Within each window, we identify the peak by fitting a Laplace distribution over the data. We greatly improve the computation cost by proposing a linearization scheme (Section 1.2.2).
3. *Adaptive smoothing.* As a post-process, we present an adaptive and feature preserving smoothing scheme based on a per pixel confidence value and Moving-Least-Squares like regressions (Section 1.2.3).
4. *Quantitative evaluation.* We objectively evaluate the performance of our method by generating synthetic focal stacks from which the reconstructed depth maps can be compared to ground truth (Section 1.2.4).

1.2.1 Overview of the pipeline

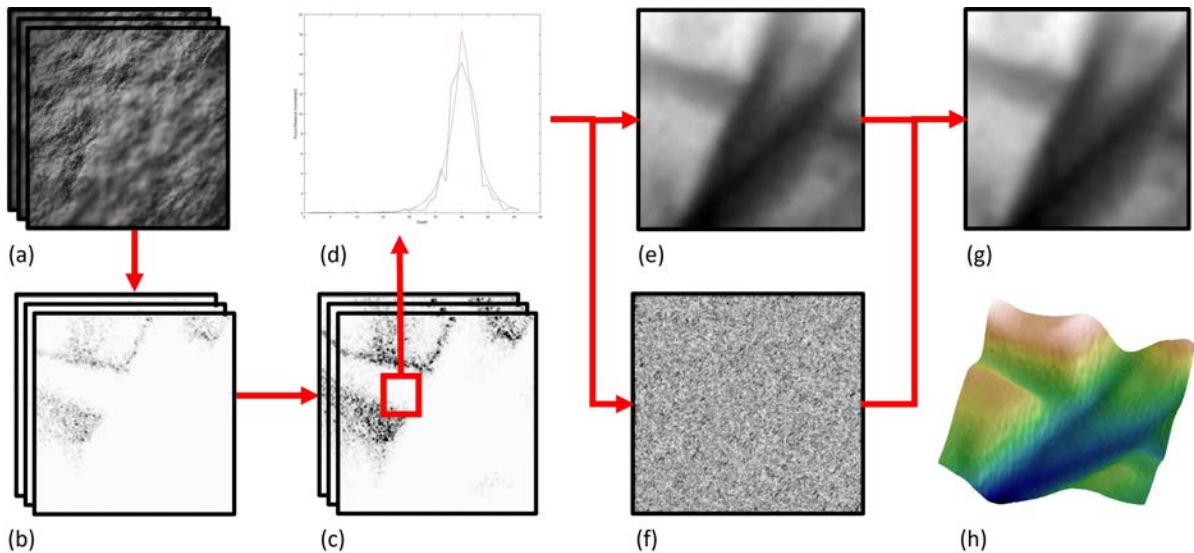


Figure 1.6: Our depth from focus pipeline. Starting from the acquired focal stack \mathcal{I} (a), we compute a Focus Measure map \mathcal{F} (b) and a low-pass version $\tilde{\mathcal{F}}$ (c) through Gaussian smoothing. For each pixel, we perform our Peak Estimation procedure (d), yielding a depth map \mathcal{D} (e) and a reliability map \mathcal{R} (f). Finally, we apply an adaptive MLS filter from these two maps to obtain our final depth map (g) giving a 3D profile (h).

As most depth-from-focal-stack methods, our approach relies on the three main steps identified in the section 1.1.2: focus measure, peak extraction, and post-processing. Our pipeline is

depicted in Figure 1.6. More precisely, we start from a stack \mathcal{I} of n images where each pixel is denoted $\mathcal{I}_{i,j}^k$, with i, j the pixel indices within the slice number $k \in [1, n]$, each slice corresponding to a focal plane depth z_k . We assume that the focal planes are uniformly sampled, with a constant sampling interval $\delta_z = z_{k+1} - z_k$. Each slice of the stack is also assumed to be pre-aligned such that for each spatial indices i, j , the 1D profile $\{\mathcal{I}_{i,j}^k\}_k$ corresponds to the same point of the scene.

Then we apply a Focus-Measure (FM) independently on each pixel of the stack yielding a FM stack \mathcal{F} (Fig. 1.6b). The rest of our pipeline can work with any FM producing sharp profiles resembling to a Laplace distribution (Fig. 1.6d). This is the case of all Laplacian-like differential operators, including the recent composite FM [187]. From our experiments, we found the 5×5 Ring Difference Filter (RDF) [196] to be the best performing, especially in areas of low contrast, while being simple and fast. To reduce the sensitivity to spatial noise, we apply a 2D Gaussian filter on each slice of the FM stack, yielding a smoothed version denoted $\tilde{\mathcal{F}}$ (Fig. 1.6c).




Figure 1.7: Demonstration of the importance of the Gaussian filter. **Left:** curve of FM of one pixel as a function of depth z , on an FM where the Gaussian filter is not applied. **Right:** curve of FM of the same pixel as a function of depth z , on an FM where the Gaussian filter has a width of 20. For each case, the blue curve is the acquired FM and the orange curve is the Laplacian distribution fitted on the FM.

The crucial step then consists in extracting from $\tilde{\mathcal{F}}$ the depth $\mathcal{D}_{i,j}$ of each pixel (i, j) . In our approach, this step is carried out independently for each pixel (i, j) , so that it boils down to a continuous peak estimation problem from the discrete 1D signal $f_k = \tilde{\mathcal{F}}_{i,j}^k$ (Fig. 1.6d). This is the main part of our contribution and it is detailed in Section 1.2.2. In addition to the depth map \mathcal{D} (Fig. 1.6e), in this step we jointly compute a reliability map \mathcal{R} (Fig. 1.6f) that is used to drive an adaptive denoising procedure detailed in Section 1.2.3 (Fig. 1.6g).

1.2.2 Peak estimation using Linearized Sliding-window

This section focuses on computing, for each pixel (i, j) , its depth value z from its 1D discrete focus measure profile f_k . In the following, it is assumed that z can be recovered as the location of the maximum value of some continuous reconstruction \tilde{f} of f_k , that is: $z = \operatorname{argmax}_x \tilde{f}(x)$. Sakurikar et al. [187] showed that the Laplace distribution

$$L_{\mu,b}(x) = \frac{1}{2b} e^{-\frac{\|x-\mu\|}{b}} \quad (1.2)$$

was a good proxy model to reconstruct \tilde{f} for a variety of focus measures. They compute the mean μ and average absolute deviation (AAD) b parameters through the global minimization

of the following non-linear energy:

$$E_0 = \sum_k (L_{\mu,b}(z_k) - f'_k)^2, \quad (1.3)$$

where f' is the normalized version of the f_k so that it integrates to 1. Then, we directly get the sought for depth as $z = \mu$, and the AAD value b , or rather its inverse b^{-1} , gives us an estimation of its reliability: the finer and higher the peak of the distribution, the sharper and more reliable it is.

This non-linear regression procedure exhibits two major issues. First, it involves an iterative and computationally expensive minimization. Second, since the regression is carried out on the whole signal in a simple least-squares sense, it is highly sensitive to outliers and noise far away the expected peak. Here we are referring to outliers and noise that might be produced by the focus measure, and not directly to outliers that would be present in the images themselves. For instance, in the presence of strongly visible *bokeh* artefacts, the focus measure signal f_k might exhibit two peaks leading to an estimated depth located in-between the two peaks.

We alleviate both issues through linearization (Paragraph *Linearization as a Triangle function*) and localized regressions (Paragraph *Sliding-window estimation*) respectively.

Sliding-window estimation




Figure 1.8: Illustration of our sliding window strategy. The FM profile is shown as blue crosses and three position of the window are depicted in yellow, red, and green respectively. The window size n_w is constant, and each window starts at a position z_k .

To deal with outliers, we localize the Laplace regression through a basic sliding window strategy described in Algorithm 1, and illustrated in Figure 1.8. The idea is to perform a full peak estimation (e.g., through a Laplace fitting) independently for each sub range $\{f_k, \dots, f_{k+n_w-1}\}$ of the stack, with n_w the width of the window. Depth estimations falling outside the current window are rejected, and finally we retain the one associated with the best reliability.

Algorithm 1 Sliding window algorithm

```
 $r \leftarrow 0$  ▷ reliability  
for  $k \in [1, n - n_w + 1]$  do  
   $\mu, b \leftarrow \text{Laplace\_regression}(f[k, k + n_w - 1])$   
  if  $b^{-1} > r$  and  $z_k < \mu < z_k + n_w \delta_z$  then  
     $r \leftarrow b^{-1}$   
     $z \leftarrow \mu$   
  end if  
end for  
return  $z, r$ 
```

Linearization as a Triangle function

Running a non-linear regression on each sub-window would be prohibitively expensive. To speed up computation, we thus linearize the energy E_0 by taking the logarithm of both the focus measure f and the Laplace distribution. However, doing so directly on the Laplace distribution yields a triangle function of the form: $\log(\frac{1}{2b}) - \frac{\|x-\mu\|}{b}$, which unfortunately still exhibit a non linear relationship between the slopes and constant terms. To address this issue, we relax the somewhat arbitrary constraint of L being strictly a probabilistic distribution by making its amplitude $\frac{1}{2b}$ in Eq. 1.2 a free parameter, say a , yielding a triangle function with three degrees of freedom of the form: $\log(a) - \frac{\|x-\mu\|}{b}$. This equation is advantageously rewritten as two piecewise linear functions:

$$T_{\alpha, \beta, \gamma}(x) = \begin{cases} \alpha \cdot x + \beta & \text{if } x < \mu \\ -\alpha \cdot x + \gamma & \text{if } x \geq \mu \end{cases}, \quad (1.4)$$

where the position μ of the line intersection and AAD parameter of the initial Laplace distribution are recovered by:

$$\mu = \frac{\gamma - \beta}{2\alpha}, \quad b = \frac{\gamma + \beta}{2}.$$

To efficiently deal with the absolute value, or piecewise nature of T , we first enforce the window size n_z to be an even number, and then assume that μ is located somewhere in-between the two slices surrounding the window center $z_c = (z_{k+n_w/2-1} + z_{k+n_w/2})/2$. This way we can partition the current window's samples into two groups, $W_1 = [k, k + n_w/2 - 1]$ and $W_2 = [k + n_w/2, k + n_w - 1]$, each group being associated to one side of the triangle.

The three parameters α , β , and γ are then easily obtained by minimizing the following quadratic least-squares energy:

$$E_1 = \sum_{k \in W_1} (\alpha z_k + \beta - \log(f_k))^2 \quad (1.5)$$

$$+ \sum_{k \in W_2} (-\alpha z_k + \gamma - \log(f_k))^2. \quad (1.6)$$

Let us emphasize that thanks to the sliding window strategy, our *a priori* partitioning with respect to the center of the window is not really an issue because all possible locations of this cut will be tested anyways. Nonetheless, nothing in E_1 prevents the two lines to cross at a position μ within the two slices surrounding z_c . Whereas this constraint could be achieved through a pair of linear inequalities, this would not be desirable as this would simply snap the eccentric estimations on discrete values. Instead, we rather tighten the acceptance test in Algorithm 1 to accept only the candidate μ within the sub range centered on the window and of width $3\delta_z$. We chose an acceptance width of $3\delta_z$ to be tolerant to slightly eccentric estimates to prevent the risk of entirely missing a peak.

In the rest of this section, we will refer to this method as the *Triangle* method. Let us note that this linearization would not be possible without our sliding window strategy, and inversely, our sliding window strategy would be prohibitively expensive without our linearization.




Figure 1.9: *Illustration of the Triangle method. In log space, the current window [28, 31] is partitioned within two parts. Each part is approximated by a piece-wise linear function having opposite slopes.*

1.2.3 Post-correction using MLS

From the depth map \mathcal{D} and reliability map \mathcal{R} , we apply an adaptive low pass filter in order to both denoise the depth map and fix low reliability values from its surrounding. To preserve the depth map features, we designed a filter inspired by the Moving Least Squares filter [35]. For each pixel (i, j) , the idea consists in approximating its neighborhood by a bivariate polynomial g , and replace its depth value by the value of g at the pixel location $\mathbf{q}_{i,j}$. We take into account the reliability map \mathcal{R} by introducing it as additional weights within the regression. More formally, the polynomial g is obtained as the minimization of the following quadratic energy:

$$\sum_{x,y \in N_{i,j}} \mathcal{R}_{x,y} \theta(\|\mathbf{q} - \mathbf{q}_{x,y}\|) (g(\mathbf{q}_{x,y}) - \mathcal{D}_{x,y})^2, \quad (1.7)$$

where θ is a compactly supported weighting function of radius h , and $N_{i,j}$ denotes the neighborhood of the pixel (i, j) within a distance h . In our implementation, we used $\theta(t) = \left(1 - \frac{t^2}{h^2}\right)^2$, and quadratic polynomials g . Compared to a simple weighted average, the use of a degree two polynomial allows to much better preserve the features of the data. In the presence of uneven sampling or holes as implied here by the reliability map, such high order polynomials also allow to much better preserve the slopes, whereas a simple weighted average would be strongly biased toward the more reliable and more numerous depth values.

1.2.4 Experiments and Results

We implemented a prototype of our pipeline within Matlab. For the spatial smoothing pass of the focus measures, we used a Gaussian filter size of 7 pixels, and a MLS filter size of 5 pixels for the post processing pass. For the sliding window, we found $n_w = 6$ to be the best performing according to our experiments (see Figure 1.10).

Evaluation on synthetic stacks

To objectively evaluate the performance of our different contributions, we procedurally generated image stacks from several reference depth and texture maps. To this end, we implemented a




Figure 1.10: Average absolute error over 20 synthetic stacks as a function of the window size n_w .

simple splatting approach spreading the contribution of each pixel of the input depth and texture maps to each image of the stack as a Gaussian splat approximating the true Airy disk. The radius of the splat is adjusted with respect to the size c of the circle of confusion computed as:

$$c = A \cdot \frac{\|S_2 - S_1\|}{S_2} \cdot \frac{f}{\|S_1 - f\|},$$

with A the aperture, f the focal length, S_1 the distance to the focal plane, and S_2 the distance to the current scene point [176].

With this procedure, we generated 80 stacks from 4 base depth maps, each being perturbed with 5 levels of low frequency noise mimicking surface details. Some combinations are shown in Figure 1.12, first column. Those 20 depth-maps are then combined with 4 different grey-level textures shown in Figure 1.11. Each stack is formed by images of 400×400 pixels.

We compare our *Triangle* method against two other peak estimations borrowed from the literature. The first one, called *Laplace*, is the global non-linear regression of a Laplace distribution coming from the the method of Sakurikar et al. [187] and that we recalled in Section 1.2.2. Implementation-wise, we used the `lsqnonlin` Matlab’s function parametrized to use a Levenberg-Marquardt algorithm with low tolerance values (i.e, `TolX=TolFun=1e-4`) to make the multiple non-linear solves as fast as possible. The second one, called *3 pts*, comes from the three-points method of Billiot et al. [31]. Their method estimates the depth as the location μ of the peak of a generalized Gaussian $a \cdot e^{-\frac{x-\mu}{2\sigma}}$ whose three parameters are computed such that it passes through the sample having the maximal focus measure, and its two neighbors. Likewise, we compute the reliability as the height of the respective normalized Gaussian. We also added results from the RDF method [196] using the implementation provided by the authors, and the full pipeline of Sakurikar et al. [187] which is similar to the “Laplace” pipeline described above but using their *composite* focus measure and a bilateral filter instead of our MLS pass.

Tables 1.1, and 1.2 show global statistics for the 80 stacks of the *Mean Absolute Error* (MAE) computed against the reference depth maps for each of the three methods. The reported MAE




Figure 1.11: The four textures used to generate our 80 stacks.




Figure 1.12: Qualitative comparison of 3-pts, Laplacian, and Triangle methods on some generated image stacks.

is normalized such that it should be interpreted as if the distance between two focal planes δ_z is 1.

Table 1.1: Statistics of absolute error prior to MLS post-correction. The 4th to 6th columns report the worst error for the given percentiles.

Method	Mean	Std-dev	25%	50%	75%	Time
3 pts	0.360	0.354	0.135	0.288	0.497	0.6s
Laplace	0.219	0.256	0.075	0.163	0.292	90.1s
Triangle	0.196	0.257	0.065	0.141	0.259	1.5s

Table 1.2: Statistics of the absolute error and running time for the full pipelines (i.e., including MLS post-correction).

Method	Mean	Std-dev	25%	50%	75%	Time
3 pts	0.237	0.223	0.089	0.189	0.331	17.7s
Laplace	0.171	0.181	0.060	0.130	0.232	107.2s
Triangle	0.159	0.184	0.054	0.117	0.213	18.6s
RDF[196]	0.827	0.921	0.246	0.534	1.067	14.8s
Comp.[187]	0.236	0.414	0.072	0.157	0.288	94.5s

One can notice that the RDF pipeline produces the worst statistics: this is expected as it




Figure 1.13: Reconstruction of a very smooth initial 3D profile (*top-left inset*), before and after our MLS adaptive post-correction.

produces discrete depth values whereas all others produces continuous ones. Not surprisingly, Sakurikar et al. [187] pipeline exhibits statistics in par with our *Laplace* pipeline.

Now focusing on our three pipeline variants, we can see that the MLS post-correction step provides an average of a 20% to 30% reduction of the MAE. Its visual effect is depicted in Figure 1.13. Furthermore, all precision-related indicators lead to the same hierarchy, with the *3-pts* method being significantly less accurate than the other two peak-estimators, and our *Triangle* method being the most accurate. Overall, compared to the non-linear *Laplace* method, our *Triangle* method provides a 10% reduction of the MAE while being more than $\times 60$ times faster, and only three times more expensive than the simplest *3-pts* method.

Those 10% gains, however, does not tell everything about the quality improvement in general. In Figure 1.14 we report the repartition of the computed depth values for each of the three methods, assuming $\delta_z = 1$. Those histograms reveal a clear bias toward a discrete set of predetermined values corresponding to the discrete input sample depths for the *3-pts* method, and their middle for the other twos. Our linearized *Triangle* method exhibits a much lower bias though. In practice, this bias produces depth maps with *plateau*-like artefacts that can be observed in the results of Figure 1.12. Our *Triangle* result thus appears significantly less noisy even-though the MAE differs by 10% only with *Laplace*. From those observations, one might though about correcting this bias by applying to the decimal part of the estimated depths a precomputed transfert function. As show in this figure (right column), this systematic bias is well captured by a cubic polynomial function constrained to pass trough $(0,0)$ and $(1,1)$. Applying this post-correction to the raw depth outputs (i.e., prior to MLS smoothing) permits to reduce the MAE of the *3 pts* and *Laplace* methods to 0.322 and 0.211 respectively which are both still clearly larger than the MAE of our *Triangle* method. This post-correction has no effect to our *Triangle* method as the magnitude of this bias effect is already very low, and lower than the amount of noise.




Figure 1.14: For each of the three methods, the top histograms show the distributions of the computed depth values for the 80 depth maps \mathcal{D} , and the bottom scatter plots show the ground truth decimal depth values as a function of the decimal of the computed depth values. Those plots also include least-squares cubic approximations of this function (blue curve), and the ideal identity relation in red.

Depth from focus at micrometer scale

To evaluate our approach in a real-world context, we strived to acquire the relief of prehistoric engraved stones exhibiting groove widths ranging from 50 to 100 μm . The precision of the depth-from-focus pipeline is greatly influenced by the design of the acquisition setup [33]. To achieve a $\times 2$ magnification while minimizing the depth of field, we stacked together a 100 mm lens with an inverted 50 mm lens (see Figure 1.16) playing the role of a close-up lens (Figure 1.17). To minimize the depth of field, both lenses are fully opened (respectively 1:2.8 and 1:1.4).

With this setup we found that a step size δ_z of 20 μm to be a good tradeoff. Automatic staking is achieved through a motorized rail, while the resulting images are aligned with the *Enfuse* software prior to depth estimation with our pipeline. Figure 1.15 shows one result at the crossing of three strokes. Despite the depth of the strokes (about 40 μm) being of same order as the sampling depth δ_z , we can clearly see the stroke order.

As we can see in Fig. 1.15, the engravings are visible and visually separable by their depths. In addition, we notice the different grooves in the same engravings, made by the fact that the tool used has been passed several times. Our method therefore makes it possible to produce a sufficiently robust 3D profile on macroscopic images to observe the details of engravings that are impossible to digitize with a laser scan or by photogrammetry.



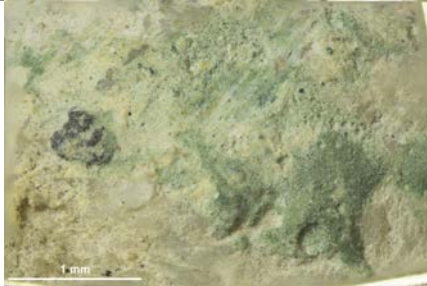
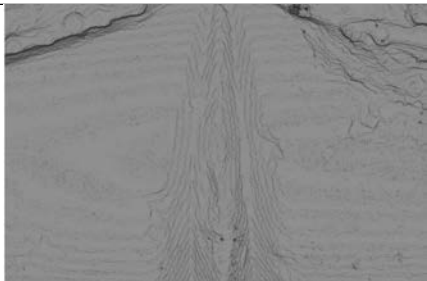
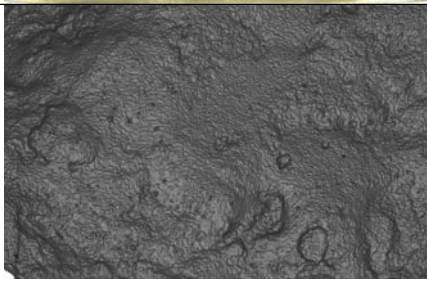
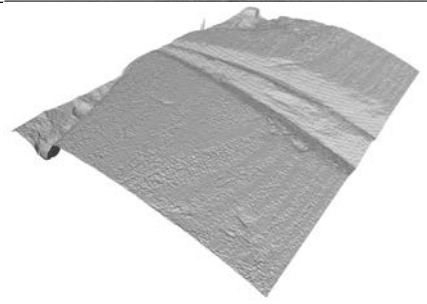
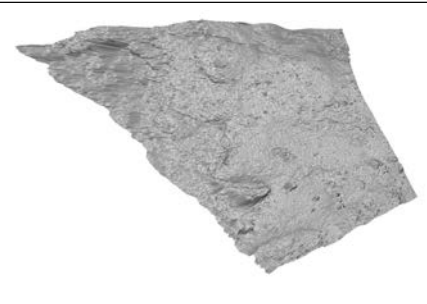

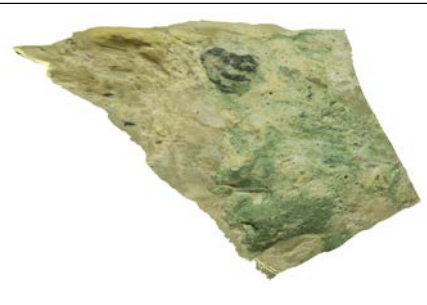


Figure 1.15: Results get on a real-world stack from an engraved stone. **Top:** 3D profile of the acquired engraving area. **Bottom left:** blurry image obtained on the image stack. **Bottom right:** reconstituted clear image of the engraving.

To go further, acquisition have been made on site in Delphi, Greece. Blocks of stone were captured outdoors (see Table 1.3). For the first line, acquisition has been proceed with a Nikon Z 100 mm lens ($1\times$ magnification) and step size δ_z of $200\ \mu\text{m}$. The second one has been made with a Mitutoyo $10\times$ and a tube lens of 200 mm ($10\times$ magnification) with a step size of $10\ \mu\text{m}$. We note a progressive offset of the shots for the second case because the alignment of the optical axis with the translation axis is not perfect. However, this offset can be corrected rather well by aligning the images on *Hugin* [4]. The results demonstrate that our method can also be used in-situ to obtain a relatively robust level of detail on macroscopic images.

Other practical results can be found in the next section, where our method has been used to accurately measure the absolute thickness of thin glaze of ancient ceramics in a non destructive manner.

Table 1.3: Example of of on site results get on a in-situ real-world stacks from stones in Delphi, Greece.

	Nikon Z 100 mm	Mitutoyo 10×
All-in-focus picture		
3D model in ortho view		
Other 3D view		
3D model with texture		

1.3 Application to monumental heritage: measuring ceramics glaze thickness

With a wealth of archaeometric approaches at our fingertips, ceramics are now widely studied with a variety of methods to answer an array of historical and prehistorical questions, such as questions of provenance, dating, method of manufacture, and trade, to name a few. In this regard, physiochemical techniques such as Scanning Electron Microscopy (SEM) [204, 71], X-ray diffraction (XRD), and Particle Induced X-ray Emission (PIXE) [125] are commonly used for determining numerous characteristics of glazed ceramics. One ceramic attribute that can be particularly informative is the determination of the thickness of a glaze. Traditionally, this measurement has been obtained by SEM examination of polished thick sections. As well as being time-consuming, this method requires the partial destruction of an object, to be able to take it to the laboratory and produce the required analytical sample. Furthermore, in the case of intact objects or architectural ceramics in place, it is inconceivable if not impossible to extract a sample. In this context, the prospect of a mobile, portable, rapid and above all non-invasive method has become an especially appealing alternative. Overall, however, the measurement of the thickness for ceramic glazes have been sparsely exploited. As a notable exception, in some works [203, 73], it was possible to observe an evolution of the thickness in relation to the chronology. In other ways, the thickness of the glaze can be used as an additional factor for characterization [147], or to discriminate between glazed ceramic productions processes [121]. Namely, it can reflect not only the chemical composition of the glaze, but also the method of application, and firing protocol. The variation in the thickness of the glaze also plays an important role in the perception of the color of the glazed ceramic, as increased thickness results in greater absorption of light by the glaze. With so much information produced through the measurement of glaze thickness, and the aforementioned limitations of various methods for its measurement, the growing influence of 3D analytical techniques in archaeology provides a particularly interesting opportunity for researchers.

Recently, the techniques of computational photography and 3D acquisition have played an increasingly important role during excavations and historical research. For example, photogrammetry [91] and 3D lasers [34] can increase the possibilities of post-excavation observations. Quick to install and portable, only requiring a camera and associated equipment, these methods have the advantage of providing valuable information relatively easily. Consequently, the development of techniques based in computational photography is particularly compelling approach for the study of glazed ceramics. Further suggesting the potential of this method, the advanced state of the art in this field also makes it possible to obtain depth information (i.e., thickness) from photographs. Specifically, we demonstrated previously that Depth from Focus strives to recover 2.5D depth information accurately. Consequently, this simple and global measurement of the average thickness of the transparent glaze can then be easily associated with physicochemical techniques, and may therefore be an excellent option for the on-site study of ceramics.

This aim of this section is twofold:

1. We develop a non-destructive and non-invasive method for the on-site analysis of glazed ceramic tiles. This is achieved by measuring in a simple, fast, and inexpensive way the thickness of the glaze while generating a clear surface image at the macroscopic scale (based on DfF).
2. We compare the results of measurements from different types of ceramics to assess the reliability of the method and to determine any limitations and drawbacks.

For this second goal, fragments of various origins have been chosen.

1.3.1 Non-invasive on-site method for thickness measurement

As mentioned above, we apply Depth from Focus algorithm. This advantage of it is that we only uses a single camera with a lens system, a ring light, and a tripod to obtain 2.5D pro-files of scenes in the form of a depth map in a monoscopic way. As it was shown by our previous section and confirmed by [33], this system makes it possible to obtain profiles with better depth accuracy than even photogrammetry, at less than ten micrometers.

This technique has certain limitations, however. Focus measures are based upon variations in the grey level of the image. It is therefore necessary to have details present on the image. It is indeed necessary to obtain sharp areas on the image so that texture may allow a clearly recognizable depth for each area. In addition, transparent materials cannot be accurately modeled in 3D. It is simply not possible to focus on a perfectly transparent surface without adding removable paint on the whole surface of the glaze. However, depth measurement is done on the focal plane. Consequently, the depth detection can take place on top of a transparent surface if it has defects that generate a slight local opacity, or the bottom if it is fully transparent in the given direction.

Glazed ceramics presents is one germane example of the impossibility of precision for 3D acquisition because they have a transparent glazed thickness on the surface. In the case of old ceramics in particular, these glazes are often scratched, chipped, etc. These effects generate a surface that is almost transparent but with roughness. The depth map obtained by DfF is there-fore heavily altered in these cases. The clear zones will sometimes appear on the surface of the glaze, while other time below, directly on the bottom. In our case, we have used the defaults of this method to get results. Obtaining depths over and under the glaze in such a way can nevertheless allow for non-invasive glaze thickness measurements. Specifically, two levels of depth were obtained by taking the measurement above opaque areas (stripes, lines made with a felt tip marker line, etc.), which can then be extracted and processed to obtain a thickness of the glaze.

Acquisition set-up

The analytical protocol requires a level of precision to just a few micrometers to enable precise measurements of thickness of the glaze. To accomplish this, it is necessary to reduce the depth of field as much as possible to optimize the precision. The full apparatus set-up can be seen in Figure 1.16. Specifically, we have used a Canon EOS 5D II camera (Figure 1.16a). The camera has a body sensor of 20 megapixels (pixel size, 6x6 microns) and was mounted on a motorized rail on a macroscopic scale which can be controlled either by computer or by hand (Figure 1.16b). This allows for an incremental movement of 10 cm with steps that can go as small as 2 μm so that images can easily be stacked at the desired scales. The optical capacity has a great influence on the precision of the acquisition of data [33]. As we require a set of lenses allowing $\times 2$ magnification and reduced depth of field, we used a 100 mm lens, and an inverted 50 mm as a close-up lens (Figure 1.16c and Figure 1.17). Both were fully opened (respectively 1:2.8 and 1:1.4). A light ring was attached to the end of the lens as a light source, allowing diffuse and grazing illumination at a working distance of < 1 cm, while not occluding the lens (Figure 1.16d). The focal stack was acquired by connecting the motorized rail to a *Helicon Remote* for automating the method. Finally, we then precisely aligned the photos by using *Enfuse* software for a precise alignment while limiting the degradation of the image.



Figure 1.16: Used acquisition device. (a) Canon EOS 5D II camera; (b) motorized rail; (c) optical system using a 100mm lens and an inverted 50mm lens; (d) light ring with diffuse filter.

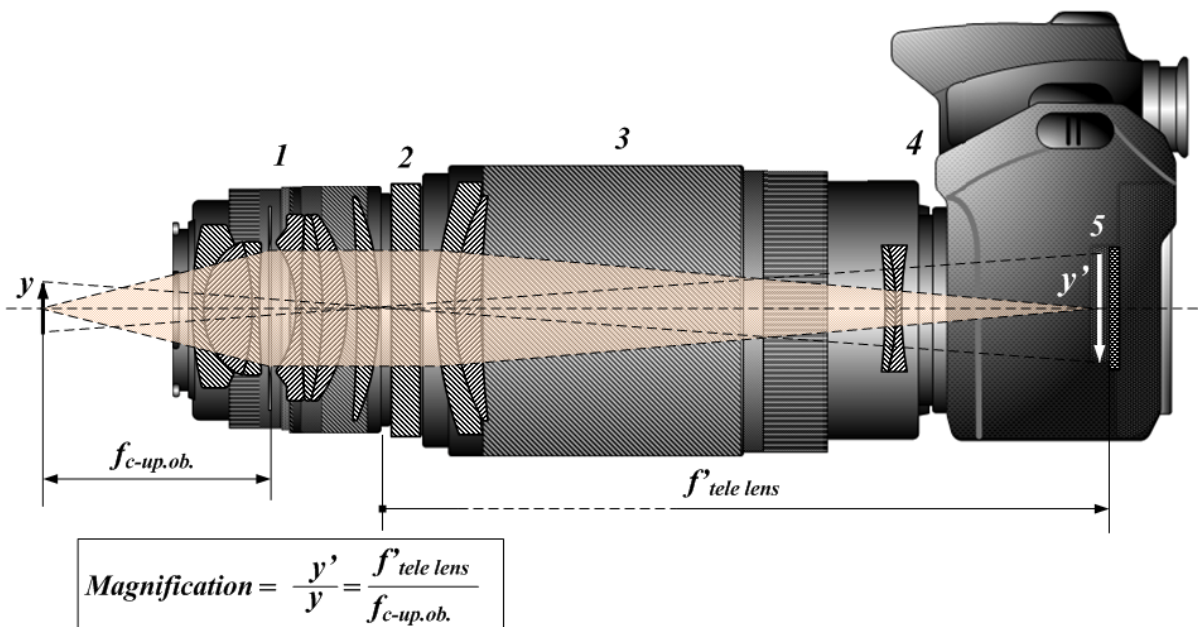


Figure 1.17: Illustration of an inverted lens as a close-up lens. Image taken from Wikipedia.

Determining the thickness

After aligning the stacked images, the Region of Interest (Figure 1.18a) was clipped out, (i.e., the area containing texture and stripes or felt on the surface), followed by the application of the DfF algorithm (the only treatment made on *MATLAB*). We first made an initial very coarse depth map for measuring the tilt of the ceramic with respect to our device (fitting with a polynomial surface of degree 1). If the sample was not flat, we could then further determine its shape (fitting with a polynomial surface of degree 2). We then produced a second depth map. High levels of precision are necessary for this second map so that information on both side of the glaze may be detected. Using the surface shape established with the coarse depth map (untilting for a flat ceramic, correction of the volume otherwise), we then transformed our depth map and obtain the planes above and under the glaze at constant depth (Figure 1.18b). Though the depth map obtained is often quite noisy, the two levels of detection are in fact yet visible.

By displaying the histogram of the values of depths of the map, two distinct peaks are evident (Figure 1.18c). These peaks correspond to the two depth levels above and below the glaze. The peaks are wide for two reasons: the variability of the measurement and the fact that the ancient ceramics are hand-made and therefore have variations in depth (which can be seen in Figure 1.18f). By subtracting the two peak values, one can obtain a thickness value d' , corresponding to the advance of the camera. However, this value does not by itself reflect the thickness of the glaze. Rather, as shown in Figure 1.19, it is necessary to consider the refractive index of the glaze to obtain the thickness $d = d' \cdot n'$. Given this, the method therefore requires having an a-priori knowledge of the refractive index. An example will be given in section 1.3.2.

Using these values, we can obtain an absolute value of the thickness of the glaze in a non-invasive way and determine the overall uncertainty from the uncertainties $\Delta n'$ and $\Delta d'$ (respectively the uncertainty on the refractive index n' , and the uncertainty on the depth values d' at the top of the peaks) according to the formula:

$$\Delta d = d \sqrt{\left(\frac{\Delta n'}{n'}\right)^2 + \left(\frac{\Delta d'}{d'}\right)^2} \quad (1.8)$$

In the following, we present the result as:

$$d \pm \Delta d \ (\mu m)$$

The result is therefore the thickness with approximation of the transparent glaze. However, this thickness is not a localized measurement but a global measurement for the region of interest studied. It is an average measurement on a sample which can reach more than 1 cm^2 .

In the end, Python is used to code all the presented processing. For the entire process, the time required for the thickness measurement is very short. Once the equipment is installed, it takes no more than 5 min to create the image stack. The processing of the Depth from Focus on *MATLAB* and subsequent *Python* coding [49] can then be done within 15 min after setting the parameters using our non-optimized implementation.

1.3.2 Samples

These measurements were carried out on several ceramics reflecting different cultural origins, uses, and eras to highlight the precision and the limits of the method. Commonalities between all samples studied include a Mediterranean basin origin and the presence of a transparent glaze allowing for the thickness measurement by DfF. All the glazes of the samples studied are transparent lead glaze with PbO between 23 and 57 Wt% [27, 14, 15, 26, 21, 23], the predominant method of the 2nd millennium AD around the Mediterranean Sea. While conventional glass tends to have a refractive index of approximately 1.5, the inclusion of lead in glass tends to significantly increase the refractive index of glass [158]. As a result, we estimated the refractive




Figure 1.18: (a) macroscopic photo of the surface reconstructed after DfF on the Iznik tile Bdx 6502; (b) obtained depth map and (c) its depth histogram for the detection of the two levels: depth under-lined in red on the depth map (d) at $-601 \mu\text{m}$ and (e) at $-467 \mu\text{m}$. (f) is a section BSE image made by SEM for comparing the method.




Figure 1.19: Optical diagram of the camera on the glaze showing that the advancement of a distance d does not correspond to the measurement of a thickness d of the glaze, and the need to take into account the refractive index n' . Advancing the camera between (a) and (b) by the distance d does not provide a clear image of the bottom of the glaze. While advancing the camera the distance d' between (a) and (c) allows focusing at the two desired levels.

Table 1.4: Description of different glazed ceramics studied.

Style and object type	Inventory number	Glaze Color	Decoration ^b	Provenance	Date (century AD)	Reference
Iznik Tile	Bdx 6493 Bdx 6501 Bdx 6505	Blue-green	B, Bl, G, RB	Tunis, Tunisia	17th	[143, 15]
House stove tile	Bdx 16621	Green		Berg Armo, France	16th	[21]
Maiolica dish	Bdx 2591	Colorless ^a	B, G, RB, Y, Be	Ravello, Italy	16th	[168, 14]
Zellij (ceramic mosaic)	Bdx 6522 Bdx 6524	Green Blue		Meknes, Morocco	14th	[26]
“Green and brown decoration” dish	Bdx 5502	Honey	Br, G	Raqqada, Tunisia	9th - 10th	[52, 27]
Earthenware	Bdx 21066 Bdx 22625	Colorless	Bl R	Bordeaux, France	19th	[22]

^a This glaze covers the decorations and a white glaze opacified with tin oxide.

^b B, blue; Be, beige; Bl, black; Br, brown; G, green; R, red; RB, red-brown; Y, yellow.

index at 1.75 ± 0.25 [158] for all the samples was due to their common high lead component. The different glazed ceramics are presented in Table 1.4.

Each sample has specificities making it possible to observe the use of DfF in various cases. First, Iznik ceramic from Tunisia presents a simple case: the glaze is flat and transparent and has sharp cracks that easily allow two levels of depth to be achieved. Like the Iznik ceramics, the Maiolica sample (Bdx 2591) is flat, but the white color of the ceramic on which the red-brown and yellow decorations have been applied (before being completely covered by a colorless

transparent glaze) exhibits little detail and texture, making depth harder to detect by DfF. The house stove tiles in our sample are all engobed and covered with a green glaze of varying tones. Like the other studied samples, transparent glazes found on these tiles are strongly lead-bearing (47 to 66% in PbO) [21] and have many cracks on the surface. Nonetheless, we focused on a flat area for our measurement. The Zellijs samples, on the other hand, have no surface defects and are a good example to assess the methodological limitations in cases where fragments are absent surface defects. For our purposes, the Raqqada fragment features yet another advantage when compared to the others, in that it is not perfectly flat. Specifically, these are pieces of curved ceramic (probably dishes). This fact makes the measurement of the thickness of the glaze more complex and sheds light on possible additional limits of the DfF. To round out our sample, the study of white earthenware by DfF is by far the most complicated case presented here. Indeed, earthenware has three limiting characteristics for our method: a very specific body due to its whiteness which makes it very difficult to detect texture, a relatively thin glaze thickness and a non-planar shape as in the case of the sample Bdx 5502.

1.3.3 Results and discussion

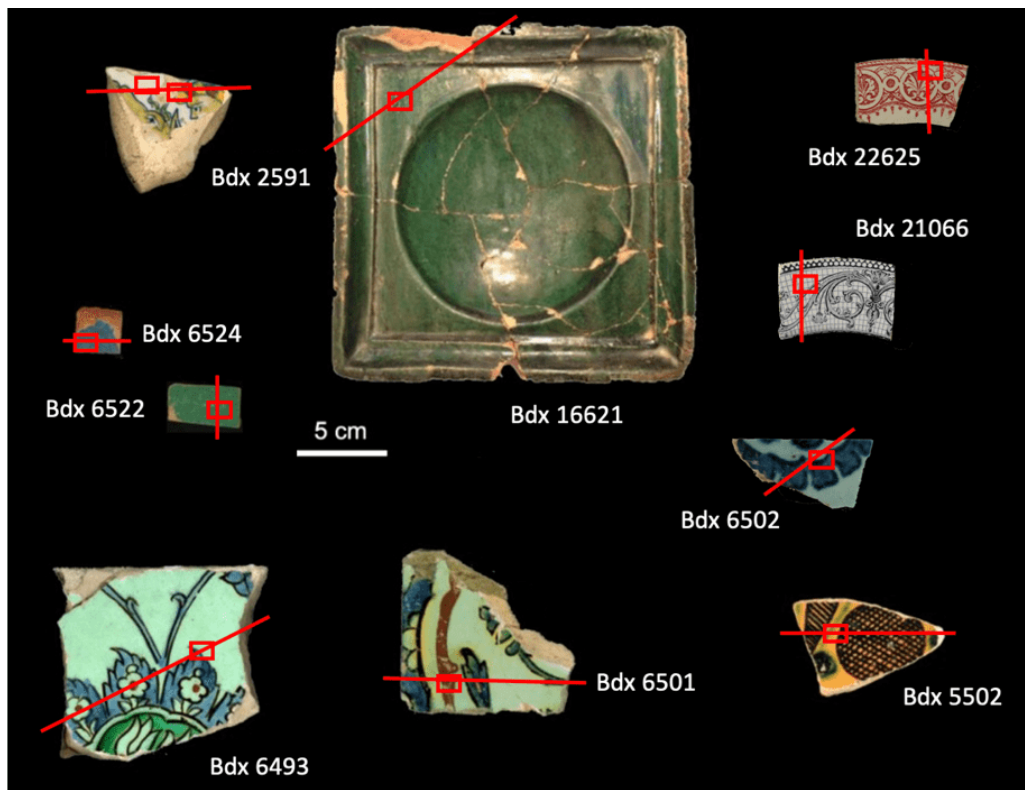
To assess the reliability of our method, we tested the DfF on fragments of glazed ceramics in the laboratory, we then observed these sections on the Scanning Electron Microscopy (SEM, JEOL JSM-6460) and compared the values obtained. On each section, we realized a minimum of 15 depth measurements to obtain an average glaze thickness.

Measurements from Depth from Focus method

For the Depth from Focus (DfF) methodology, we first measured the thickness of the different tiles using our set-up, we then made sections for the SEM observation for the aforementioned reliability assessment. Iznik's tile samples present the ideal case for this test. In addition to presenting very clear scratches on the surface and details in depth, the thickness of the glaze is important. As detailed in section 1.3.1, we measure $233 \pm 55 \mu\text{m}$ for Bdx 6502, $199 \pm 49 \mu\text{m}$ for Bdx 6493 and $156 \pm 27 \mu\text{m}$ for Bdx 6501 (the results are visible in Figure 1.20)). Regarding these samples, the depth histogram is clear in this case (see Figure 1.18c): the two peaks are readily identifiable.

The study becomes more complicated for the Berg Armo house stove, however. This case is illustrated in Figure 1.21. The scratches are in fact clear and easy to identify on the depth map (Figure 1.21b). But the thickness is much thinner than in the case of Iznik ceramics. Consequently, the two peaks on the histogram mix and to form a single, larger peak (Figure 1.21c). It is necessary to develop a visualization highlighting in red on the depth map the pixels for a chosen depth (Figure 1.21d and 1.21e). By combining the depth histogram and the visualization, we may then obtain an exploitable depth of $58 \pm 26 \mu\text{m}$ for the sample Bdx 16621. We estimate that for less than approximately $150 \mu\text{m}$, it becomes harder to measure the glaze thickness from the histogram. The Raqqada ceramics present a similar situation. With these, a complicating factor for the performing measurements is that the samples are not flat. Using the volume correction (section 1.3.1) we measured a glaze of $102 \pm 38 \mu\text{m}$ for this sample.

To determine the thickness of the yellow pigment of Ravello maiolica Bdx 2591, we made measurements at two points: the glaze-yellow and glaze-white thicknesses. As with the Berg Armo ceramic, the glaze thickness is quite thin, and is observable with the histogram coupled with the visualization. Specifically, the measurement for the glaze-yellow thickness is just $29 \pm 10 \mu\text{m}$. However, the absence of scratches or degradation on the surface above the white part creates difficulties for measurement. To combat this, we drew a red dot on the surface with an erasable marker. The contours of this red point on the surface of the glaze were then measured as being on the surface of the glaze. Using the above visualization, we achieved a thickness value of $86 \pm 21 \mu\text{m}$.



Sample	Depth (µm)	
	SEM	DfF
Iznik tiles		
Bdx 6493	215±38	199±49
Bdx 6501 (glaze to red)	167±11	156±27
Bdx 6502	265±64	233±55
House stove tiles		
Bdx 16621	72±16	58±26
Maiolica dish		
Bdx 2591 (yellow part)	36±7	29±10
Bdx 2591 (white part)	99±6	86±21
Zellij		
Bdx 6522	352±24	267±48
Bdx 6524	255±15	306±49
Raqqada ceramics		
Bdx 5502	113±7	102±38
Earthenware		
Bdx 21066	51±3	57±18
Bdx 22625	75±2	69±15




Figure 1.20: Image of the fragments measured by DfF as part of the method validation (*top*). Analyzed areas are represented on the samples by red lines (sections made for the SEM) and red rectangles (measured areas by DfF). Numerical results (*bottom left*) and visual comparison (*bottom right*) of measured depth by SEM and by DfF on samples illustrated on image. In numerical results are the distance between the exterior surface of the glaze and the surface of the body, excepted when we precise “glaze to red”: in this context, we measure the distance between the exterior surface of the glaze and the upper surface of the red pigment.

Zellijs from Meknes features the same disadvantage as the maiolica because the surface does not include any defects. For this sample, we again employed the red dot technique. In this case, measurements revealed a very thick glaze, and as such, the histogram is sufficient without further visualization, despite the fact that the Zellij glaze includes many bubbles and unmelted crystals. Due to these imperfections, the main peak of the bottom is very wide, and a lot of noise is visible on the histogram. We nonetheless measured a thickness of $267 \pm 48 \mu\text{m}$ for Bdx 6522 and $306 \pm 49 \mu\text{m}$ for Bdx 6524.

Finally, the Bordeaux earthenware is the most complicated case yet, given that it features the complicating factors of all the previous fragments: very fine glaze, no alterations on the surface of the glaze, no details on the white background (only the patterns can be measured) and a non-planar volume. Using our methods, however, we were able to overcome these difficulties and yet measure a thickness value by using the red mark on the surface and by coupling the histogram and the visualization after correcting for volume. The glaze was thus estimated to be $57 \pm 18 \mu\text{m}$ for Bdx 21066 and at $69 \pm 15 \mu\text{m}$ for Bdx 22625.

All-in-focus image of the samples, corresponding depth map and histogram are visible in Annex I.

Comparison with SEM

The values measured by SEM are visible in Figure 1.20. Overall, the results are very consistent with those obtained through the DfF method. Ten of the eleven values measured by SEM are within the thickness uncertainty interval obtained by DfF. Only the measurement results for one Zellij does not match. For this sample, Bdx 6522, we suspect that the large presence of unmelted inclusions and bubbles in the glaze may be reducing the value of the measurement. On the other hand, it is also possible that the assumptions about the refractive index are incorrect. However, the orders of magnitude of the Zellijs nevertheless give a thickness scale of 200-400 μm .

Notably, we obtained similar estimations of thickness, as indicated by the overlapping intervals (Figure 1.20, by reducing the thickness to 30 μm despite the difficulty in reading the histogram and the need for visualization. In the end, the resulting interval of thickness is quite broad, which can be explained by the large uncertainty in the refractive index, but also by wide thickness variation within the sample we observed through SEM.

1.3.4 Conclusion on the thickness measurement method

Through comparison of our measurements to those acquired through SEM, we have demonstrated that an acquisition of a focal stack using a simple camera is an effective and accurate method for measuring the thickness of a transparent ceramic glaze. Furthermore, the advantages of this method are numerous. First, there are the non-invasive and ultra-portable aspects of the DfF method and the possibility of measuring directly on site without any contact with the ceramic. Moreover, the method is also comparatively fast to complete and does not require significant equipment. Our analyses using different types of ceramic also demonstrate that it is possible to adapt the method according to sample-specific factors, including the shape of the sample, the tint of the transparent glaze, the state of preservation of the glaze (i.e., presence or absence of deterioration), and the sample's thickness. This is all critical, as the variability in the thickness of the glaze effectively influences the perception of the color of the transparent glaze. Finally, the DfF has the advantage of giving an average measurement over a larger area than that observed by SEM (which is constrained in terms of analysis area).

The Depth from Focus glaze measurement method nonetheless has certain limitations. It requires strong a priori assumptions, such as those regarding the type of glaze. As such, it is most useful as a compliment to the overall study of the sample, rather than a wholesale replacement of the SEM. The DfF, for its part, offers a different visualization from the microscopy, with a frontal view (and not in section). This does not allow for as clear of view of the state of the glaze




Figure 1.21: (a) macroscopic photo of the surface reconstructed after *DfF* on the house stove tile from Berg Armo Bdx 16621; (b) depth map and (c) its depth histogram for the detection. As we can see, the two peaks blend. It is therefore necessary to use the visualization by depth to obtain the depth (d) under the glaze at $-925\ \mu\text{m}$ and (e) over the glaze at $-892\ \mu\text{m}$. (f) is one of the SEM BSE im-ages made to compare.

itself. In addition, unlike SEM, the DfF cannot provide information regarding the components of ceramic materials. Finally, though methods for correction are available as discussed above, measurements regarding the thickness and state of the glaze can yet be difficult to obtain, regarding to the thickness and the state of the glaze. In the ideal case (as on the Iznik tiles), the measurement can be quite simple, and apparent through the histogram. But below 150 μm , it requires a more intensive analysis to locate the depth of the two planes. In certain cases, such as when the glaze is not completely transparent (i.e., in the case of the Zellijs), it can become downright impossible to obtain a perfect result. Yet, despite any difficulties or limitations, the use of the DfF method for measuring glaze thickness remains encouraging.

1.3.5 Results on the tiles from Saint-Maurice residence




Figure 1.22: *Illustration of our process applied in Saint-Maurice residence in Cairo.*

After confirming the accuracy of our approach for measuring glaze thicknesses, we proceeded to apply it to the Iznik tiles of Saint-Maurice residence. Due to time constraints and limited access to certain ceramics at elevated heights, we conducted measurements on only 13 instances across 9 tiles. Specifically, we recorded 7 measurements on P2 group tiles, and 6 others that were distributed across tiles 1 and 4 of group C2. This second set of measurements aimed to explore the significance of variations within a single tile.

The measurement is made with the same set-up as before (as depicted in Figure 1.16), with each measurement being taken on a surface of approximately 2 cm in length.

Of the total 13 measurements taken, we obtain results for 12 of them. The measurement for tile 89 P2, however, produces a depth map that is inoperable because of a glaze without any scratch, resulting in a fully transparent surface that offer no discernible differences between the glaze exterior and the body surface. But the remaining measurements show clear histograms featuring two distinct and sharp peaks, as depicted in Figure 1.23.

Results and discussions. The results are visible in Fig 1.23. We can quickly highlight two points: (1) the thickness of the glazes of the Cairo tiles have the same order of magnitude as the Iznik ceramics from Tunis measured with the SEM. (2) Tiles show fairly significant variations, between 87 P2 and 90 P2 for example. But these differences are hard to elucidate because we notice that a single tile may vary significantly as it is underlined in the case of C2 group. In addition, we observed in the previous section that ceramics can present significant variations in the same area locally. Both the body and surface are not smooth and fluctuates. That cause larger uncertainties of measurements. Finally, these variations seems not necessarily

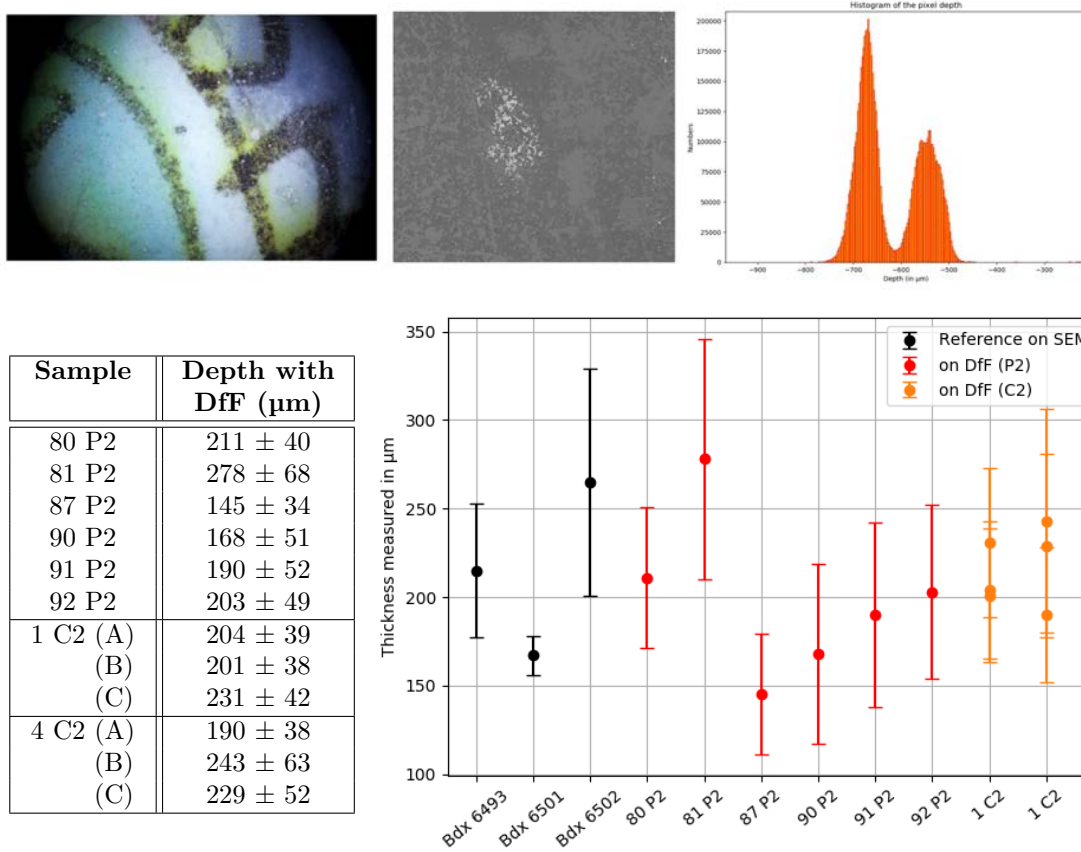


Figure 1.23: All-in-focus image of the 80 P2 tile, corresponding depth map and histogram (top left to right). These informations are visible for the other measurements in Annex I. Numerical results (bottom left) and visual comparison with reference results on Iznik ceramics from Tunisia (bottom right) of measured depth by SEM and by DfF on samples illustrated on image. In numerical results are the distance between the exterior surface of the glaze and the surface of the body.

meaningful because they can be caused by disparities during the confection (handmade), so it can be explained by fluctuations in the amount of glaze applied before firing.

According to these remarks, the measurement of the thickness of the glaze by DfF therefore remains a good way of characterizing ceramics as being Iznik ceramics within the framework of a global study.

Comparing different groups has proven to be of great interest as it allows us to ascertain that **the order of magnitude remains consistent from one group to another and between those from Cairo and Tunis. This similarity indicates a shared methodology of conception.** As previously observed, other methods can result in significant variations in thickness, as exemplified by the varying thicknesses of Zellijs glazes (thicker) and Raqqada ones (thinner). However, these methods were unable to distinguish between individual tiles in our case. That's why, we must examine other characteristics of ceramics, such as pigments and their composition, to differentiate between them.

This chapter allowed to study in more detail how to acquire the shape of a 3D object. Precisely, the development and improvement of a Depth from Focus method was therefore relevant and made it possible to acquire 3D profiles with accuracies of the order of ten μm . Thanks to

this, we were able to study the transparent glazes thickness of ceramics from various periods and origins. After validating our method, we applied it to Iznik ceramics from Saint Maurice residence and compare our results with values obtained on other late Izniks from Tunis. If the measurements did not make it possible to compare the ceramics of the residence between them, they showed that the thicknesses were comparable to the Tunisian tiles.

However, we have seen that the study of the glaze of the tiles alone is not enough for a complete non-invasive characterization. This is why it is interesting to cross-reference the results obtained with other data, such as pigments information for example. In this sense, a cross-study with the reflectance spectra of the pigments used seems to be a track for the future.

Chapter 2

Hyperspectral Imaging for Reflectance Spectra Similarity Studies

We have elaborated in the previous chapter a portable and non-invasive method to acquire the shape of objects. Its application on the Iznik ceramics is unfortunately not conducive to obtain the expected results, that is to classify and compare the ceramic tiles. The variations due to hand-making remain too great and the measurement of the glaze thickness only leads to the conclusion that the method used to create these ceramics tends to be the same as on the Iznik ceramics from Tunis.

This is why observing the pigments instead of the glaze can give more information. For this, the study of the reflectance spectra of the pigments seems indicated. Indeed, the pigments have characteristic reflectance spectra which allow their identification by identifiable bands (which can be bumps, slopes, peaks, etc.). In addition, the study and processing of spectra acquisition is part of our logic of studying the digital acquisition of objects as a whole: after being interested in the 3D aspect, the present chapter centers on the acquisition and process of color as a spectral information.

2.1 The interest of hyperspectral imaging for pigments characterization

As explained by Cucci *et al.* [51], Hyper-Spectral Imaging (HSI) has emerged over the last 30 years as a non-invasive analysis technique of artworks surface for Cultural Heritage. It is effectively used for the identification and the mapping of materials used such as pigments or dyes, and their state of conservation. This is why we focus in this chapter on the HSI study of the Iznik ceramics pigments of Saint-Maurice residence.

2.1.1 Differences between Hyperspectral and RGB Imaging systems

RGB imaging systems

First of all, it is important to define how a classic camera works in order to understand the challenges to overcome to develop a hyperspectral imaging system. Very schematically, a classic digital camera is composed of a photo-sensitive sensor and an optical system (lens, diaphragm, shutter). The optical system of a camera is mainly composed of a lens assembly. Its role is to select and guide the light rays on the surface of the sensor. Sensors are separated into two groups: CCD (Charge-Coupled Device) and CMOS (Complementary Metal Oxide Semiconductor). Most cameras have CMOS sensors these days because they are cheaper, faster, and less

energy-consuming. A sensor is made up of millions of photo-sites (i.e., pixels) that capture photons and transfer them as an intensity value [101]. However a sensor only delivers an overall intensity: it can only output a grayscale image. This is why manufacturers place on the sensor a Bayer matrix, which acts as a spatial filter (see Figure 2.1), composed for each group of 2 pixels with 2 green pixels, one red and one blue [38]. The result is that the photo-sites corresponding to each of the components are not aligned, and the final image must therefore be reconstructed by an interpolation process. So we observe at the level of the digital camera that the RGB color image requires the Bayer filter as an additional object because the sensors only collect one piece of information per pixel.

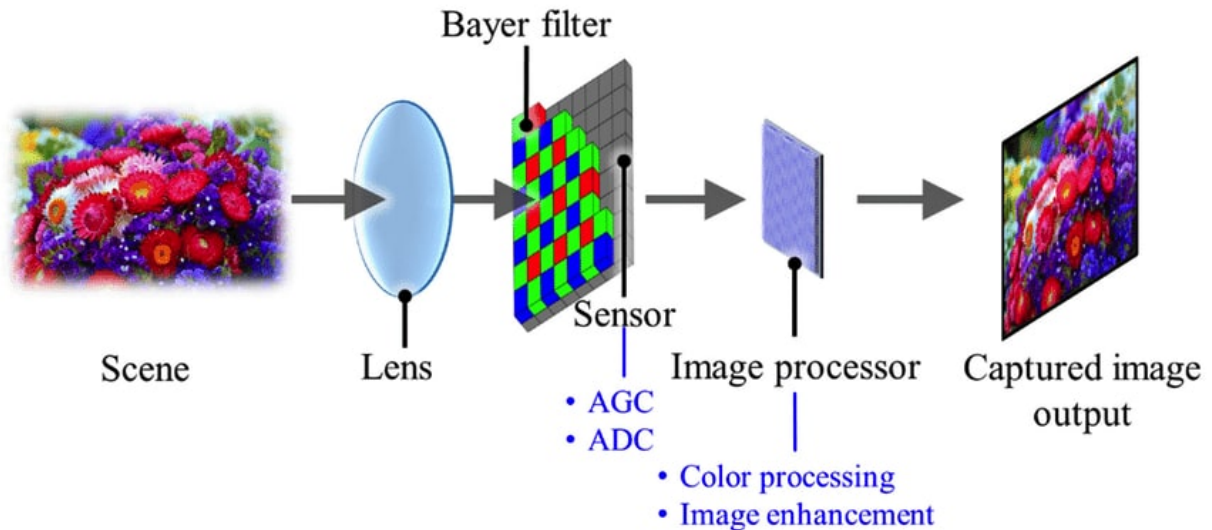


Figure 2.1: Illustration of a digital camera. Image taken from [38]. ADC means Analog-to-digital Converter. AGC means Automatic gain Control.

Hyperspectral imaging systems

In order to obtain hyperspectral images instead of RGB, HSI has to put together spectroscopy and imaging: data acquired involves capturing a large number of images with spectral continuity, acquired in narrow and contiguous bands across an extended range of wavelengths. This range can be limited to the Visible and Near Infrared (VNIR – 400-1000 nm) or extended to the Short-Wave Near-Infrared (SWIR – 1000-2500 nm). Note that the term hyperspectral is usually used but the difference between hyperspectral and spectral is quite ambiguous.

The desired result of the acquisition is a data-cube (also called hyperspectral image) composed by 2-D spatial information and a third, spectral dimension, thereby providing a complete spectrum for every pixel in the imaged scene [93]. This results in the association of a high-resolution reflectance spectrum with each pixel of the imaged scene instead of an RGB triplet for classic cameras.

Like a classic camera, a HSI system combines an optical system and a sensor. Contrary to sensors, which remain substantially the same, the optical systems of a HSI must differ a lot. To obtain a spectral system, it is effectively not feasible to employ simply a classic optical system followed by a filter grid on the sensor in a similar manner to the conventional Bayer matrix. The reason behind this is the requirement of a large number of values to obtain a spectrum (generally several hundreds), which cannot be accommodated by the filter grid approach. We then need to create an entire new system. In this way, reflectance hyperspectral imaging systems were first developed during the 80s as avionic systems for the remote observation of the earth [80].

Although aeronautic application remains important [160, 153], it gradually expanded to other areas such as industry [198] and food [83] control, pharmaceuticals [96] or agriculture [201].

Nowadays, many technologies were developed like snapshot [109], tunable filters [76], or Fourier transform infrared (FTIR, [194]), but pushbroom technology [59] persists in being the most commonly used (see Figure 2.2).




Figure 2.2: Illustration of a pushbroom imaging system. Image taken from [59]. For each pixel of a pixels line, the light enters the pushbroom camera through an entrance slit. Then, a dispersive element disperses the beam into its constituent wavelength components. Following this dispersion, a focusing optics system is employed to project the quasi-monochromatic beams onto a matrix detector, such as a CCD or CMOS detector. After it, the system moves in order to scan the next pixel column. At the end, the image is reconstructed.

Tunable filters systems consist in scanning the wavelengths. One frame is acquired for each required wavelength in the spectral domain. A system of tunable filters is generally placed upstream of the optical system (or before the sensor) to select only the desired wavelength. Between each frame, we therefore change the filter then redo an acquisition, until the data-cube is filled. The main advantage is that we obtain a spatial resolution identical to that of an RGB system. But we cannot obtain a large number of wavelength values because it is tough to manipulate hundreds of filters. To partially counteract this problem, another type of tunable filter system exists: liquid crystal systems whose resolution and loading time can be a good intermediary. The acquisition time can still be very long, and it is necessary to make sure to keep the same scene while ensuring that the lighting remains uniform.

Snapshot systems (like those developed by *Imec*) therefore reduce the time required for acquisition. They separate the sensor into group of pixels and cover each of them with different dichroic filters that represents a single spatial pixel (like would do a Bayer matrix). Where the acquisition can be done in one shot, the spatial and spectral resolution are strongly impacted because it is necessary to sacrifice highly one of them (the higher the spectral resolution is, the larger the groups are and the more the resolution space decreases). Moreover, as for the tunable filters camera, the use of filters limits the quantity of photons received on each pixel. So there

is a big loss of intensity.

FTIR systems work by associating each pixel with a complete spectrometer. Thus, each pixel measures an interferogram from a different spatial position of the scene in a single acquisition. This interferogram works by summing cosines generated by monochromatic light interference. Then, the data of each pixel is processed by Fourier transform to obtain the spectrum. So, in one acquisition we get a complete and high-resolution data-cube. FTIR cameras are however very expensive and imposing. They are therefore limited to lab measurements.

Finally, **pushbroom** technology is the most versatile one. An optical system images a horizontal line which is then dispersed vertically by a spectrometer on the sensor (more details in Figure 2.2). We therefore measure one spatial dimension (horizontal one) and the spectral one. Adding a mechanical scanning system generates a line-to-line scanning method and recover the vertical spatial dimension: the data-cube is generated. Although the acquisition time of a data-cube is very long (line-to-line scanning can be exceedingly long for a high-resolution data-cube), pushbroom cameras hold many advantages. At first, they are not limited by their resolution (spectral as well as spatial) like snapshots and tunable filters systems. Then, using a spectrometers instead of transmissive filters prevents the loss of intensity on the sensor. Finally, they are much cheaper than FTIR cameras and much more maneuverable due to their size.

2.1.2 Ultraportable systems for on-site measure

To answer our questions about Iznik ceramics, the HSI system must not only be maneuverable but also portable. It must be able to be used *in-situ*. In this context, *Specim Ltd.* company developed an ultraportable pushbroom camera: the **Specim IQ** camera, commercially presented as the first ultraportable compact hyperspectral camera (see Figure 2.3). The instrument (mainly employed for agriculture and food analysis applications [25]) weighs 1.3 kg and measures $207 \times 91 \times 74$ mm. The camera is covering a wavelength range between 400 and 1000 nm, with a spectral resolution of 7 nm (204 values) and a spatial resolution of 512×512 pixels per image. Thanks to its mobility, it can be used both indoors and outdoors, using controlled lighting with halogen lamps or acquiring images in full sunlight.



Figure 2.3: Vectorial visualization of the *Specim IQ* with annotations and dimensions on (left) and commercial pictures (right). Figure taken from [25].

This camera has many advantages, due to its compact appearance and its simple use which makes it easily transportable and usable outdoors. Pushbroom technology guaranties a great spectral resolution but the spatial resolution falls short of expectations.

As it is underlined by Behmann *et al.*[25], the images acquired in the violet/blue band (from to 400–450 nm) and in the near infrared band, between 900 and 1000 nm, are quite noisy. In the range 900–1000 nm, the observed noise is attributed to the lightning/scattering effect, which seems to be caused by the atmospheric absorption bands in the range 920–980 nm (c.f. [81]). Since other CCD-based hyperspectral cameras (designed for both airborne and stationary applications) do not suffer the same problems, this might be considered structural to the IQ camera and its detector sensitivity. Behmann *et al.* [25] proposed a solution that utilizes a spectral flattening filter with increased integration time when using the IQ camera under natural light.

It is interesting to note that the poor quality of the NIR spectrum between 900 and 1000 nm depends a lot on the lighting environment. This was confirmed by a test I carried out on glazed ceramic tiles from Iznik, analyzed under different conditions to compare spectra acquired in both artificial and natural light (c.f. [192]). Photographs of an Iznik tile were taken in a completely closed room (which could be compared to controlled laboratory conditions) with two halogen lamps. Despite the fact that the selected tile was not the same as the one photographed outdoors in Cairo, the pigments were similar and present almost the same spectra, but the outdoor acquisition was characterized by noise in the corresponding IR bands.

Figure 2.4 shows that, in the open air and natural light, the spectrum is very noisy, even after spatial averaging over the region of interest of 5×5 pixels. We can also notice a sharp characteristic peak of H_2O absorption between 925 and 970 nm [25, 81]. For measurement in a dry environment, the H_2O absorption peak disappears [25], and the noise in the IR region is attenuated: it is possible to eliminate it by carrying out a spatial averaging, even if it remains slightly visible in the spectra of the single pixels.

However, even in the most favorable conditions, it is important to temper the effect of spatial averaging. Indeed, the spatial resolution of the IQ camera is already very low. For attempting a classification of the materials with coarse details, we can tolerate a relatively low signal-to-noise ratio and even lose details by averaging over a 5×5 area (thus reducing the resolution of the camera to a mere 0.01 Mpixels). On the other hand, when the details are important, spatial averaging and noise in the measurement may result in less tolerability. In that case, the low sensitivity at the extremes of the spectral range can be reduced by using an illuminant providing a lot of blue and IR radiation. Then the reflected signal would be higher, and the quality of the images would be better.

2.1.3 Notion of metamerism on colors

Despite having a low spatial resolution with the Specim IQ camera, it remains more suitable than a classic camera for the Iznik tiles study. Indeed, hyperspectral imaging provides indubitably more information than RGB imaging in the context of automatic image processing because RGB imaging suffers from **metamerism**.

Metamerism is when two materials with different spectral distributions result in two identical colors under a particular illumination [218, 106] (see Figure 2.5). This happens when the projection of the two spectra into the three fundamental bands of blue, green, and red result in an identical RGB color.

Two different materials may have the same color but a different spectrum and while it would be impossible to differentiate the two materials automatically with an RGB image, the process can be done by adding more information with a set of hyperspectral images.

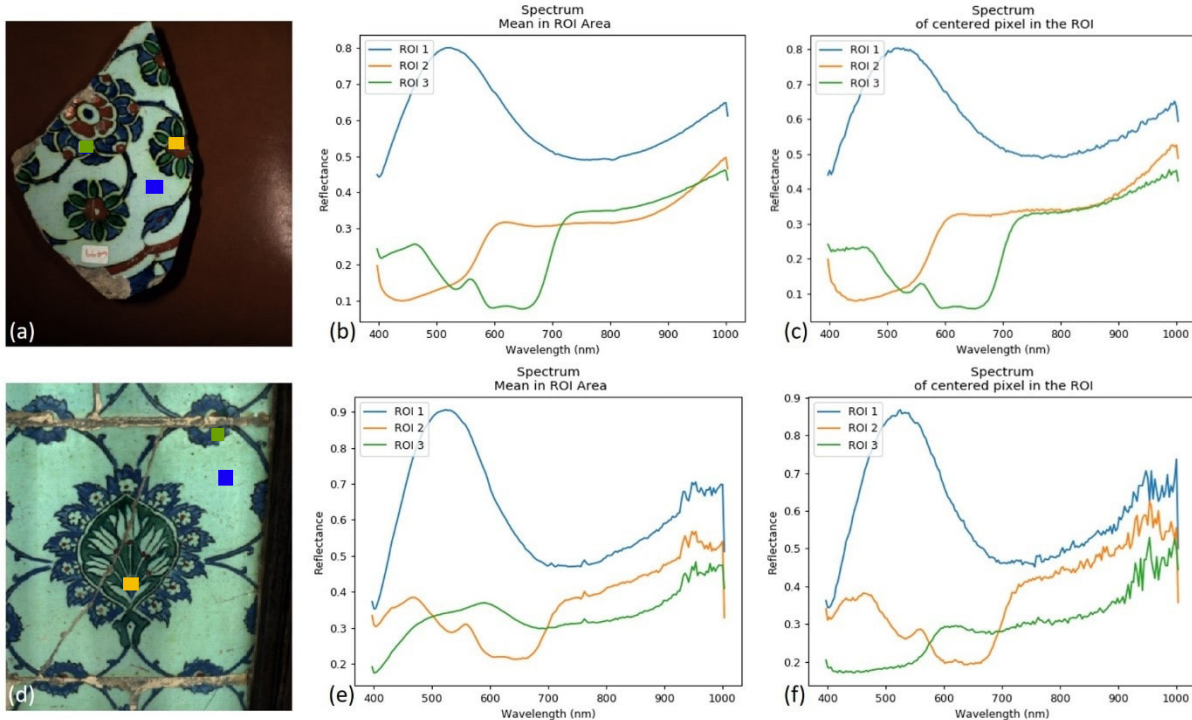


Figure 2.4: Comparison of the spectral image of a fragment of Iznik ceramic in (a) a controlled dry environment and under artificial light and (d) a tile containing the same pigments in the open air, under natural light. For the regions of interest containing the blue and red pigments, as well as the background, we observe the spectrum averaged over 5×5 pixels in a dry, controlled environment under (b) artificial and (e) in the open air under natural light. Likewise, we observe the spectrum of these elements (c) in a controlled dry environment and under artificial light and (f) in the open air under natural light in the case of a single pixel (no averaging). These results were published in *Journal of Field Archaeology* (c.f. [192]).

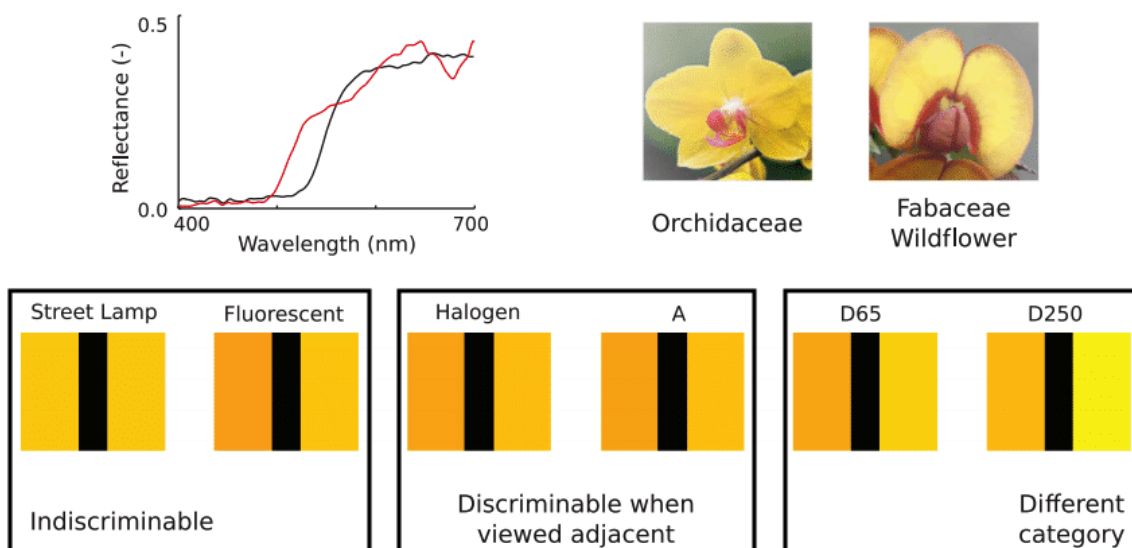


Figure 2.5: Illustration of one metameric pair under different illuminants. Color differences are reported for the pair of reflectance spectra under different illuminants. Figure taken from [10].

2.1.4 Hyperspectral imaging data processing

HSI appears to be a better technique for the study of materials (recognition, segmentation, etc...). However, hyperspectral image processing seems to be very challenging due to the large amount of data to process. That is why it is important to separate in two parts the process: dimensionality reduction and (frequently) spatial segmentation.

Dimensionality reduction

HSI processing differs from classic image processing because of the large amount of data. An RGB image can be simplified as a data-cube with 3 spectral values. By comparison, a small pushbroom camera like the Specim IQ consist of 204 spectral values, which is 68 times more. As a consequence, HSI processing requires specific stages that can be mostly sum up to reduce the amount of data without losing spectral accuracy.

The most common procedure for HSI processing is Multivariate Image Analysis (MIA) [77, 85]. In this method, the data-cube is separated along 3 axis: x and y coordinates for the spatial dimension and z for the spectral bands (in wavelengths) covered by the device. Image processing then consists of reducing the spectral information along z-axis by highlighting the relationships between variables and compressing the dataset. This reduction is usually done pixel-by-pixel or, in some case, for the whole image [174].

The most popular MIA model is the Principal Component Analysis (PCA), which is computed on the entire data matrix [65] transformed into a 2D dataset (one spectral dimension to reduce, and a spatial one). PCA is generally applied for dimension reduction by means of a linear transformation. This transformation results in a new coordinate system that effectively captures most of the variance in the data with fewer dimensions than the original dataset. The new coordinate system can be seen as a new orthogonal basis of components where different individual dimensions of the data (i.e., the spectra in our case) are linearly uncorrelated. It is frequently used in HSI to detect general patterns in the data and identify outliers.

It can also be used to reduce the z-dimension to a 2D plot where single pixels are classified according to their spectral features and the relationships among them, for simple visualization purpose (see Figure 2.6). In this case, pixels presenting similar spectra can form more or less scattered clusters of points.




Figure 2.6: PCA model calculated on the image of Roman frescoes fragments from an excavation in Pisa (Italy). On the left the Original RGB image, in the middle the PCA scatter plot in which shows clusters corresponding to the background, the white reference, labels and the frescoes. On the right, the graph presenting the two main components as the x and y axis. Image given by Claudia Sciuto (University of Pisa) for personal use.

Spectral segmentation

Applying some spectral segmentation on the data-cube is also a very common practice with HSI processing. It is pretty useful for materials recognition because each material of the scene is highlighted and summed up to one representative spectrum.

Traditionally, segmentation is done after applying PCA. It is applied on selected components obtained by PCA, while the wavelengths involved as variables for each component can be visualized as cluster centers results. The number of components used as well as the number of clusters depends on the the specific target/object investigated.

The image is then visualized as a posterized image in which the number of colors is the number of clusters. A new value is associated with each pixel according to the information of the clusters centers obtained (see Figure 2.7).

Process enhancing

The difficulties are often found in the choice of the segmentation algorithm as well as the choice of the reduction ones. If K-Means [140] remains the most common and fast to use, it is not always the most efficient. PCA also shows limitations by simply trying to reduce the dimensionality without dealing with the illumination level or the wavelength value of each frame.

That is why Zhao *et al.* [223] add an additional step for pigments segmentation application with pigments mapping. First, data reduction is made by transforming the spectra in the CIELAB color space before a classic K-Means segmentation. Then, each pigment within each cluster is mapped with a database of spectra applying Kubelka-Munk theory [119] to obtain a concentration map. The problem encountered here with this pipeline is that reducing into a 3-dimension color space (like RGB) can generate metamerism as well. Moreover, PCA remains more interesting for data reduction because the bases used are based on the components obtained from the hyperspectral image itself. So it can handle more metamerism-like situations.

In some cases like [55], the reduction and K-Means are replaced by Spectral Angle Mapper (SAM) [118]. It consists in a spectral classification that computes an angle between the spectrum of each pixel considered as a n -dimensional vector, and some reference spectra (endmembers) chosen from an internal or an external database. The lower the angle, the better the similarity between the spectrum and its endmembers. At the end, the clusters are based on each of the endmember, and the segmentation is made by choosing for each pixel the best angle (i.e. the best endmember). The low sensibility to illumination variation and a more physically based similarity computation appear as clear advantages of this method, but the difficulty becomes visible in the choice of the references as they are a very sensible key for segmentation.

Improvements are then also done on the reduction part. Studies seek to replace PCA with more suitable methods. For example, Pouyet *et al.* [173] demonstrates that a reduction method based on t -distributed stochastic neighbor embedding (t -SNE) [209] gives a much better mapping than with PCA.

Aletti *et al.* [12] then proposes to improve both with a semi-supervised 3-steps method. A Regularized Linear Discriminant Analysis (RDLA) reduces the data-cube dimensionality instead of a PCA. Then, a special spectral/spatial similarity metric is developed to establish a similarity between points. Finally, segmentation is executed with a Random Walker method [84]. The main interest of this algorithm is that it takes into account the distance between pixels: two identical but distant pixels and two close pixels with a less similar spectrum present a similar weight. The global process is maintained as a dimensionality reduction followed by a segmentation.

Finally, with the rise of Deep Learning solutions, some hyperspectral mapping process are composed with neural networks. For example, Han *et al.* [98] develop a classification model based on deep learning (3D-CNN and GLCM) and spectral-spatial joint feature for sea ice avionic remote images. Ma *et al.* [139] present a quite complete review of Deep Learning for remote imaging applications including HSI. Their study presents the different existing neural models for




Figure 2.7: *Illustration of segmentation results from [12]. Result of the segmentation of an aerial view. From left to right: image of the landscape in false colors, labels of the ground truth, final result.*

image fusion, registration, segmentation or even object-based image analysis. We therefore find even in the deep learning models the desire to solve mapping problems for a lot of applications.

2.1.5 Applications of hyperspectral imaging on Cultural Heritage

These pipelines are very appealing for Cultural Heritage application as they unite imaging methods with non-invasive analytical aptitudes to investigate and map materials properties [67]. That is why we can observe a large number of examples and adapted pipelines for each of them.

A large case of applications and methods

Following the evolution of HSI coming from airborne systems, many of the first cultural heritage applications were made for remote observation of the earth integrated into airplanes [145] or even satellites [40, 164]. At this time, the main role of HSI was akin to exploration on vegetation and soils to find potential archaeological sites. However, some applications shows the use of avionic hyperspectral devices employed for identification and mapping of pictorial materials like [188] on painted walls in the Pompeii site.

Slowly, cameras have become more and more accessible and have begun to be used for smaller objects like paintings [72, 51], illuminations [149, 151], textiles [55, 11, 172], ceramics [26], graphic documents [50], etc. Identification is now made from the simplest set-up with few tunable light sources [114] to the most complex HSI technologies like FTIR devices [221].

Many of the cases have the recognition of material (e.g., pigments or dyes) as their objective. Identification can then be made with HSI alone [51] or combined with other non-invasive methods such as X-Ray fluorescence [57] and Raman [149], or spectrofluorimetry [152]. These different imaging technologies can be compared to each other [149] or they can be observed together with

data fusion [68]. Although HSI and X-Ray fluorescence (XRF) imaging spectroscopies alone lack of robustness, combining both improves results using complex data processing [72].

Then, the range of wavelength has been studied for more results: SWIR domain has demonstrated the interest of the hyperspectral in addition to the simple spectral by providing signatures of molecular vibrations for pigments and dyes [50], but also wool and silk fibres [58]. Even going deeper by adding the mid-infrared (350 to 25 000 nm) presents excellent mapping with SAM methods [72].

Adapted methods of visualization

The growing use of HSI for cultural heritage leads to the implementation of many data visualization methods adapted to their various and specific goals. Salerno *et al.* [188] actually presents mainly 4 types of visualizations developed in cultural heritage: *Blind separation methods*, *False color*, *Chromatic derivative imaging* and *Neural Networks Analysis*.

Blind separation methods are a generalization to what we presented about dimensionality reduction (see Figure 2.6). Indeed, a data-cube has too many frames (i.e., an image for a given wavelength) and it is impossible to view it frame by frame. So, we try to decrease its dimensionality down to a few values, like what we explained for the PCA. An alternative to the PCA is the Independent Component Analysis (ICA) [206] for which components are no longer uncorrelated but independent. Then, two visualizations can be developed: the one we presented in Figure 2.6 and simply displaying each new generated frame corresponding to each component as a grayscale image.

However, the display of images taken from PCA components has two shortcomings. First, several images still have to be observed despite the reduction in dimensionality. Second, the images can be difficult to understand physically because the components have no physical or visual values. That is the reason why **False Color** has been developed. The idea is to display a classic color image similar to an RGB by changing the 3 channels in order to highlight peaks or variations for specific values. Typically, we start from four bands: Blue (B: 400-450 nm), Green (G: 450-550 nm), Red (R: 550-650 nm) and Infrared (Ir > 700 nm). Then, each band is shifted: the Blue channel displays Green, the Green channel displays Red, and the Red channel displays Infrared. It is what we call IrRG false color imaging (see Figure 2.8).

Numerous other false color images were created such as IrGB False color (the channels are respectively Infrared, Green and Blue), or $\bar{I}rRG$ in which the Infrared band is substituted by its negative. Grifoni *et al.* [87] even proposes to enhance the False color technique by merging Ir and RGB data with a Gradient Transfer method, in order to keep RGB and Ir data in the same False Color image. Accessibility and adaptability are great benefit of this kind of visualization but the importance of the spectrum tends to be under exploited.

Another variation of False color including the four bands (Blue, Green, Red, Infrared) is the **Chromatic derivative imaging** (ChromaDI) introduced by Legnaioli *et al.* [124]. The four bands information are included into the False color image by subtracting each consecutive couples of spectral images corresponding to each band (see Figure 2.9). ChromaDI enhances classic False color imaging by including 4 bands in a 3 channel image: this is a reduction to 3 or 4 dimensions as for RGB. The chosen dimensions can thus still be subject to forms of metamerism and the key information of the pigments could remain unclear. However, the False color drawbacks still persist.

Finally, Salerno *et al.* [188] mention **Neural Networks Analysis** that can be used for dimensionality reduction but, above all, for automatic segmentation.

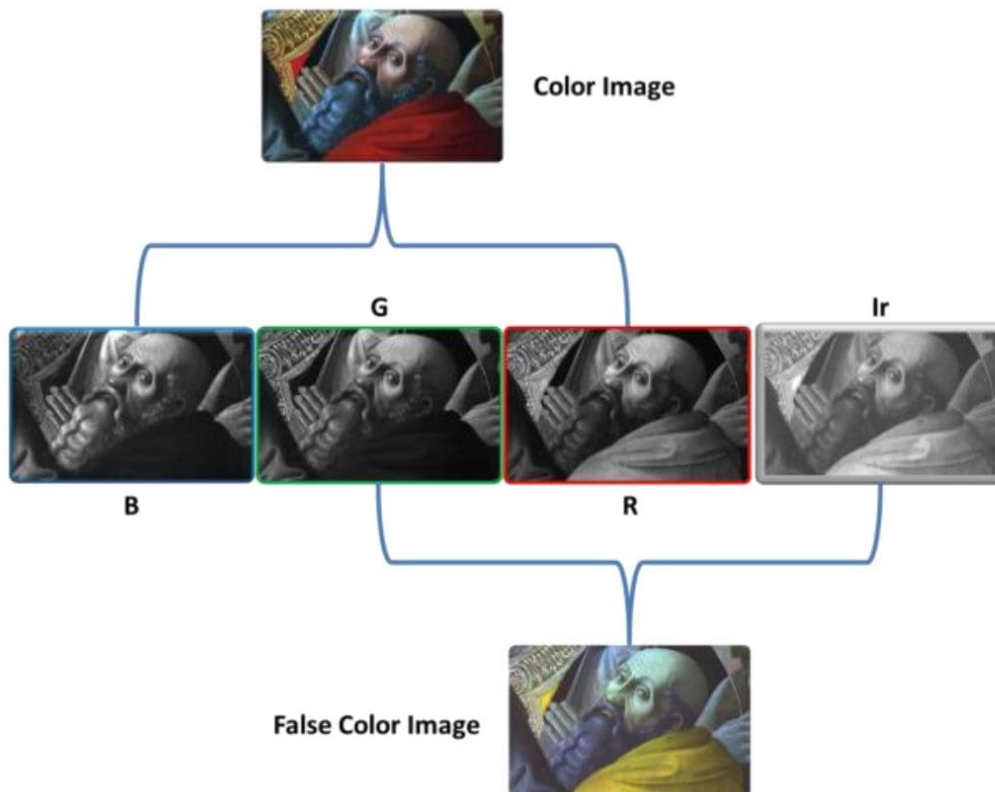


Figure 2.8: Color (RGB) and False Color Images ($IrRG$). Figure taken from [188].

Machine Learning for automatic segmentation

With the emergence of Deep Neural Networks and Machine Learning, many perspectives soar for Cultural Heritage applications. Making the identification and mapping methods fully automatic seems very appealing. In this spirit, Kleynhans *et al.* published two articles of automatic mapping hyperspectral images of paintings. In [115], they compare SAM based-method using a manual pipeline (ENVI-SHW software [2]) with an automatic one. The process can be divided in three parts. First, the endmembers are automatically determined by an algorithm called maximum distance (MaxD), which is based on the convex hulls theory. Then, the classification is done identically to the SAM method. Finally, the recognition is also automatized with a spectral library containing the spectral signatures of 48 pigments used with some spectral matching algorithms. The results show that mapping and identification coming from the traditional method and the automatic one are very similar. Automatic methods, which are faster and easier to use, seem therefore promising.

For this reason, Kleynhans *et al.* goes further with a neural network method for identification and mapping in paintings [116]. To simplify, they train and use a one-dimensional (spectral) convolutional neural network (see Figure 2.10).

Results presented are really good for the mapping as well as the identification. They are clearly better than in the previous presented methods. However, authors admit that “*existing spectral databases are small and do not encompass the diversity encountered*”. This is the biggest limitation of deep learning methods. They need a huge amount of data to feed the network, that are representative of all the cases we can encountered. A simple colour chart is not enough and for this publication, authors had to go to schools of artistic practices in order to get painting data. It needs precise intuitions on materials and perhaps may be incomplete because of the possible degradation of pigments and dyes throughout the time. Moreover, a trained CNN model works only within the limitations made by the input data used to feed and train it. In




Figure 2.9: *ChromaD Image construction. The Ir, R, G and B channels are separated. 3 new channels are made by subtracting B from G, G from R, and R from Ir. The image is then visualized by projecting the 3 channels as False Color Imaging. Figure taken from [188].*

other words, a network trained for the types of paintings presented here cannot be used for a different type of painting.

At this point, we can declare that despite being very appealing, automatic methods are not so perfect because of the requirements to use them. A pipeline created for one specific case will surely not work for a close but different one. It can even be more frustrating not to control any parameters at each step. For these reasons, we decided to focus more on classic pipelines composed of dimensionality reduction and segmentation, which can be more or less semi automatic.

Pigments studies on ancient ceramics

After getting a global look of HSI for cultural heritage, we focus on the case of ancient ceramics in order to study our Iznik tiles. We have seen in this part that the glaze thickness has a too wide standard deviation (even within a single tile) to carry out comparisons of similarities. At the same time, we assume that the drawings were made by hand. The differences visible on the patterns cannot be used to discriminate the tiles then. Finally, we make the assumptions that variations in the use of pigments seem to be the best way to classify or compare ceramic tiles. The recipes were indeed much more similar on the same production because the pigments produced will be used on a large number of ceramics in series contrary to the pattern and the thickness that are applied tile by tile.

Furthermore, we globally know the type of pigments used for each of the colors visible on Iznik tiles. Ben Amara [26] gave a lot of information about the composition of each of the pigments. First, the superficial glaze is slightly tinted by copper (Cu^{2+}) that can modifies the




Figure 2.10: Workflow for the creation, validation and testing of the trained 1D-CNN model. Image taken from [116].

perception of the different colors of the decoration. This tint may be due to an artistic choice or a slight degradation with the diffusion of blue or green. These two pigments are also visible in the drawings with an higher concentration and are mainly composed of copper (Cu^{2+} , *Copper Blue* – light blue – and *Copper Green* in the following) and cobalt (Co^{2+} , *Cobalt Blue* – dark blue –) respectively. The black decor used to outline the drawing lines is made of chromite and does not diffuse but forms a distinct layer of the glaze. At last, the red one is formed by a distinct layer sandwiched (like the black) between the siliceous engobe and the superficial glaze. It presents a heterogeneous texture that would correspond to ferruginous sand. In the following, we will define it as *Iron Red*. In our case, it is worth noting that the red/brown color is less red than those from the first productions of Iznik, and it reflects what is known as the technological decline of this ceramic.

In summary, we know that the tiles possess the same pigment families, but we want to identify the disparities between them so that we can observe their similarity. Even so, it is noted that the pigments are analyzed by spectral measurement or by HSI in the literature. For example, Cosentino [47] uses Fiber Optics Reflectance Spectra method to produce a database of 54 spectra of historical pigments commonly used in art work (list of pigments in the database is visible in Figure 2.11). Khampaeng *et al.* [114] go further and develop a tunable lighting imaging system with multicolored LED used for the identification and the segmentation of 357 Kremer pigments [79]. They assert that twelve well optimized tunable single-wavelength LEDs between 400-700 nm are sufficient. And we effectively show on our own investigations that about ten to fifteen values on the spectrum are enough for the hyperspectral analysis of ceramics (see Annex II).

On the one hand, some publications use hyperspectral data in order to generate databases while on the other hand, some others focus more in detail on a particular pigment. This is the case of cobalt blue for example, because it has a very attractive color but remains expensive to produce in addition to a non-negligible toxicity. So we try to replace it with other blues, such as Co-olivine blue pigment [134] or lanthanum-strontium copper silicates pigments, which are toxic metal free blue color inorganic pigments [110]. The precise comparison can then be completed associated with SEM [134], or with multispectral techniques alone in order to see if the color remains the same according to any illuminant.

Finally, the work of Farjas *et al.* [66] seems very interesting to us because it combines spectral data with a 3D model of Nazca polychrome ceramic vessel, which links our two chapters together. In addition to perform on hyperspectral 3D mapping of high precision, their results encourage us to use hyperspectral methods on ceramics for recognition and comparison purposes.

White	Yellow	Red	Green	Blue	Brown	Black
Lead white, 46000	Cadmium yellow, 21010	Alizarin, 23600	Cadmium green, 44510	Azurite, 10200	Burnt Sienna, 40430	Ivory black, 12000
Zinc white, 46300	Cobalt yellow, 43500	Cadmium red, 21120	Chrome green, 44200	Blue bice, 10184	Burnt umber, 40710	Vine black, 47000
Lithopone, 46100	Lead Tin yellow I, 10100	Red lead, 42500	Cobalt green, 44100	Cobalt blue, 45730	Van Dyke, 41000	Bone black, 47100
Titanium white, 46200	Lead Tin y. II, 10120	Red ochre, 11574	Green earth, 11000	Egyptian blue, 10060	Raw Sienna, 17050	Lamp black, 47250
Gypsum, 58300	Massicot, 43010	Vermilion, 10610	Malachite, 10300	Indigo, 36005	Raw umber, 40610	
Chalk, 58000	Naples yellow, 10130	Madder lake, 372051	Phthalo green, 23000	Maya blue, 36007		
	Orpiment, 10700	Lac dye, 36020	Verdigris, 44450	Prussian blue, 45202		
	Saffron, 36300	Carmine lake, 42100	Viridian, 44250	Smalt, 10000		
	Yellow ochre, 40010	Realgar, 10800		Ultramarine, 10510		
	Yellow Lake, 36262			Phthalo blue, 23050		
	Gamboge, 37050			Cobalt violet, 45800		

Figure 2.11: *Distribution by color of the 54 historical pigments (Kremer Pigments) studied in this paper and respective pigment code. Table taken from [47].*

To sum up, HSI techniques turn out to be perfectly adequate for an *in-situ* portable use in the case of Cultural Heritage applications. By observing spectra that goes further than RGB information, it allows the analysis, the identification, and even the segmentation of materials using adapted algorithms for data reduction or visualization. Finally, HSI can also finally be applied to pigments and therefore seems to be very suitable for studying the reuse of Iznik ceramics from Saint-Maurice residence.

2.2 A new dimensionality reduction based on Pearson correlation

The previous section has underlined that many applications of the HSI in cultural heritage revolve mainly around two methodologies: the development of reduction/segmentation and new visualizations.

The hyperspectral data processing methods mentioned above give satisfactory results, but remain imperfect. We can notice it on complex materials like textiles for example. The Aubusson Tapestry (see Figure 2.16, c.f. [55]) is a good illustration: classic methods of reduction/segmentation such as SAM looks not accurate enough for a perfect dye identification. In order to improve it, I seek to improve the results by focusing on dimensionality reduction in this section.

Indeed, we can have as a starting point that the amount of data is enormous. One single hyperspectral image taken by the *Specim HS-XX-V10E* has a resolution of $2300 \times 1600 \times 840$, which means 3.09 billions of values (though this camera is a common HSI device for laboratory measurements). Therefore reduction methods appear as essential to process precise data correctly.

Nevertheless they still exhibit imperfections. We already know that reduction methods in color spaces pose problems because of metamerism. Thus, switching from hyperspectral data to RGB clearly cancel the advantage given by the HSI. This is exactly the same in $L^*a^*b^*$ space. The main problem here is that the reduction components do not take into account the data specificity.

This is the advantage of PCA, that is based precisely on this specificity to generate its components and carry out its reduction. However PCA can also give poor results because it takes into account the values of the spectrum and discards their positions (e.g., the wavelength values, see Figure 2.12). This poses two major problems. First, the illuminance on the sample matters enormously because the general intensity is taken into account as much as the slope

variations. Second, the difference calculation does not take into account the spectral distance, which makes it very vulnerable to measurement noise and generates physically unrealistic results.

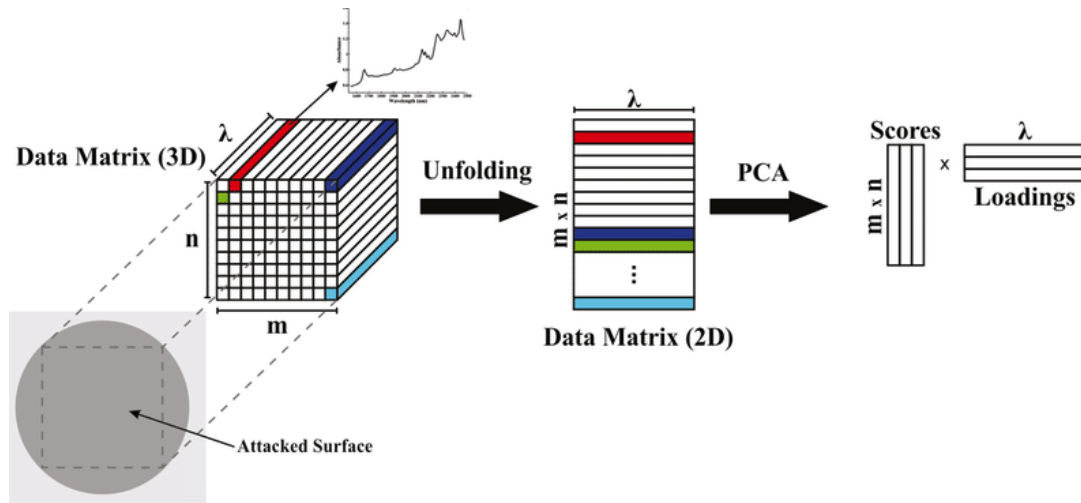


Figure 2.12: Schema representing the application of PCA in HSI. The data-cube is unfolded into a 2D matrix and then reduced into a new matrix. As we can see, the wavelength values are not taken into account. Figure taken from [163].

SAM, the last widely used method we have spoken about, seeks precisely to perform calculations that are more physically accurate. The differences are then more precise. Moreover, the idea of using endmembers from internal or external base is very good. However the choice of these endmembers still remains extremely sensitive on the results. In addition, illuminance is also a bit of a problem here.

In any case, it shows us that the reduction must be done using an internal or external basis. So we have chosen to start from this point and develop a more accurate way to compare 2 spectral curves, as a whole, without taking into account their overall intensity (i.e. illuminance).

In this section, we present a novel technique for the similarity study of pigments and dyes spectra with three axis of reflection:

1. A look about noise pre-processing on the hyperspectral images with some low-pass filtering,
2. Some interactive visualizations that make easy to identify and compare spectra or group of spectra,
3. A new dimensionality reduction method using a set of characteristic spectra and based on the computation of Pearson correlation coefficients (PCC).

2.2.1 Description of the pipeline

Our approach relies on three main steps: the set generation, the pre-filtering of the studied data, and the dimensionality reduction using Pearson correlation. The complete pipeline is presented in the figure 2.13.




Figure 2.13: Our pipeline: Top: the set of characteristic spectra is obtained from hyperspectral images. Bottom left: our studied image is filtered for denoising. Finally, the studied spectrum is extracted and a Pearson correlation coefficient is computed between it and each of the characteristic spectrum.

Set of characteristic spectra generation. As the first step, we create a set of characteristic spectra. The role of the characteristic spectra is to serve as a comparison to the spectra measured. They must be as representative as possible of what is being studied. For example, if we want to observe cobalt blue pigments as it is examined in [134, 110], the best set would be a set including a wide variety of cobalt blue, by their composition or their concentration in the binder. Moreover, these spectra can come from an internal or an external database. This means that it is possible to create particular color charts with specific pigments (external database) or we can select spectra from our data that look representative (internal database).

Usually, a characteristic spectrum c is extracted from a small area of several pixels (it can be 5×5 like 100×100 pixels) from a hyperspectral image \mathcal{C} . Hence we extract N characteristic spectra $c_{n=[1:N]}$ from T hyperspectral images $\mathcal{C}_{t=[1:T]}$. Of course, the extraction can be achieved from one ($t = T$) or more images. At the end, we get a set of N characteristic spectra that will be used for the reduction.

Filtering for noise correction. Usually, the studied spectrum s is extracted directly from a pixel of the studied hyperspectral image \mathcal{I} . It then presents significant variations due to measurement noise (i.e. the set of spurious signals that are superimposed on what we are acquiring), but also to inhomogeneities present on the surface corresponding to the projected pixel. The graph on the left of Figure 2.14 illustrates it. Noise can be very problematic because it induces significant variations and differences, which distorts the reduction, and then the segmentation.

So rather than extracting directly from \mathcal{I} , the spectrum s should instead come from a filtered image $\tilde{\mathcal{I}}$ (see the right graph on Figure 2.14). This hyperspectral image $\tilde{\mathcal{I}}$ is obtained from \mathcal{I}




Figure 2.14: *Noise reduction.* Here, the main interest is the comparison of the tapissery spectra in red. **On the left**, the red spectrum from a given pixel obtained from the original image. **On the right**, the red spectrum of the same pixel from the filtered image.

by applying some low-pass filtering. These denoising processes are generally based on weighted averaging the pixels values (i.e. the spectra) close to the analyzed one. Several spatial processing methods can be used, and this is the great advantage of this step. Some examples are Gaussian filtering, Bilateral filtering (see section 2.3), or even simply extracting an homogeneous area of several pixels manually and average them (see section 2.4). After this step, $\tilde{\mathcal{I}}$ is denoised and we obtain the studied spectrum s we want to reduce.

Dimensionality reduction with Pearson correlation. As a last step, we carry out the dimensionality reduction on s . For each n of the N characteristic spectra c_n , a Pearson [167] correlation coefficient ρ_n is computed between s and c_n :

$$\rho_n(s, c_n) = \frac{\text{cov}(s, c_n)}{\sigma_s \sigma_{c_n}} \quad (2.1)$$

where:

- $\text{cov}(s, c_n)$ is the covariance between the spectra s and c_n .
- σ_s is the standard deviation of spectrum s
- σ_{c_n} is the standard deviation of spectrum c_n

The obtained values of $\rho_{n=[1:N]}$ are between -1 and 1. The closer it is to one, the more similar are the two spectral curves.

We finally get N coefficients ρ_n . Based on our observations, a set of around 20 characteristic spectra is more than enough for good results. In comparison, *Specim IQ* has 204 wavelengths measured (so we have 10 time more values) and *Specim HS-XX-V10E* has even more with 840 (42 time more) for each pixel. So data is well reduced with regularized values in our cases, and more relevant.

2.2.2 Reducing the data for a better relevance

Our Pearson Correlation Coefficient Reduction (PCCR) method should generate better results than a classic PCA because similarity with Pearson correlation is more consistent for curves comparison whereas PCA is usually working on decorrelated data. We can focus on mathematics to explain it.

Note: Mathematical explanations

Indeed, given that our spectrum has a total number of m values, correlation between s and a given characteristic spectrum c_n can be rewritten [167]:

$$\rho_n(s, c_n) = \frac{1}{m-1} \sum_{j=1}^m \left(\frac{s_j - \bar{s}}{\sigma_s} \right) \left(\frac{c_{n,j} - \bar{c}_n}{\sigma_{c_n}} \right) \quad (2.2)$$

where:

- $m, s_j, c_{n,j}$ are respectively the number of wavelengths used, the individual sample points indexed with j of s (resp. c_n).
- \bar{s}, \bar{c}_n are respectively the mean of the spectrum of s and c_n
- $\sigma_s = \sqrt{\frac{1}{m-1} \sum_{j=1}^m (s_j - \bar{s})^2}$ is the sample standard deviation of s (and similarly σ_{c_n} for c_n)
- $\left(\frac{s_j - \bar{s}}{\sigma_s} \right)$ will be next called the standard score of s (and similarly for the standard score of c_n).

A standard score is computed for s_j and $c_{n,j}$ for each measured wavelength j . This standard score is basically a regularized difference between s_j and the mean of s (similarly $c_{n,j}$ and the mean of c_n). Finally, a Pearson correlation coefficient computation works by adding term by term the product of the standard scores of s and c_n . For each wavelength j , the closer the standard scores of s and c_n are, the bigger will be $\rho_n(s, c_n)$. In addition, the normalization, thanks to the division by the sample standard deviations, ensures to keep a value between -1 and 1.

Concretely, it is a very interesting approach for the comparison of two spectra. First, we can focus on a slope/peak shift. A reflectance spectrum of pigments shows characteristic “peak” and “slope”. It is these peaks that we want to look at: we see how they shift or grow between two samples for example. Using PCC, a slight shift of a slope on the studied spectrum s will create a shift in the differences of its standard score term by term. The overall coefficient will drop a little. If this shift increases, the coefficient will decrease accordingly (see Figure 2.15A). This coherent evolution between the position of a peak or a slope and the value of the correlation coefficient is very useful for our study of spectra. We actually know that for the same pigment, manufacturing variations could lead to shifts in these characteristic peaks and slopes [79].

The second big advantage is the fact PCCR compares variations. The standard scores are based on differences and the total value is normalized (by σ_s and σ_{c_n}). Therefore the coefficients of two spectra i presenting the same variations, peaks, or slopes will have the same coefficients $\rho_{n=[1:N]}$ (see Figure 2.15B). A variation of their average will not change anything. If we have a case with two identical pigments but a different global illuminance (or a different intensity for some cases), the final result remains unchanged.

The use of Pearson correlation coefficient for the similarity studies of pigments appears as ideal, or at least as better than previous methods.




Figure 2.15: Illustrations of the advantages of PCCR method. **A:** A growing shift of a slope will make the Pearson correlation coefficient decreases accordingly. **B:** The PCCR method is not sensitive to variations of global illuminance.

2.3 Experimentation of the method on dyes on Aubusson tapestry

It is necessary to validate the PCCR method before applying it on the Iznik ceramics of Saint-Maurice residence. That is why we used the hyperspectral images of Aubusson tapestry coming from the work and data of De la Codre *et al.* [54, 55, 56] and then compare our results for a confirmation.

2.3.1 Context

Tapestry is a historic form of textile art that is crafted by hand on a loom. It involves weft-faced weaving, wherein the warp threads remain concealed within the final work, setting it apart from other woven textiles where both the warp and weft threads are typically visible.

This art get its Golden Age in France and Europe during the seventeenth and eighteenth centuries. At the end of the seventeenth century, Jean-Baptiste Colbert gave the title of Royal tapestry Manufactories to three French manufactories in a context of flourishing activity: Gobelins, Aubusson and Beauvais. At this time, they are two main weaving techniques: high weaving practiced at Gobelins, and low weaving, at Aubusson and Beauvais. The three royal tapestry factories had each been schematically assigned to a position in terms of quality and subjects represented. To simplify: the Gobelins weave high-quality tapestries is for the king, Beauvais for the aristocracy and the Aubusson workshops are reserved for a less well-off provincial bourgeoisie [29].

In order to uphold the exact standards of the royal court, quality control procedures were implemented under the supervision of Crown inspectors and accomplished cartoon painters who were dispatched to the Manufactories [29]. Among the established regulations was the utilization of a meticulously curated selection of approved dyes, known as the *Grand Teint* dyes. It is in this context that De la Codre *et al.*[54, 55] seeks to identify the dyes used in some of the Aubusson tapestries: the tapestry studied by De La Codre in [54] and in this section is a *Grand Teint Verdure* (a large landscape representation) of 3.39×3.48 m commissioned by Count Brühl, Prime Minister of August III, King of Poland (see Figure 2.16).

Several non-invasive methods have been applied to materials (fibers, dyes, and mordants) on both faces, at different degradation states [43, 150, 55]. The tapestry was then acquired through a VIS-NIR CCD camera (Specim HS-XX-V10E) with a 1600×840 pixel resolution (pixel size: $8 \times 8 \mu\text{m}$), a 2.8 nm spectral resolution, and a spectral wavelength range from 400–1000 nm. For the camera, the projection of one pixel on the tapestry corresponded to 0.38 mm (surface mapped: $61 \times 32 \text{ cm}^2$). The illumination was made with two halogen lamps oriented at 45° at a working distance of 1.1 m for a global image. The hyperspectral image used below is also visible in Figure 2.16.

In the following, we focus on the identification of red and yellow dyes for analytical purposes.

First about the yellow dyes, De la Codre [55] demonstrated that the fine recognition (probably weld or dyer's broom) seems to be not clearly achievable in HSI because significant bands disappear or shift with degradation.

Second about the red dyes, madder and cochineal were the two most common red *Grand Teint* dyes. Madder was an affordable one whereas cochineal was very expensive and widely used in Gobelins. An identification of one or the other would then make it possible to make historical assumptions about the value of the tapestry and its owner.



Figure 2.16: *Left*, Entire tapestry (3.39 m \times 3.48 m) with the arms of Count von Brühl. Image taken from [54] *Center*, RGB image of the hyperspectral image studied in the following. *Right*, the part we focus on.

2.3.2 Method

As detailed in section 2.2.1, the method can be split into 3 parts: (a) the set generation, (b) the noise filtering and (c) the PCC reduction. Here we will detail the unfolding specificities on the tapestry, and we will add a last step: (d) the visualization of our results.

Set of characteristic spectra generation. One of the major elements presented by De la Codre in [54] is the creation of color charts composed by more than 600 samples of dyes. These samples were made on many colors: red, yellow, indigo, green, brown... They differ from each other in the type of pigments and fabrics used, as well as in their concentration and variation (see Figure 2.17).

These color charts were acquired through the same VIS-NIR CCD camera (Specim HS-XX-V10E) as the tapestry. The illumination was also made with the two same halogen lamps in the same configuration.

We acquire characteristic spectra from each of the samples on the color chart by simply averaging the spectra of the squares of textile (around 150×150 pixels). Two set are made (see Figure 2.17) since we decided to focus on red and yellow dyes to validate our pipeline. One set is a set of red dyes (30 spectra: 11 spectra of madder, 19 spectra of cochineal). The second set is a set of yellow dyes (34 spectra: 11 spectra of weld on wool, 12 spectra of weld on silk, 11 spectra of dyer's broom). This type of set is specified as from an **external database**.

Filtering for noise correction. Then, a low-pass filter is applied to the studied image \mathcal{I} visible in Figure 2.16 for the noise correction. We decide to use a **bilateral filter** [205] to

preserve contour while denoising. When the noise correction is done (see Figure 2.18) and, then the dimensionality reduction is applied with PCC method **on all pixels of the filtered image image $\tilde{\mathcal{I}}$** .

Note: Bilateral filtering

In image processing, it is quite common to use a low-pass filter to denoise an image. The most common filter is the gaussian filter. We replace each pixel \mathbf{p} at the position (x, y) on the image \mathcal{I} with a weighted average of intensity values from nearby pixels computed as follows:

$$\tilde{\mathcal{I}}(\mathbf{p}) = \frac{1}{W} \sum_{\mathbf{p}_i \in \Omega} \mathcal{I}(\mathbf{p}_i) g_s(\|\mathbf{p}_i - \mathbf{p}\|), \quad (2.3)$$

where:

- $\tilde{\mathcal{I}}$ is the filtered image
- \mathcal{I} is the original image
- W is a normalization term
- \mathbf{p} are the coordinates of the current pixel
- Ω is the window centered in \mathbf{p} , so $\mathbf{p}_i \in \Omega$ is another pixel
- g_s is the spatial kernel for smoothing differences in coordinates (a Gaussian function).

This results in what is called Gaussian blur (or Gaussian smoothing) in the image (see Figure 2.18). The filtered image is denoised by blurring, but the main drawback is that details and contours are also affected. That is why an intensity weight f_r is added in the bilateral filter [205]:

$$\tilde{\mathcal{I}}(\mathbf{p}) = \frac{1}{W} \sum_{\mathbf{p}_i \in \Omega} \mathcal{I}(\mathbf{p}_i) f_r(\|\mathcal{I}(\mathbf{p}_i) - \mathcal{I}(\mathbf{p})\|) g_s(\|\mathbf{p}_i - \mathbf{p}\|), \quad (2.4)$$

where f_r is the kernel for smoothing differences in intensities (in our case, it is also a Gaussian function).

The added weight is dedicated to the intensity variation. The filter denoises then large areas while keeping the contours (see Figure 2.18). However, this filtering has two shortcomings. It is first more complicated to configure because it is necessary to find the parameters of two kernels instead of one, which can be tedious by hand. In addition, it takes much longer to calculate.

Visualizations. The most appropriate visualization method for our case is a blind separation method (see paragraph *Adapted methods of visualization* in section 2.1.5), especially a PCA visualization visible in Figure 2.6. However, in our case, the reduction to 2 dimensions for the points density scattering is done on the PCC values $\rho_{n=[1:N]}$ instead of the intensity values from the spectra (i.e., s). It makes more sense because $\rho_{n=[1:N]}$ are more relevant and present decorrelated data. Obtained clusters are visible in Figure 2.19.

We developed ann interactive visualizations. The user can click on the PCA graph. Pixels corresponding to an area around the click are selected (called “selected area” for simplification). They are then highlighted in red on the RGB representation of the data-cube, and an average spectrum is extracted. It is therefore possible to find which cluster corresponds to the dye sought (see Figure 2.19 for an illustration of the interactive visualization tool).




Figure 2.17: RGB images of the hyperspectral images of color charts. Acquisition made by De la Codre [54]. **Red dyes:** left – madder, right – cochineal. **Yellow dyes:** left – weld on wool, center – weld on silk, right – dyer’s broom




Figure 2.18: *Top:* Comparison of the effect of applying gaussian and bilateral filtering. **Left:** a gaussian filtering is applied to the image: result is denoised but blurred. **Right:** a bilateral filtering is applied: result is denoised and contours are preserved. **Bottom:** Application of the bilateral filtering to the tapestry.

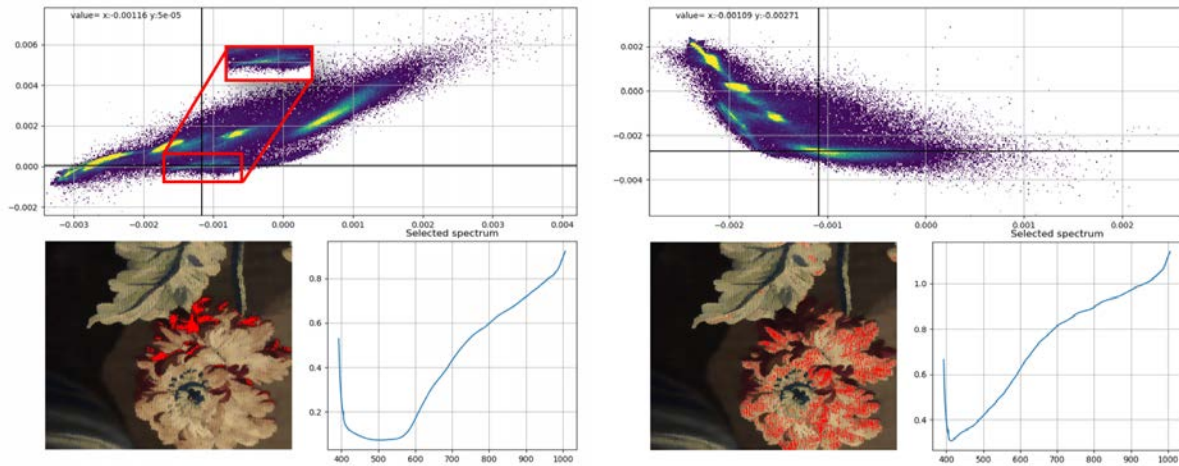


Figure 2.19: Illustration of two interactive visualizations for Aubusson tapestry (left: red dyes; right: yellow dyes). **Description of each visualization.** The upper part is the PCA graph. The black cross shows where the user has clicked. On the bottom left, the RGB representation of the data-cube with red pixels coming from the selected area. On the bottom right, the average spectrum obtained from this area. A detailed red window of the selected area is shown for the red dyes case.

2.3.3 Visualization on the color charts

Figures in section 2.3.2 illustrate the direct application on the tapestry. However we first tested the method on the hyperspectral images of color charts (visible in Figure 2.17) in order to have a simpler case. These data-cubes present appealing advantages for a first consideration as they are composed of homogeneous samples of dyes. That is the reason why we focus on them in a first time.

For the validation purpose, we compare the PCA graphs obtained with our method and a case where the PCA graph is obtained directly on the spectrum (i.e. a classic simple PCA reduction). For each case, a color chart hyperspectral image is chosen, associated with its corresponding set of characteristic spectra (red or yellow). The characteristic spectra are also placed on the PCA density graphs in order to locate and identify each possible cluster. This way to visualize can be used here because characteristic spectra are partially from an internal base: some of them correspond exactly to the spectra of the color chart. Some results are presented in the Table 2.1.

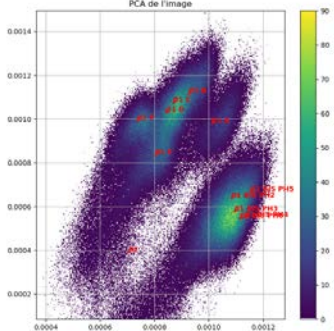
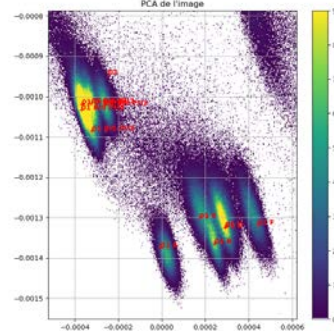
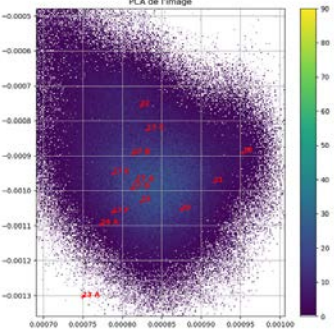
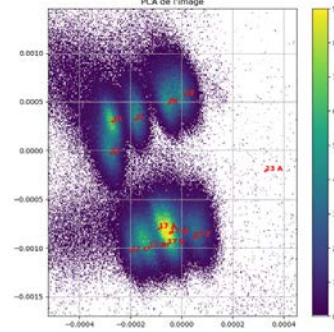
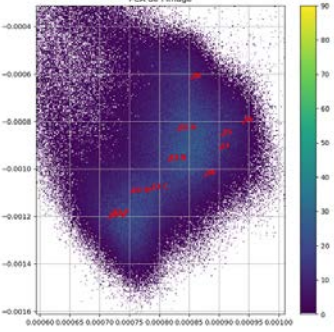
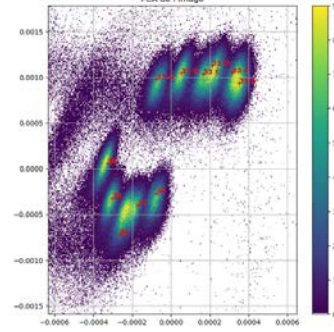
In each case, our method provides much better results. First, we see on the madder color chart hyperspectral image that a PCA on spectra (first row, third column) can generate global clusters of some different samples. With our, each of the samples are much more clustered. Only the *01 BIS* samples are mixed, but it can be explained by a very common composition.

Results are even more significant with the weld on wool and the dyer's broom color charts (c.f. second and third row of Table 2.1). In the case of a classic PCA reduction, points are scattered and no clusters are discernible. Whereas with PCCR, clusters are really visible and can be differentiated.

2.3.4 Dye identification

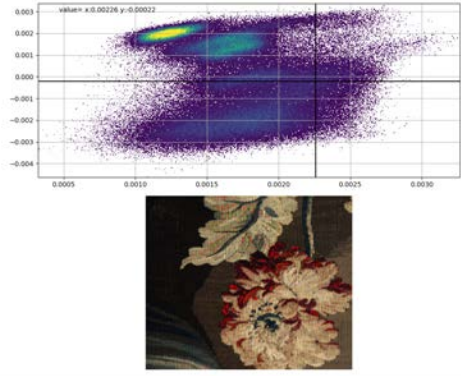
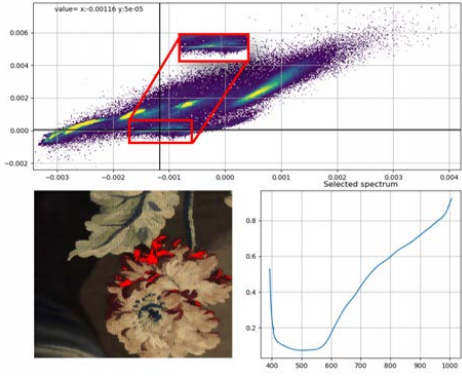
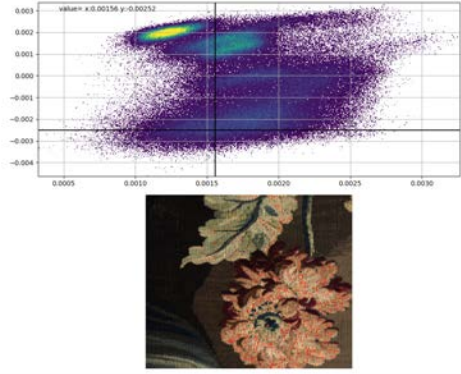
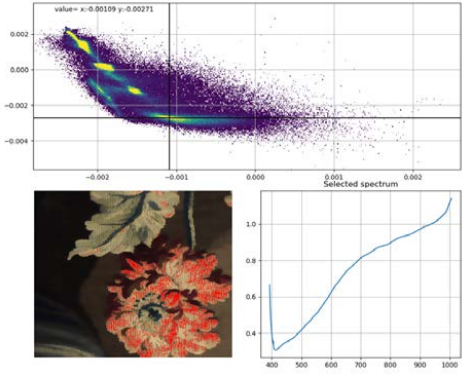
A similar comparison is achieved on the studied hyperspectral image of the tapestry. On one hand we apply our PCCR method as explained in section 2.3.2. On the other hand, a classic PCA reduction for visualization as comparison is done as presented in Figure 2.6. On this part, the characteristic spectra are also placed on the PCA density graphs. In both case, the visualization is interactive: the user can click on a desired area on the PCA graph, and selected

Table 2.1: Example of comparison between a PCA density graph directly on the spectrum or after PCCR on color charts hyperspectral image. As we can see, the points are much more clustered on the right. In addition, we can recognize each cluster by the dye used. Conversely, the points are very scattered on the left.

HS image	Set	PCA on spectra	PCA after PCCR
Madder	Red		
Weld wool	on Yellow		
Dyer's broom	Yellow		

pixels are highlighted in red on the RGB representation of the data-cube. Obtained results are presented in the Table 2.2.

Table 2.2: Comparison between a PCA density graph directly on the spectrum or after PCCR on the tapestry. **Left column:** PCA is done directly on spectra. Results are quite scattered. **Right column:** PCA is done after PCCR. Results are more clustered.

Set	PCA on spectra	PCA after PCCR
Red		
Yellow		

The first thing to notice is how the results underline the relevancy of the PCCR method. Indeed, we can see that visualization changes depending on the kind of studied dyes with PCCR whereas it remains the same for classic PCA. Hence we can adapt our data with PCCR for different purposes.

More importantly, when we look at the classic PCA, we can see that red and yellow spectra differ a lot. Their corresponding points are quite scattered along the graph. It is even complicated to highlight which part is corresponding to which color with the interactive visualization, but the PCA graph after our PCCR draws clear clusters for yellow and red. The dye identification is then much easier and more reliable.

Yellow identification. An average spectrum of yellow dye can be then obtained from the visualization. However, as explained by De la Codre *et al.* [56], significant bands have disappeared on the spectrum (disappearance of the chlorophyll absorption band at 670 nm). Therefore we cannot make the difference here between dyer’s broom and weld. But the generation of a representative spectrum on the tapestry appears as a good way to confirm the degradation.

Red identification. The average red spectrum extracted from the hyperspectral image presents mostly a strong broad absorption band from 410 to 560 nm. This is usually visible in the higher madder concentration dyes, which masks (c.f. [55]) the two characteristic absorption bands of madder (510 and 545 nm). We can then recognize typical shapes of madder dye instead of

cochineal. Crossed methodology operated by De la Codre [56] on fluorimetry hereby confirm the presence of madder.

In conclusion, the experimentation of the PCCR method on yellow and red dyes on Aubusson tapestry, and its comparison with a traditional PCA reduction, demonstrates its robustness and reliability to identify dyes in complex cases.

2.4 Application to groups of pigments on Iznik ceramics

The experiment on the Aubusson tapestry demonstrates the reliability and the robustness of the PCCR method for data reduction. Beside, dimensionality reduction is a key for the hyperspectral study of our ceramic tiles. Indeed, we are not studying a data-cube of one tile, but several dozen. The problem here is therefore more complex because it is about the similarity of tiles and not just pixels. Hence it is necessary to reduce the dimensionality as most as possible of each of them. For this, we adapt our method (as presented below). We made the hypothesis that each studied tile presents a so-called representative spectrum per pigment. It is these spectra that are compared in this section.

To achieve the similarity study, the acquisition have been proceed with the Specim IQ camera presented in Section 2.1.2 during 3 days in outdoor conditions. The daytime images were taken under ambient lighting (sunny) while the nighttime ones were taken under the illumination of a B10 Profoto light spot. Because of the location of Saint Maurice residence in an arid climate, the lighting condition during daytime was considered to be constant. In order to facilitate the comparisons of the results, the white balance was carried out on site using the same white reference. Each taken hyperspectral image shows one tile because of the low spatial resolution of the camera. We acquired 55 hyperspectral images of only 55 ceramics (8 tiles coming from P1 patterns, 36 from P2, 5 from C1 and 6 from C2 – see *Context* for more information).

We were not able to acquire all tiles because of the time constraints (very long exposure time) and the reduced distance from the camera to the ceramic necessary caused by its spatial resolution. To solve this problem, the use of a tunable filters camera seems to be a good solution. However, this requires a study of the number of minimum spectral bands before the results deteriorate. The results of this study are presented in Annex II.

In the following, we present our method and our comparison results on the tiles. For a better understanding, the RGB images and the data of each tile are visible in Annex III.

2.4.1 Method

We adapt the order of the method in relation to that of the tapestry. Indeed, the set is extracted from an internal base (the spectra of the tiles). It is therefore necessary to carry out the denoising and the extraction of the representative spectra of the tiles beforehand. These two steps are done simultaneously as explained below.

Tile representative spectra extraction. First, we generate a map for each tile acquired, that is we extract a representative spectrum of each pigment present on each tile. We apply PCCR on the extracted spectra. For this, on each tile and for each pigment, we manually select an homogeneous area presenting the pigment. The spectra included in the zone are averaged in a **representative spectrum**. This operation removes the measurement noise. We do this for the areas with pigments, as well as for the blank drawing areas (which we will call “background”).

The first problem encountered is that the spectrum does not actually represent the pigment, but the color of the pigment and the colored glaze present above (see Figure 2.20).

In order to solve this obstacle, we make several simplifications of the ceramic model. First, we assume that the glaze has a relatively homogeneous spectrum and thickness on the tile.




Figure 2.20: Schematic illustration of a ceramic tile with our assumptions. We see that the light passes through the glaze before being reflected by the pigment. The spectrum measured does not correspond of the pigment alone.

Furthermore, we simplify the ceramic by assuming that the pigment and the glaze are on two separate layers. This assumption is not true in practice but remains relatively correct in light of our results. Finally, we assume that the bottom of the ceramic is perfectly white.

The homogeneity hypothesis allows us to say that the spectrum of the glaze will be the same everywhere. So we have on the whole tile:

$$s_{measured} = s_{pigment} \cdot s_{glaze}, \quad (2.5)$$

where $s_{measured}$, $s_{pigment}$ and s_{glaze} are respectively the measured spectra, the representative spectra of the pigment, and of the tinted glaze. The assumption on the white background thus makes it possible to affirm that our representative spectrum of the background is in reality the spectrum of the tinted glaze:

$$s_{glaze} = s_{background} \quad (2.6)$$




Figure 2.21: Representative spectra mapping for the tile C2-1. **Left:** picture of the corresponding tile. **Right:** representative spectra obtained for the corresponding tile. Red, green, copper blue and cobalt blue are the representative spectra after the glaze correction. Glaze is the the spectrum of the tinted glaze.

From this, we finally get the representative spectrum of each pigment:

$$s_{pigment} = \frac{s_{measured}}{s_{background}} \quad (2.7)$$

After this operation, we get our mapping composed of the spectra of each pigment and the glaze for each tile (see Figure 2.21).

Set of characteristic spectra generation. Representative spectra identify the pigments used in the manufacture of Iznik ceramic tiles. We recognize the signatures of **copper blue**, **copper green**, **cobalt blue**, and **iron red** according to Ben Amara [26]. They are in agreement with common pigments used by Iznik production [26]. For each of these pigments, we then manually choose spectra to generate the set of characteristic spectra. These spectra are chosen because they appear to best represent the diversity within each pigment. For example, we only take one copper blue because its spectrum is enough for a global representation. We also add 3 spectra of the tinted glaze because it seems that the color of the glaze can also be used to discriminate the tiles. The set is visible in Figure 2.22.




Figure 2.22: *Set of characteristic spectra from internal database*

However, this choice may not be perfect because it is purely visual on the shape of the spectra. We can nevertheless try to partially validate it by observing if its chromaticity (its representation in a color space) seems representative of all our colors. We project the set of characteristic spectra in the XYZ color space, and observe their positions according to the representative spectra.

In Figure 2.23, we see that the characteristic spectra globally represent our values correctly. The representative spectra of cobalt blue lie on a line, taken into account by the corresponding characteristic spectra. The copper blue spectra are visually similar, justifying the use of a single characteristic spectrum. On the other hand, the copper greens show stronger variations, which is represented by 4 characteristic spectra evenly spaced. The characteristic reds are perhaps a little bit too similar visually but their positions generally follows the representative spectra globally distributed on a line on the XYZ space. Finally, the characteristic spectra of the glazes surround the representative spectra of the tiles in a triangle.

Of course, this visualization representation is not a perfect validation. It may suffer from metamerism. But this gives a good indication on the choice of spectra for the set generation before the PCCR.

Pearson correlation coefficient reduction. The reduction by PCC can then be made. For each tile, we calculate the Pearson correlation coefficient between the representative spectrum

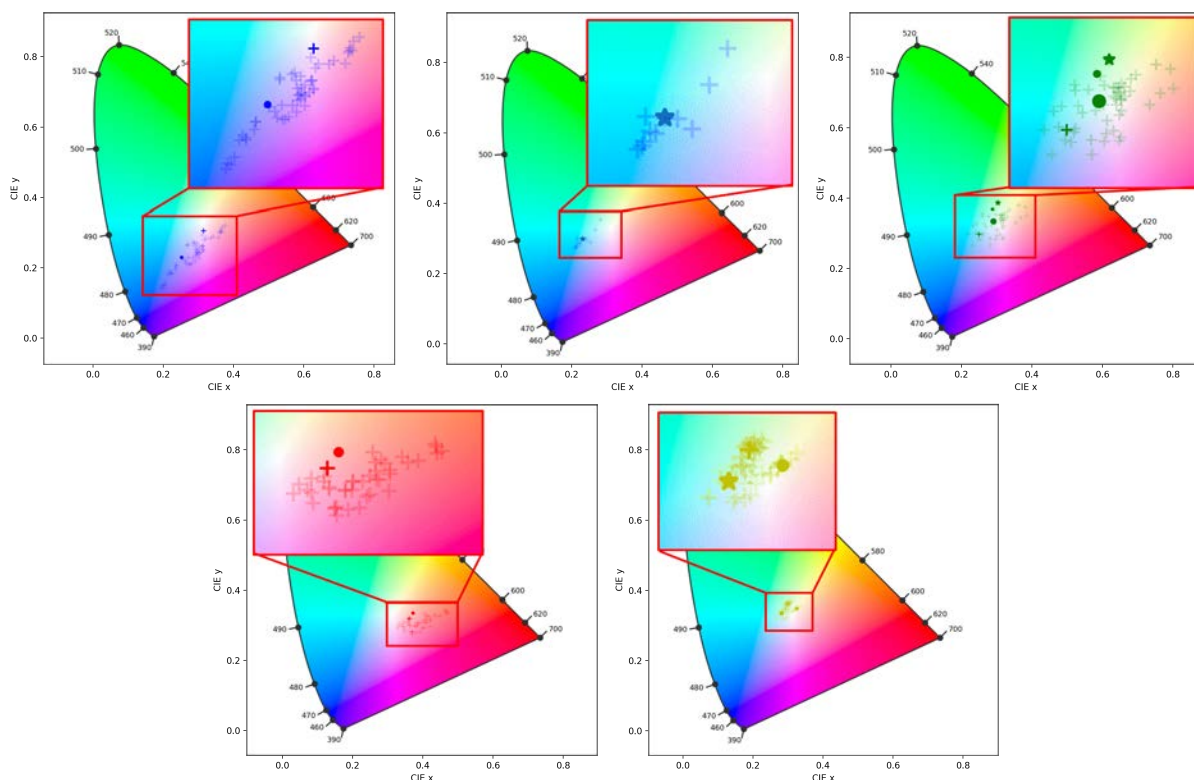


Figure 2.23: CIE 1931 chromaticity diagrams - CIE 1931 2° standard observer. On each graph, the crosses in transparency represent the chromaticity in XYZ of the representative spectra of each tile. The other patterns are the characteristic spectra corresponding to each pigment. In order: cobalt blue, copper blue, green, red, glaze.

of a pigment (accordingly a glaze) and its corresponding characteristic spectra. For example, we compute the PCC of the red pigment spectrum and the characteristic spectra Red 1 and Red 2, that of cobalt blue with Cobalt Blue 1 and 2, etc... If a tile does not have one of the pigments, then we allocate the value of 0. The information on a tile is thus reduced to 12 values, visible in Table 2.3

Visualizations. The goal here is not to carry out an identification as in the case of tapestry, but to study the degree of similarity of the tiles within the same pattern. For this, 3 processes are put forward: visualizations of PCA graphs, confusion matrix and dendrograms seem the most readable (see Figure 2.24).

As for the tapestry, a **PCA graph** can be made on the data. Since it is decorrelated. The two dimensions of the PCA are obtained after reduction of the 12 values and are displayed on the graph (see top left of Figure 2.24). So we have a point corresponding to a tile. For a same pattern, the more the points are clustered, the more the tiles are similar.

However, there may be a loss of information. In this case, a **confusion matrix** [117] may be preferred. For each tile, we calculate its similarity (with the Silhouette coefficient [184]) with the others, and we display the global table (see top right of Figure 2.24). Silhouette coefficient is a comparison metrics regularized between -1 and 1. Like Pearson correlation coefficient, the bigger is the coefficient, the closer are the two tiles. This visualization is interesting because it really allows to observe the similarity between two particular tiles.

However, it may seem very cumbersome because it contains too much information. In this case, it is possible to keep the similarity information of the Silhouette coefficient as metrics, and generate a **dendrogram** from the confusion matrix (see bottom left of Figure 2.24). This is then simply a tree easy to read. In addition, some other methods make it possible to obtain

Table 2.3: Table of the values of each tile of P1, C1 and C2 after PCCR.

Tile	Cob. Blue 1	Cob. Blue 2	Cop. Blue 1	Cop. Green 1	Cop. Green 2	Cop. Green 3	Cop. Green 4	Red 1	Red 2	Back 1	Back 2	Back 3
P1-27	0.981	0.899	0.000	0.738	0.953	0.588	0.949	0.000	0.000	-0.156	0.898	0.992
P1-29	0.951	0.717	0.000	0.629	0.989	0.444	0.865	0.000	0.000	-0.463	0.210	0.422
P1-31	0.966	0.748	0.000	0.798	0.957	0.672	0.976	0.000	0.000	-0.550	0.445	0.701
P1-25	0.986	0.823	0.000	0.780	0.922	0.646	0.956	0.000	0.000	-0.539	0.604	0.854
P1-23	0.957	0.728	0.000	0.751	0.937	0.561	0.897	0.000	0.000	-0.782	0.212	0.539
P1-6	1.000	0.845	0.000	0.785	0.898	0.635	0.930	0.000	0.000	0.022	0.955	0.979
P1-7	0.981	0.762	0.000	0.632	0.995	0.450	0.883	0.000	0.000	0.289	0.977	0.887
P1-8	0.958	0.758	0.000	0.662	1.000	0.509	0.919	0.000	0.000	0.715	0.842	0.575
C1-1	0.988	0.784	0.982	0.000	0.000	0.000	0.000	0.000	0.000	0.956	0.531	0.144
C1-2	0.997	0.849	0.795	0.000	0.000	0.000	0.000	0.000	0.000	-0.157	0.871	0.965
C1-3	0.995	0.824	0.995	0.000	0.000	0.000	0.000	0.000	0.000	0.884	0.614	0.277
C1-4	0.994	0.818	0.993	0.000	0.000	0.000	0.000	0.000	0.000	0.971	0.335	-0.055
C1-5	0.987	0.777	0.971	0.000	0.000	0.000	0.000	0.000	0.000	0.978	0.382	-0.029
C2-1	0.957	0.692	0.999	0.931	0.882	0.807	0.954	-0.965	-0.854	0.985	0.111	-0.307
C2-4	0.975	0.742	0.989	0.962	0.797	0.918	0.950	-0.957	-0.860	0.984	0.432	0.026
C2-5	0.989	0.766	0.991	0.830	0.957	0.709	0.983	-0.937	-0.807	0.981	0.436	0.026
C2-11	0.944	0.668	0.997	0.992	0.726	0.931	0.879	-0.968	-0.859	0.970	0.053	-0.364
C2-12	0.980	0.730	1.000	0.662	0.662	0.936	0.837	-0.973	-0.871	1.000	0.271	-0.144
C2-18	0.985	0.760	0.996	0.935	0.853	0.825	0.953	-0.952	-0.827	0.962	0.507	0.111
P2-88	0.911	0.916	0.000	-0.193	-0.121	-0.411	-0.395	0.932	0.995	0.246	0.999	0.914
P2-87	0.681	0.928	0.000	-0.643	-0.942	-0.580	-0.916	0.953	0.956	-0.327	0.809	0.981
P2-89	0.909	0.865	0.000	0.426	0.665	0.147	0.426	0.814	0.931	0.481	0.971	0.784
P2-95	0.896	0.886	0.000	-0.018	0.200	-0.260	-0.104	0.874	0.978	0.358	0.994	0.863
P2-94	0.850	0.935	0.000	0.079	-0.090	-0.099	-0.252	0.859	0.968	0.244	0.998	0.918
P2-93	0.690	0.949	0.000	0.345	-0.442	0.437	-0.124	0.950	0.967	-0.072	0.934	0.997
P2-81	0.600	0.893	0.000	-0.548	-0.928	-0.481	-0.870	0.921	0.951	-0.323	0.812	0.982
P2-82	0.864	0.955	0.000	-0.164	-0.184	-0.365	-0.423	0.938	0.995	0.267	0.999	0.907
P2-83	0.805	0.905	0.000	0.199	-0.027	0.029	-0.155	0.913	0.990	0.336	0.996	0.875
P2-84	0.842	0.917	0.000	0.187	0.332	-0.046	0.076	0.730	0.885	0.404	0.988	0.838
P2-85	0.884	0.916	0.000	0.452	0.146	0.289	0.078	0.873	0.981	0.416	0.985	0.829
P2-86	0.830	0.992	0.000	-0.166	0.011	-0.450	-0.287	0.936	0.969	0.062	0.974	0.973
P2-92	0.911	0.979	0.000	0.721	0.062	0.666	0.247	0.983	0.982	0.248	0.996	0.915
P2-91	0.850	0.968	0.000	0.295	0.050	0.124	-0.064	0.937	0.997	0.309	0.999	0.890
P2-90	0.840	0.880	0.000	0.326	0.296	0.113	0.112	0.800	0.946	0.376	0.993	0.855
P2-96	0.851	0.982	0.000	0.127	-0.090	-0.061	-0.235	0.955	0.996	0.044	0.970	0.981
P2-97	0.834	0.886	0.000	0.122	-0.004	-0.050	-0.174	0.749	0.896	0.331	0.997	0.879
P2-98	0.861	0.996	0.000	0.368	-0.348	0.328	-0.165	0.959	0.977	0.048	0.971	0.978
P2-78	0.867	0.940	0.000	0.302	0.107	0.115	-0.026	0.932	0.996	0.293	0.999	0.898
P2-72	0.880	0.883	0.000	-0.038	0.183	-0.290	-0.127	0.814	0.947	0.382	0.992	0.846
P2-66	0.888	0.921	0.000	-0.049	0.192	-0.287	-0.122	0.856	0.967	0.471	0.974	0.789
P2-67	0.908	0.970	0.000	0.143	-0.135	-0.029	-0.254	0.735	0.898	0.271	1.000	0.906
P2-73	0.832	0.873	0.000	0.035	-0.057	-0.155	-0.255	0.864	0.979	0.402	0.990	0.841
P2-79	0.875	0.895	0.000	0.124	-0.060	-0.052	-0.212	0.889	0.976	0.304	0.999	0.893
P2-80	0.841	0.998	0.000	0.973	0.568	0.898	0.754	0.933	0.989	0.056	0.973	0.977
P2-74	0.721	0.968	0.000	-0.646	-0.887	-0.636	-0.917	0.912	0.958	-0.093	0.929	0.997
P2-68	0.759	0.985	0.000	-0.460	-0.515	-0.617	-0.724	0.953	0.993	0.009	0.962	0.987
P2-65	0.855	0.980	0.000	-0.447	-0.396	-0.616	-0.653	0.894	0.987	0.234	0.999	0.922
P2-71	0.857	0.906	0.000	-0.178	-0.315	-0.350	-0.506	0.821	0.951	0.359	0.995	0.867
P2-77	0.858	0.979	0.000	-0.132	-0.477	-0.232	-0.563	0.941	0.999	0.163	0.993	0.949
P2-76	0.896	0.960	0.000	0.103	-0.391	0.040	-0.386	0.890	0.990	0.195	0.995	0.939
P2-70	0.823	0.987	0.000	0.926	0.430	0.831	0.612	0.959	0.988	0.143	0.987	0.953
P2-64	0.801	0.990	0.000	-0.011	-0.727	0.105	-0.508	0.964	0.988	0.070	0.974	0.975
P2-63	0.648	0.936	0.000	-0.652	-0.932	-0.601	-0.913	0.945	0.973	-0.157	0.901	0.999
P2-69	0.609	0.915	0.000	-0.428	-0.942	-0.270	-0.759	0.947	0.971	-0.144	0.906	1.000
P2-75	0.720	0.939	0.000	-0.042	-0.728	0.068	-0.466	0.953	0.960	-0.178	0.890	0.999

a dendrogram. We can start from a confusion matrix or apply a hierarchical agglomerative clustering algorithm [161] directly to the values from the PCCR (see bottom right of Figure 2.24). Their results are usually presented in this way too, so we can compare the different trees for a better understanding.

2.4.2 The question of tiles segmentation

Within the framework of a classical HSI study, we see that the classic pipeline consists in carrying out a segmentation after the dimensionality reduction. In our case, this seems interesting at first glance: we can classify the tiles into subgroups within the same patterns.

At first, we realize a simple segmentation mixing all the ceramic tiles of different patterns to see if our reduction already makes it possible to segment the 4 groups. We used K-Means method [140] with 4 clusters. As visible in Figure 2.25A, reduced values are enough to separate clearly the 4 pattern groups.

Then, we want to separate the 4 pattern groups into subgroups using K-Means for each of them. Here, the main constraint of K-Means is the number of clusters that must be defined upfront, which is not known at this step. Therefore we use the “elbow method” [200] in order to obtain the ideal number of clusters for each pattern. The method entails creating a plot that represents the explained variation [112] as a function of the number of clusters, and subsequently identifying the point on the curve where a distinct elbow or bend occurs. This elbow point is then selected as the optimal number of clusters to be employed in the analysis. This technique is now somehow criticized for its performance, but since our case is really simple with a low amount of data, it is appropriate.

We note in Figure 2.26 that if in some cases, the elbow is easy to find (C1, C2), we obtain rather an order of magnitude within 2 or 3 in others (P1, P2). By trial and error and manually, we refined our values in order to obtain 2 clusters for C1, 1 for C2, 4 for P1 and 7 for P2.

K-Means is then applied to each of the 4 patterns and the results of the groups are visible in Figure 2.25B. It is then difficult to be able to confirm or invalidate the subgroups obtained. The first idea would be to cross-reference the results obtained with other segmentation methods. We again use the agglomerative algorithm results used for the dendrogram visualization and we apply it by taking the same number of clusters for each pattern. The results obtained are visible in Figure 2.25C. Note that the subgroups are almost identical. Only 3 tiles of P2 do not show the same results. Again, it is hard to tell if our tiles are in the center of the subgroups or at the border with another.

To estimate the validity of our partitioning, it is then necessary to use evaluation formulas. There are two forms of them: internal and external evaluation. The difference is that the external assessment relies on test results to determine success, while the internal assessment will instead determine how clean each cluster is, compared to the others. Having no test sample, we can only use internal evaluation methods. We then first apply the different indices to our global results (i.e. one coefficient for the entire segmentation):

- Silhouette Coefficient [184]: 0.56
- Davies-Bouldin Index [53]: 0.63
- Calinski-Harabasz Index [39]: 37.9

Since it is not easy to interpret directly these indices, we may not that:

- Silhouette coefficients, as presented in section 2.4.1, are between -1 and 1 (-1 meaning that the clusters are bad, 1 that they are perfect, and 0 that they present overlapping).
- Davies-Bouldin is best when it tends to 0
- Calinski-Harabasz needs to be as large as possible.

Therefore the first 2 indices indicate that our results are rather satisfactory. Thereafter, we

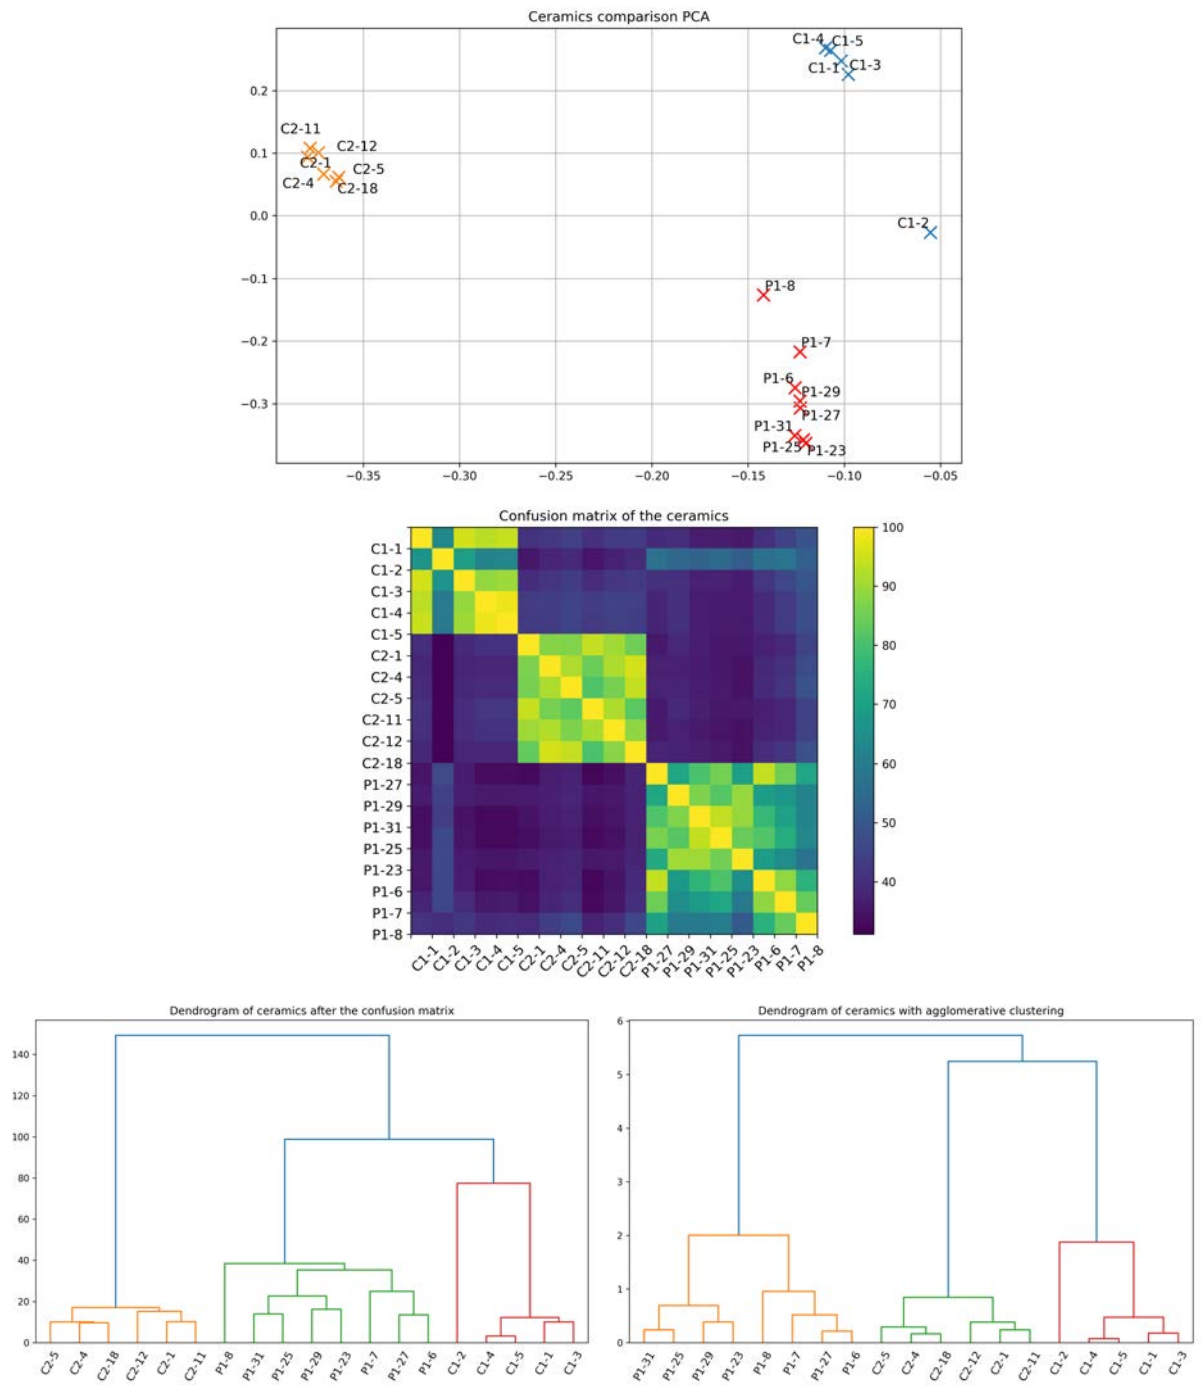


Figure 2.24: The different visualizations made. In the following, we have chosen to focus on the results obtained on the tiles from P1, C1 and C2 for a better understanding. The results obtained with P2 are visible in Annex III. **Top:** PCA graphs. **Middle:** Confusion matrix. **Bottom left:** Dendrogram from the confusion matrix. **Bottom right:** Dendrogram from agglomerative clustering.




Figure 2.25: Segmentations applied on tiles presented in Figure 5. **A:** K-Means applied to the total tiles see. The schema represents the tiles as they are represented in Context. Each color represents one pattern group. **B:** K-Means applied to each of the pattern in order to form subgroups. **C:** Agglomerative clustering applied to each of the pattern in order to form subgroups. The 3 tiles out of 55 presenting different results are squared in red. **D:** Silhouette coefficients used for every tiles in order to measure the segmentation reliability. A percentage of reliability is computed and presented on a colorbar.




Figure 2.26: *Graphs presenting the elbow method results for P1, P2, C1, C2.*

will use the Silhouette coefficient more precisely on individual tiles as it was used in section 2.4.1 for two reasons: by a simple re-scale, it can easily define a percentage of membership in a group, and it is easy to use in order to observe the probability of a particular tile to belong to one or more groups. The interest of the Silhouette coefficient is that it estimates the probability of being part of a subgroup for a tile by comparing its proximity to the other tiles in and out of the subgroups. We can then observe individually the probability of each tile to be part of a subgroup obtained by the K-Means method (see Figure 2.25D).

We rescale Silhouette coefficients from -1 to 1 as a percentage, and we see that obtained results are quite acceptable: a minimum coefficient of 55% on a tile, a maximum of 100%, and an average of 82%. Therefore, our method gives encouraging results, by showing that the segmentation is coherent, while indicating the tiles for which a redefinition by hand is not necessary excluded.

Moreover, the very binary aspect of the segmentation makes it impossible to know the similarities of different subgroups. The segmentation does not only really bring new information, but also hide the global similarity of tiles in pattern groups. Finally, it seems more useful to focus only on visualizations after PCCR method.

2.4.3 Similarity study results and discussions from visualizations

Finally, we concentrate on results presented in section 2.4.1. Indeed, it is more interesting to discuss about the global similarity within a pattern group than the possible existing subgroups. From a historical point of view, the similarity study of pigments allows to make much better assumptions about the common (or not) origins of tiles. So, the following discussion comes from the results presented in Figure 2.24 and in Annex III.

Below we will discuss about each pattern group one by one. Of course, for each case, conclusions are raised from the small samples acquired on site. That is why **results are presented as hypothesis** and assumptions are made. Indeed, these results have been discussed with specialists and are only the beginnings of leads or hypothesis that can be made by crossing with current historical knowledge. Note that the number of samples is taken into account: P2 group is not studied with the same view as C1.

C2: In the PCA graph of Figure 2.24, we can clearly see that C2 is the most clustered pattern group. The confusion matrix and the dendrograms confirm it: all the Silhouette coefficient distances between the tiles in the confusion matrix exceed 90% and the first separation in the two dendrograms are really low.

It is interesting to notice that C2 tiles are the most complicated as it is the only pattern comprising iron red, cobalt blue, copper blue and copper green pigments. And all pigment remains very similar.

Therefore we can make some assumptions about this important similarity. Indeed, it seems that the pigments are too similar to come from different manufactories. It is therefore probably

a pattern where the tiles come from **a common batch**.

P1: The analysis process of P1 tiles is pretty close to the C2 one. The PCA graphs in Figure 2.24 shows that the P1 tiles pigments are more scattered. They do not present different clusters as their scattering is globally homogeneous. The dendrograms and the confusion matrix confirm this impression, but underline a slightly bigger difference of P1-6, P1-7 and P1-8. The tiles coming from around the doors are finally a little bit more similar and clustered. However, this similarity is a way less important than the C2 similarity and it is hard to clearly separate panel and door P1 tiles into two groups.

Coming from this information, we can make the assumption that P1 tiles may clearly **not coming from a common batch**. The tiles might have totally different origins, or they simply come from different orders before their reuse in Saint-Maurice residence. In addition, the tiles of the panel and the door could perhaps have been recovered from different places because of the variety seen. Or they could have been placed this way in order to bring less different tiles together in the same place to create greater consistency. However, I would tend to think that the second hypothesis would be more plausible since the differences are not marked enough.

P2: P2 group is the most difficult group to study because of the number of tiles acquired and their variations. Figures in Annex III demonstrate that P2 tiles present a low global similarity. The tiles are indeed really scattered on the PCA graph: some tiles are even closer to P1 pigments than some of their own. This similarity with P1 can be confirmed by the dendrogram obtained with the agglomerative clustering (but not the one coming from the confusion matrix).

However, it stay difficult to highlight clear trends: the points are scattered but it remains homogeneous. Similarly, dendrograms do not generates branches that are very distinct from others. These observations corroborate with the visual impressions when looking at the tiles individually: a great diversity of color is visible between tiles.

It is then quite simple to make the hypothesis that P2 tiles have a very wide range of origins. They do not come from a common batch and were probably **obtained from various places and situations** when they were brought to Saint-Maurice residence. Some may even have been brought to the residence or the embassy at different times. Their positions on the wall were also probably thought to hide the colors variations because we notice that the close numbers rub shoulders from time to time on the PCA graph.

C1: The case of the C1 tiles is probably the most interesting. We can see at first sight that C1-1, C1-3, C1-4 and C1-5 are very clustered on the PCA graph on Figure 2.24 while C1-2 is quite far. This gap is clearly visible on the confusion matrix: all the tiles of C1 have a resemblance of more than 90%, except C1-2 which has less than 70% with all the others. Moreover, we see that this difference is the first encountered within the groups on the dendrograms. Even before the differences within P2 (see Annex III). We even see that in the PCA graph including the tiles of P2 (see Appendix III), C1 is in the middle of the P2 tiles.

This difference can be seen visually in Annex II. Looking closely, we see that C1-2 has a light blue center, a color that no other has. However, the results are obtained without taking this specificity into account. They mean that C1-2 is also different on the other pigments, and not only in the center.

These data, crossed with the strong similarity of the 4 others tiles, suggest that the C1 tiles would come from **a common batch** like C2, **but that the C1-2 tile would have been recovered later from a completely different place**. The reasons can be varied (breakage, loss during the move between the residence and the embassy, lack from the start to fill, etc).

The development of a robust dimensionality reduction technique was therefore necessary and

relevant in the hyperspectral study of Iznik ceramic tiles. The similarity study of their pigments made it possible to have a solid basis for the development of hypotheses on the reuse of tiles. This method also appeared to be more useful than Depth from Focus for our framework. But in a more global context, the combined use of the two can be perfectly adapted to an in-depth on-site and non-invasive study of ceramics. The complete observation of the thickness of the glaze and the pigments spectra also allows a good characterization of the ceramics of Saint-Maurice residence as being Iznik ceramics similar to those of Tunis observed in the lab [26].

In a more global framework of the reflection on the appearance of objects, this chapter have made it possible to take an interest in color, after having taken an interest in shape. For a more global vision of the appearance, it is finally necessary to take a closer look at the properties of light reflection described by the BRDF radiometric quantity.

Chapter 3

Capturing the in-situ BRDF

In the two previous chapters, we have focused on the acquisitions of shape and color (by its spectral definition) in the context of ceramic tiles. These elements happen to be the two traditional parts of a digital 3D object. However, as explained in the *Introduction*, a colored mesh is not enough to be accurate. This is why we aim to replace the color by a Bi-directional Reflectance Distribution Function (BRDF). Using BRDF allows, for example, to better represent the shiny aspect of gilding, reflections on polished stones, etc.

In the way, the restitution of the Saint-Maurice residence interiors presents the materials previously mentioned. This is why we want to acquire this data beyond colored mesh for a much more realistic visualization. In this chapter, we are therefore interested in the development of a portable, in-situ method for the acquisition and the measurement of the properties of BRDF on a monumental scale.

3.1 BRDF of materials

3.1.1 BRDF definition

As we mentioned above, color is not enough to represent accurately the appearance of a material. To understand this, we have to look a little closer at how color works from an optical point of view. In the following, we are only interested in opaque materials, thus we only focus on the reflection at a given point. What we see is actually the reflection of light on a shape through reflective properties. We should notice that light is not reflected uniformly, either spatially or in terms of wavelength. First, from a color point of view, a material reflects the light at different intensity depending on its wavelength λ . This curve of intensity as a function of wavelength is fixed (for most of the materials), and can be used as a material fingerprint. This is the reflectance spectrum, on which we were interested in the previous chapter.

Then, from a spatial point of view, the light is not reflected uniformly either. It is effectively a function of (globally) two parameters: the direction ω_i of the incident light, and the direction ω_o from the point of view of the observer (see Figure 3.1). For a given incident direction ω_i , the light is reflected in all directions. This makes it possible for us to see an illuminated object regardless of where the light is coming from (of course, out of shadow areas). However, the reflection intensity varies along the ω_o direction. Similarly, for a defined direction of reflection ω_o , different ω_i does not give the same intensity seen by the observer. If these spatial variations of intensity can be defined for a given object, they are also not the same from one material (which can be defined in our case as an element whose BRDF is unique and homogeneous) to another. Figure 3.2 shows three examples of reflection on the main types of different materials and illustrate this notion.

These variations of the intensity according to the directions (ω_i, ω_o) were then formalized in the form of a function by Nicodemus [159] in 1977. This is the Bi-directional Reflectance




Figure 3.1: *Graphic showing vectors used to define the BRDF. All vectors are unit length. ω_i points toward the light source. ω_o points toward the viewer (observer). \mathbf{n} is the surface normal. Image and caption taken by Michael Holroyd (2007 for Wikipedia).*

Distribution Function (BRDF), which is defined by:

$$\rho(\omega_i, \omega_o) = \frac{dL_o(\omega_o)}{dE_i(\omega_i)} = \frac{dL_o(\omega_o)}{L_i(\omega_i) \cos \theta_i d\omega_i} \quad (3.1)$$

where

- L is the radiance ($Wsr^{-1}m^{-2}$),
- E is irradiance (Wm^{-2}),
- θ_i is the angle between ω_i and the surface normal, \mathbf{n} ,
- i indicates incident light direction,
- o indicates reflected light direction.

In radiometry, radiance L is the radiant flux reflected (in our case but it can be emitted, transmitted, or received) by a given surface, per unit solid angle per unit projected area. Irradiance, as for it, is the radiant flux received by a surface per unit area.

The BRDF is therefore a function which makes it possible to obtain very easily what is the flux received by the observer for a given position with respect to a scene comprising an object and one or more light sources. Moreover, the flux obtained depends on the wavelength (or simply depending on RGB channels), therefore a color. Thus this function ensures a much better visual coherence when attached to a 3D object.

In addition, the BRDF need to fulfill some properties that ensure a realistic physical consistency:

- Positivity: $\rho(\omega_i, \omega_o) \geq 0$,
- Obeying Helmholtz reciprocity [212]: $\rho(\omega_i, \omega_o) = \rho(\omega_o, \omega_i)$
- Conserving energy: $\forall \omega_i, \int_{\Omega} \rho(\omega_i, \omega_o) \cos \theta_o d\omega_o \leq 1$

Positivity prevents obtaining a negative flux. Helmholtz reciprocity demonstrates that the directions ω_i and ω_o can be interchanged. Finally, conservation of energy assures that there is no creation of energy. The sign \leq instead of $=$ is used to take into account the energy absorbed by the material.




Figure 3.2: The figures upper show how light is reflected off the surface of an opaque object whereas rendered images show the appearance of a given BRDF on an object (here a 3D model of the Willendorf Venus). For incident an light direction ω_i (the black arrow), light is reflected in different light directions ω_o with varying intensity. The magnitude of the intensity is symbolized by the length of the gray arrows. Each material has its own optical properties. Some materials are said to be lambertian (**left**), when they reflect light equally in all directions. A material that reflects light in many directions is said to be diffuse. This is particularly the case with terracotta, matte paints, etc. Other materials have a more glossy appearance. They have a lobe around the main direction of reflection, also called specular direction (**center**). Visually, they will appear a bit more lighter in that direction. This is the case, to varying degrees, of plastics, but also of metals or polished stones. In this case, the finer the lobe, the shinier the material. Finally, some materials only reflect light in the specular direction (**right**). These are mirror materials. This last case corresponds to the application of Snell's laws.

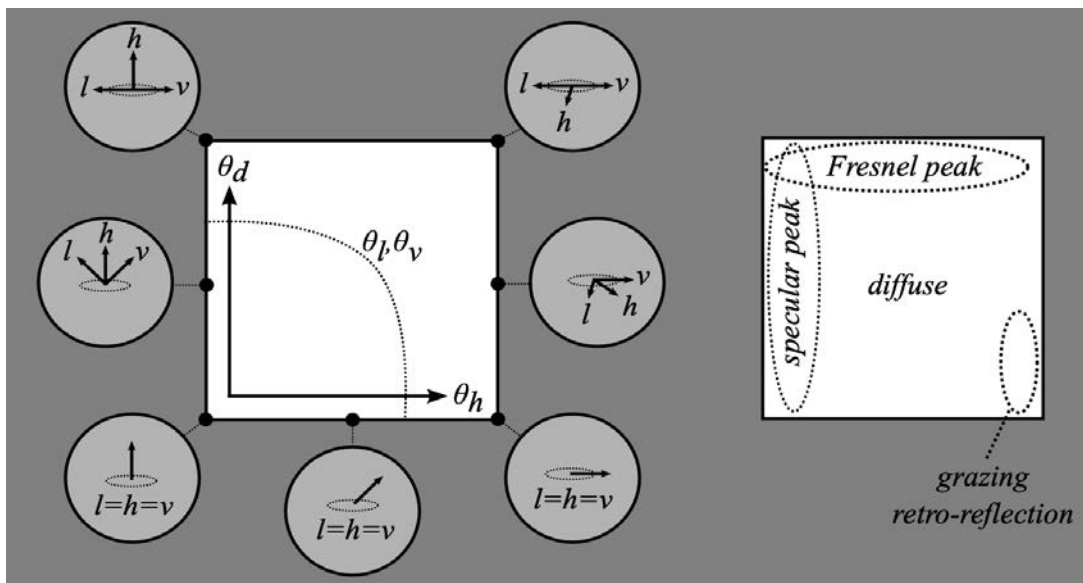


Figure 3.3: **Left:** The “image slice” view: BRDF is represented on a 2D scattering point plot along the axes (θ_h, θ_d) . Here, $\mathbf{l} = \omega_i$ and $\mathbf{v} = \omega_o$. **Right:** The slice can then be studied as a simple map of the BRDF, where each zone corresponds to a particular optical effect. Image taken from [37].

The BRDF as defined by Nicodemus [159] is therefore a real function which can be as complex as wanted. That take into account (ω_i, ω_o) which are vectors and which one will rather write in practice $(\theta_i, \phi_i, \theta_o, \phi_o)$, with θ_i the angle between ω_i and the surface normal, \mathbf{n} , ϕ_i the angle between ω_i and the surface tangent, \mathbf{t} (and analogously for θ_o and ϕ_o , see Figure 3.4 left).




Figure 3.4: Instead of treating the BRDF as a function of (θ_i, ϕ_i) and (θ_o, ϕ_o) , as shown on the **left**, Rusinkiewicz et al. [186] consider it to be a function of the halfangle (θ_h, ϕ_h) and a difference angle (θ_d, ϕ_d) , as shown on the **right**. The vectors marked \mathbf{n} and \mathbf{t} are the surface normal and tangent, respectively. Image and caption taken from [186].

Rusinkiewicz [186] then proposes a parametrization according to (\mathbf{h}, \mathbf{d}) (see Figure 3.4 right). This parametrization allows a more interesting representation of the BRDF for several points.

Already, for the sake of simplification. In the quasi-general framework of an isotropic BRDF (i.e., it is independent of ϕ_o and ϕ_i for a given $\Delta\phi = \phi_o - \phi_i$), one must generally observe the following BRDF depending on $(\theta_o, \theta_i, \Delta\phi)$. However, Romeiro et al. [183] observe MERL BRDFs [146] and conclude that the isotropic BRDF can be observed following (θ_h, θ_d) only with this new representation.

Moreover, the representation according to (θ_h, θ_d) is much more relevant. As presented by Figure 3.3, we can represent the BRDF values on a 2D plot according to (θ_h, θ_d) (an “image slice”), and we obtain a representation for which each point zone corresponds to a different aspect of BRDF: specular, diffuse, etc.

Finally, the definition of the halfangle vector \mathbf{h} takes into account BRDF models such as in the theory of microfacets explained in the section below.

3.1.2 Models for simplification and utilization purposes

In order to be able to represent the BRDF as faithfully as possible, models have been effectively developed and implemented. Of course, a BRDF can remain in the form of data, but the amount of information to represent it can become overweening.

In the following, we present a type of model widely used in Computer Graphics: the **micro-facet models**.

At first, the BRDF is oftenly separated into two distinct parts to model: the diffuse part ρ_d and the specular part ρ_s . This is akin to representing each BRDF as the sum of a diffuse base similar to Figure 3.2 left, and a specular lobe similar to Figure 3.2 center:

$$\rho(\omega_i, \omega_o) = \rho_d(\omega_i, \omega_o) + \rho_s(\omega_i, \omega_o) \quad (3.2)$$

The diffuse part is based on **Lambert's** illumination model [122]. It represents a perfect diffuser which emits a constant luminance in all directions. It is then defined by its reflectance \mathcal{R}_d :

$$\mathcal{R}_d(\lambda, \omega_i, \omega_o) = \frac{\mathcal{R}_d(\lambda)}{\pi}, \quad (3.3)$$

which can be rewritten in the case of the BRDF by

$$\rho_d(\omega_i, \omega_o) = \frac{k_d}{\pi}, \quad (3.4)$$

where k_d is an albedo of constant value (or a constant triplet of value RGB) defined between 0 and 1.

The case of the specular lobe is much more complex because it presents BRDF shapes which can vary enormously according to (ω_i, ω_o) . This is why many models exist.

Torrance-Sparrow

Torrance-Sparrow Model [207] was developed to represent metallic surfaces. Based on measures and physical comparison, it relies on geometric optics to describe the specular reflection on a rough surface composed of “**microfacets**”, that is a distribution of microscopic surfaces with different normals (see Figure 3.5). These surfaces exhibit a softer specular lobe unlike smooth surfaces which do not show a well-marked specular peak.




Figure 3.5: Diagram showing a microfacets surface. Image taken from [213].

In this way, it defines the specular reflectance \mathcal{R}_s as:

$$\mathcal{R}_s(\lambda, \omega_i, \omega_o) = \frac{F(\theta_i, \eta(\lambda))}{4} \frac{D(\alpha)}{\cos \theta_i} \frac{G(\omega_i, \omega_o, \mathbf{h})}{\cos \theta_o} \quad (3.5)$$

where

- $F(\theta_i, \eta(\lambda))$ is the Fresnel function and represents the fraction of energy reflected from the surface. It can be written as

$$F(\theta_i, n(\lambda)) = \frac{1}{2} \left(\frac{\sin^2(\theta_i - \tau)}{\sin^2(\theta_i + \tau)} + \frac{\tan^2(\theta_i - \tau)}{\tan^2(\theta_i + \tau)} \right) \quad (3.6)$$

with $\tau = \frac{\sin \theta_i}{\eta(\lambda)}$ and η which is related physically to the material as it is its refractive index.

- $D(\alpha)$ is the facet distribution of the surface with roughness α . It is defined by Torrance and Sparrow as a Gaussian distribution with α as the standard deviation :
- $G(\omega_i, \omega_o, \mathbf{h})$ is finally the geometric attenuation factor translating the phenomena of shading and masking of one facet by another:

$$G(\omega_i, \omega_o, \mathbf{h}) = \min \left(1, 2 \frac{\cos \theta_h \cos \theta_o}{\cos \theta_h}, 2 \frac{\cos \theta_h \cos \theta_i}{\cos \theta_d} \right) \quad (3.7)$$

This model was then adapted to computer science as the **Cook-Torrance** Model [45]. It is also based on a microfacets model as before, which it adapts to a BRDF:

$$\rho_s(\omega_i, \omega_o) = \frac{F(\omega_i, \mathbf{h})G(\omega_i, \omega_o, \mathbf{h})D(\mathbf{h}, \alpha)}{4|\cos\theta_i||\cos\theta_o|} \quad (3.8)$$

where F , G and D still refer to the Fresnel function, the microfacets distribution and the geometric attenuation accordingly. Improvements are made to each of the terms. Already, the model improves the rendering of metallic surfaces by showing that the Fresnel coefficient depends on the wavelength. It is rewritten in a form which is an exact formulation for dielectrics with unpolarized light:

$$F(\omega_i, \mathbf{h}) = \frac{1}{2} \frac{(g-c)^2}{(g+c)^2} \cdot \frac{1 + (c \cdot (c+g) - 1)^2}{(c \cdot (c-g) - 1)^2} \quad \text{with} \quad \begin{cases} c = \omega_i \cdot \mathbf{h} \\ g = \sqrt{\eta^2 + c^2 - 1} \end{cases} \quad (3.9)$$

Then the main contribution is made on the microfacets distribution.

Beckmann

But the Gaussian facet distribution function does not satisfy the energy conservation condition. Then this model improves it based by Beckmann *et al.* [24] and we have:

$$D(\mathbf{h}, \alpha) = \frac{\exp^{-\tan^2 \theta_h / \alpha^2}}{\pi \alpha^2 \cos^4 \theta_h} \quad (3.10)$$

The geometric attenuation factor is also modified but this does not mean that it is perfect: it is calculated thanks to rough approximations on the shape of the surface and does not correspond very well to reality because its derivative is not continuous. Moreover, it is not invariant under rotation around the normal and it is totally independent of the surface roughness.

GGX

Walter *et al.* [213] propose then to improve this model. It starts from the same F as Cook-Torrance, and improves the distribution of microfacets with:

$$D(\mathbf{h}, \alpha) = \frac{\alpha^2}{\pi \cdot (\cos^2 \theta_h \cdot (\alpha^2 - 1) + 1)^2} \quad (3.11)$$

This new distribution function is called GGX. Moreover, as the function of geometric attenuation is related to the geometry of the microfacets, and because of the problems of continuity cited above, Walter *et al.* [213] modify it by taking into account the parameter α in this one:

$$G(\omega_o, \omega_i, \alpha) = \frac{4}{\left(1 + \sqrt{1 + \alpha^2 \cdot \frac{1 - \cos^2 \theta_i^2}{\cos^2 \theta_i^2}}\right) \cdot \left(1 + \sqrt{1 + \alpha^2 \cdot \frac{1 - \cos^2 \theta_o^2}{\cos^2 \theta_o^2}}\right)} \quad (3.12)$$

This method then has many advantages over Cook-Torrance. First, the results show that GGX better represents the BRDFs on the samples measured by Walter *et al.* [213] than Beckmann. Then, it does not present the continuity defects to the derivatives of Cook-Torrance, which is more interesting within a framework of minimization problems. Besides, it simplifies the BRDF models by only having the α parameter to determine for D and G (see Figure 3.6). However, it is important to remember here that our goal is ultimately to focus on the acquisition of BRDF to best model materials. Finding a BRDF model that remains simple to fit while being realistic and physically plausible remains our priority. In this case, GGX appears as a great solution.




Figure 3.6: Four rendered images of a 3D object of Willendorf Venus. All the models are attached to a GGX BRDF model with various values of α and a fixed $n = 3.0$.



Figure 3.7: Two rendered images of a 3D object of Willendorf Venus. **Left:** Rendering made with a model associated with a BRDF model (GTR Disney [37]). **Right:** Same rendering but with a SVBRDF. Results shows that a SVBRDF is necessary for a faithful representation.

Newer and more intuitive models have of course been developed since, such as the GTR Disney model [37]. It remains based on a microfacets model, but seeks to develop more intuitive parameters like *metallic* or *sheen*. It is without any doubt more malleable, but it contains more parameters, and is therefore more complicated to fit. That is why, in the following, we will focus on the GGX model.

3.1.3 Adding spatial information

Simplifying the BRDF as a model makes it possible to acquire or represent objects digitally. However, a simple BRDF is not enough. The majority of objects around us do not have a perfectly uniform appearance. This is why a Spatially Varying BRDF or SVBRDF $\rho(\mathbf{x}, \omega_i, \omega_o)$ is required. As explained by Guarnera et al. [90], the problem of the acquisition of the SVBRDF will then be the fact of having a BRDF per vertex (or per face). Capturing an SVBRDF generally requires measurements and time-consuming processing with the help of extensive specialized laboratory equipment to capture reflectance data which may still be incomplete. This is why SVBRDF simplifications are also implemented. The most common method consists in assuming that our material is not composed of an SVBRDF which varies in each vertex, but of a list of BRDFs which compose it. For example, Alldrin *et al.* [13] assumes that an object can be represented as a linear combination of BRDFs:

$$\mathcal{H} = \mathcal{W}\mathcal{B}^\top \quad (3.13)$$

where $\mathcal{H} \in \mathcal{R}^{n \times d}$ is a discretization of the BRDF at each of n surface points, $\mathcal{B} \in \mathcal{R}^{d \times k}$ is a discretization of k basis BRDFs, and $\mathcal{W} \in \mathcal{R}^{n \times k}$ is a matrix of weights for the contribution of each basis BRDF at each point.

The problem of this method is that it is difficult to define the size of the k bases of BRDF (which does not necessarily represent real BRDFs associated to real materials) if we want to acquire an object. In a pipeline where the object contains areas with well-defined materials, it may then be interesting to segment the materials into several clusters. Each material can then be assigned a BRDF. The acquisition can in this context be carried out after the segmentation, or the segmentation can be done on the measurement of BRDFs.

Note that in the following, we will mostly speak about BRDF acquisition instead of SVBRDF acquisition: a SVBRDF measurement is finally a localized BRDF measurement.

3.2 State of the art of BRDF acquisition methods

After taking a look at the representation of the appearance thanks to the BRDF quantity, we are interested in how to acquire it. As we said before, the BRDF is a complex function and its capture is not as simple as a color acquisition (from RGB to hyperspectral). Indeed, we have seen that it is a function of ω_i , and ω_o , that is the directions of the light and the observer with respect to the surface. For its measure, it is therefore necessary to obtain a large number of measures data from a maximum of known directions of the hemisphere for ω_i and for ω_o . It means that we usually have to acquire with a detector some intensity values from various light and sensor positions around the captured sample. When acquiring the SVBRDF of an object instead of its BRDF, the computational complexity significantly escalates with the increasing number of dimensions.

By using the properties of reciprocity and assumptions (for example on the isotropy), one can further simplify the measurement, but complex set-ups remain necessary.

We should remember that in our case, the goal is to complete the acquisition of the interiors of Saint-Maurice residence, in particular sections of wall or ceilings. Therefore we have several constraints. First, it is an acquisition of an SVBRDF (the interiors present many different elements such as polished stones, paintings on wood, gilding, etc.) in-situ, in an uncontrolled environment. Consequently, this requires easily transportable equipment. In addition, the measurement is made on a time limited campaign. For that reason, we need a measurement that can be relatively fast, especially given the size of the object to be measured.

3.2.1 Acquisitions made in a controlled environment

For precision measurement, the rigor required on the positions of light, camera and object, as well as the necessary calibrations, require a very controlled environment. It is for this reasons that set-ups are mostly being developed in laboratory. Their goal is generally to obtain dense, precise, even metrological data. In the laboratory environment, the goal is of course to have a data driven approach over a model driven one.

Acquisition set-ups are numerous and Guarnera *et al.* [90] gives a broad state of the art. According to them, we can globally separate into 5 groups the existing in-lab set-ups: goniore-flectometers, image-based setups, catadioptric measurement setups, (hemi)spherical gantry and linear light sources set-ups. We are going to review this list briefly since laboratory methods by definition do not apply to our case.

The oldest set-up types, such as the ones described by Nicodemus [159], are also the ones used by Torrance and Sparrow to compare their model to data [207]. They are based on **gonioreflectometers** and are more described by Hsia *et al.* [102]. It is simply collimated light (e.g., a laser) that illuminates the studied sample on a 3D rotating sample holder. The light is then reflected and an acquisition device (a spectroradiometer for example) capture it (see Figure 3.8). This detector is usually mobile, so the light and the observer direction can be multiplied easily. Reliable, these set-ups can be precise enough to be used for reference measurement from a metrological point of view [162]. Metrological set-ups require having a sample with a known shape so as not to encounter any normal problems.

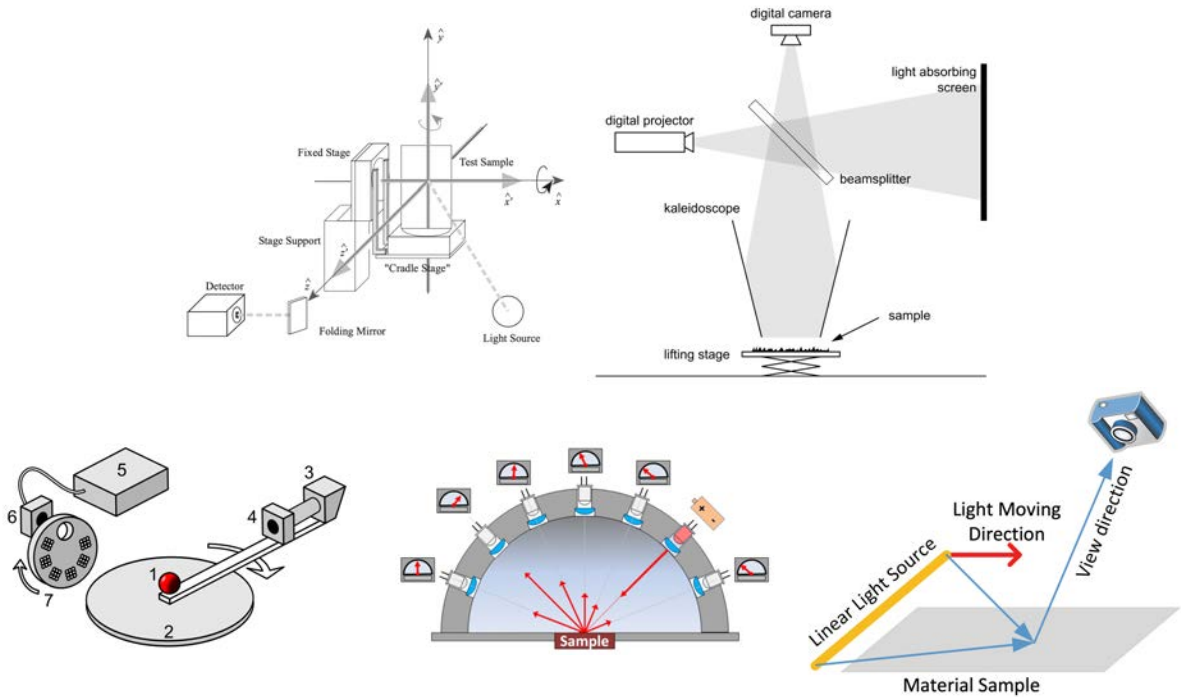


Figure 3.8: Diagrams of different techniques of in-lab BRDF acquisition . **Top left:** A gonioreflectometer. Image taken from [70]. **Top right:** The catadioptric set-up provided by Han *et al.* using a kaleidoscope. Image taken from [97]. **Bottom left:** An image-based acquisition set-up using a sphere. “A sample sphere (1) is mounted on a turntable (2), to which a digital monochrome still camera (3) is attached. The camera is equipped with a visible-spectrum tunable filter (4). The sphere is illuminated by a light guide coupled xenon light source (5) with another tunable filter (6) mounted in front; near-UV light is generated with LED that can be selected using a motorized wheel (7). On the exit aperture of (6) and the entry aperture of (4), we attach optical depolarizers” Image and caption taken from [105]. **Bottom center:** A hemispherical light dome set-up. Image taken from [28]. **Bottom right:** A linear light source set-up. Image taken from [42].

Coming from this limitation, **image-based BRDF measurement** techniques are varied and make it possible to acquire shaped samples. Marschner *et al.* [144] use it to develop a simple set-up. They basically replace the sensor of a gonioreflectometers with a camera. Matusik *et al.* [146] carry out a similar set-up, on which an image of the sphere is acquired (like in Figure 3.8 bottom left). The resulting image then gives a large amount of angle, defined by the underlying normal on the sphere for each pixel of imaged sphere. So measurements are much denser. It is interesting to note that a drawback of camera-based systems is that they are more limited in term of spectral acquisition. They indeed require fast spectral imaging systems such as tunable

filters as we saw in previous chapter, and the acquisition can be very long if we want more wavelengths. Nevertheless, these types of set-ups make it possible to go further by adapting the equipment, to capture more complex BRDF, for example associated with fluorescence [105], or polarization [18].

In addition, these BRDF acquisition systems can be enhanced with optical devices using both lenses and mirrors, so-called **catadioptric** set-ups. The objective is to play on mirrors and/or lenses to obtain more efficient capture systems. The main interest is then to limit the mobility of the parts as much as possible. For example, Gosh *et al.* [78] developed a set-up with a mirrored dome and parabola. Thanks to this, a BRDF from a simple sample can be captured from a single camera position. In the same way, kaleidoscopes make it possible to multiply the positions of view and camera. This allows to reconstruct simple 3D objects [180], but also the appearance through the BRDF from sample with a geometry [97]. The saving time is then very important. However these devices are very sensitive, it is necessary to make a very precise calibration if we want accurate results. Moreover, the set-ups presented above are all limited to uniform samples and does not fit for complex objects with SVBRDF.

In order to acquire whole objects, the most common method is the creation of a light dome, in which the object is photographed. These are the so-called **(hemi)spherical gantry**. A single camera can be fixed at the top of the light hemisphere for flat samples [142] (see Figure 3.8 bottom center) or at the extremity of a robot arm [137] for complex ones. Other systems are built replacing the lights in the hemisphere with cameras with flash [185]. Finally, Tunwattanapong *et al.* [208] simulates a complete sphere with a curved arm of LEDs. All these methods allow, after a calibration, to multiply the points of view and lights to obtain dense SVBRDFs. However, they are generally very large and very expensive.

Gardner *et al.* [75] present then a set-up based on a **linear light source**. A neon tube is translated horizontally over the surface of flat sample and the camera is moved simultaneously (see Figure 3.8). Then, the BRDF is obtained with a reflectance model used to fit the sparse measured data. Chen *et al.* [42] even enhance the technique by modulating the illumination of the linear source. This technique is then much faster than the hemispherical methods, but the output data is too sparse to be studied without model fitting.

Finally we note that if the set-ups developed in laboratory are interesting, they do not fit our requirements. First, gonioreflectometer-based methods presented are limited to homogeneous BRDF samples. Second, the main problem is that most of these devices cannot be transported and deployed in-situ. In the case of hemispherical gantries, the large number of lights and cameras makes it inoperable on a large object (e.g., a full ceiling in our case) on site. In addition, the enormous volume of data obtained by a dome makes the on-site acquisition time unmanageable. Similarly, if we wish to simplify our set-up with catadioptrics, the methods with kaleidoscopes require an extremely rigorous calibration, which does not fit with the idea of a measurement in an uncontrolled environment. The set-ups presented above and their limits point us to the transition to specific techniques for an uncontrolled environment. However, it should be noted that they can be used to calibrate these specific techniques.

3.2.2 Acquisitions under constraints

The in-situ constraints seem to show that a model driven approach may be better suited instead of data, with adapted set-ups. The state-of-the-art indeed underlines that under uncontrolled environment method show rough and imprecise measurements. To clarify, this inaccuracy can be applied to a lot of things: it can target both the intensity value obtained, but also the directions of the observer and the light(s), their calibration (e.g., the light intensity, its angular distribution,

the intrinsic camera parameters), the light environment, or even the angular representativeness of the input parameters (ω_i, ω_o) in the hemisphere. It is this imprecision on the measurement that limits the desire to take a data driven approach. That is why under constraints set-ups are radically different and try to represent as faithfully as possible the captured object rather than making a metrologic measurement.

In this way, the first applications of on-site acquisition of the appearance of materials occur with reflectance transformation imaging (RTI, see Figure 3.9) [154, 41]. This method is not strictly speaking a BRDF acquisition, but rather a **capture of reflectance information**. It works by illuminating with a spot light a studied object or scene from a new direction each time a photo is taken, taken from a fixed position. A mirror sphere placed in the field of the camera is then used in order to get the lighting environment (i.e., a direction of the light at infinity to simplify). Visualization can then be done in several ways. First, the object can be rendered from the point of view of the camera, by moving the light source in the scene. In this case, the appearance is computed by interpolating the images taken under different light directions. Second, we can reconstruct a 3D profile from the photographs with a shape-from-shading technique assuming the object is lambertian. The 3D profile is visualized with a shiny material to underline the details. As can be seen, this is a portable visualization technique and not an actual acquisition of the appearance.

In the same spirit, Li *et al.* [129] captures the 3D and global reflectance of a scene from a single point of view. Using the stereo view of two 360° cameras, and a neural network, they reconstruct on one side an illumination map, and on the other side the normal and reflectance maps of the scene. However, in this case too, the reproduction of the scene is more or less realistic scene, without the notion of physical acquisition of the appearance (i.e., a BRDF).



Figure 3.9: Pictures illustrating the RTI method. **Left:** RTI shooting on Christian monograms. **Top right:** Photograph of a hyposthènon of the theater of Miletus: traces of a graffiti. **Top left:** Snapshot of the hyposthènon using the RTI software for specular enhancement in order to read a graffiti. Images taken from [41].

We can then observe methods that will try to measure a BRDF more as simplified models. These cases seek to be simple and quick to set up, with little material to be portable. For example, some measurement techniques **from one HDR photo** (or few ones) seek to obtain the BRDF as [133] (see Figure 3.10). But this case does not fit for an object: the captured BRDF is not spatially varying. In addition, it presents a visible problem in many cases: the

techniques presented on the acquisition with a single photo rely on strong hypothesis. There is no notion of acquired object here, but a BRDF which can only be associated with normal maps for the simulation of mesostructures [133], or acquisitions on perfectly flat objects with no question about the normal directions [179]. Finally, it is necessary to start from very strong a priori on the base material [214, 9, 104], or on incident light [181].

Zhou *et al.* [225] makes this scenario more complex and creates a pipeline for acquiring an SVBRDF of a known 3D object with few images and uncontrolled light. This method still needs a strong priori; it requires a generic basic of BRDF. The 3D of the object is never modified and has to be precise. So it demands a reliable 3D acquisition before.




Figure 3.10: *Global pipeline of the one photograph BRDF acquisition of Lin et al. [133]. An HDR photo is taken (a) with a light probe (b) for the lighting environment. Then, an estimated BRDF (c) is extracted, and a mesostructure (d) is determined in order to create a model including a BRDF with a mesostructure for texturing (e). Image taken from [133].*

Other methods then also start with a small number of photos as input without calibrated and localized light, but seek to free themselves from a priori thanks to **Deep Learning**. We can see various situations. For example, we observe some previous simple cases: a single photograph and a surface plane [130, 219, 131], or a 3D object [132]. Some other cases also use more images, with different positions of light like [61]. The main interest we can see in comparison to the previous techniques is that the results obtained required less hypothesis at first sight. However, as already explained in Chapter 2, deep learning methods are also limited by the original hypothesis we made with the training data. It requires a huge amount of images from scenes which are connected to known BRDF. So it is also dependent from our train data that can be real or synthetic data. In our case, it is impossible to generate this type of data if it is coming from real images, so we may not consider it for our acquisition set-up. Moreover, in the case of synthetic images (which is the case of the works presented), the a priori becomes very important and depends on the BRDF model chosen for training. Learning is then synthetic and we move away completely from the case of real acquisition. Finally, Deep Learning as well as one-image based techniques cannot still really be called measurement methods. They are closer to visualization techniques and it is not a way we want to follow. It is also interesting to note that the deep learning methods are greatly influenced by their input data during learning, much more than traditional pipelines.

The logical next step is therefore to **increase greatly the number of photos** taken, to over 50 or even over 100. Nam, Lee *et al.* [155] present a simple portable set-up based on a camera attached with an unstructured flashlight. The acquisition then takes place by taking the object studied in photo a large number of times with the calibrated flash by varying the position of the camera, as if performing a photogrammetry (see Chapter 1). A coarse 3D model is created by SfM-MVS approach. Then, an inverse rendering algorithm is applied as a two-part iterative process. The first part is the estimation of the SVBRDF based on Alldrin *et al.* [13] whereas the second one is the normal optimization for 3D enhancement (see Figure 3.11). This method therefore makes it possible to get out of many of the hypotheses necessary before, and to work on more complex SVBRDFs associated with objects with realistic shapes. We notice

that the 3D results are very good and that the appearance obtained is realistic.




Figure 3.11: Overview of algorithm used by Nam *et al.* [155] and Hwang *et al.* [107]. They iterate estimations of reflectance F (defined as basis functions F_b with corresponding weight maps W), shading normals N , and 3D geometry X . Image and caption taken from [155].

The iterative inverse rendering is by the way quite similar to **differentiable rendering** techniques [60]. It consists mainly in simulating the light transport of a scene in order to make a complex minimization of inverse rendering. This makes it possible to retrieve BRDF properties, object shapes, etc. If there does not necessarily currently exist a total acquisition pipeline using it, it seems to be very promising. Indeed, its evolution presented in [60] tends to think that incorporating or drawing inspiration from the proposed solutions is likewise an interesting track.

Hwang *et al.* [107] propose a similar technique to Nam, Lee *et al.* and enhance it by incorporating a polarimetric component (see Figure 3.12). A first part based on the Baek *et al.* [18] explain that the BRDF can be expressed and defined by a polarimetric reflectance model. So a set-up based on a polarization machine vision camera and a unstructured flashlight surrounded by a linear polarizer is built. As before, the acquisition works on taking many images with different positions. Again, the 3D model is first found by Sfm-MVS. Then, a polarimetric BRDF is found based on a Mueller matrix [30] in order to finally get an estimation of the SVBRDF.




Figure 3.12: Presentation of the set-up and the results obtained by Hwang *et al.* [107]. From a simple portable set-up, the 3D geometry obtained is precise and the appearance is realistic. Image taken from [107].

These 2 methods and the differentiable rendering generally present the same disadvantages in our case, but also the same opportunities and ideas to exploit. Overall, we notice that in both cases, we are limited in terms of angular densities to measure a BRDF by the fixed flashlight-camera device. Precisely, the measurements are done with a calibrated angle of $\theta_a = 5^\circ$. It actually limits the measurements on many angles, which can limit the BRDF and does not fit with what we want to do. This is also why in some areas, the specular reflections of the varnish are less shiny than on the initial image. In addition, the SVBRDF is based on the weight method of Alldrin *et al.* even so, in our case, we rather want to perform a BRDF measurement associated with a real material, and then segment our object into several materials. This approach seems

more physically correct to us, and it is therefore a limit to the methods presented.

However, we can also start from several ideas presented here. Already, we see that some approaches give very good results: the multi-view for the 3D reconstruction by an SfM-SVM approach, followed by an iterative reconstruction of the BRDF and the normals.

All the presented methods demonstrate several things. First, in-situ and sparse measured data are much less reliable than in the laboratory, and requires a model fitting. Then, we remark that the constraint of on site portability imply to have a much smaller number of measurements, especially in the case of SVBRDF. This could be much more accentuated in our case, because we measure at the architectural scale and no longer at the scale of an object held in the hand. This is not necessarily a problem, as Dong *et al.* [63] shows that we can measure very densely small areas and then in a simpler way large spaces. So it would fit for a monumental measurement at the data scale of a portable acquisition. However, it seems necessary to set up a pipeline where the directions of views and lights are uncorrelated and known on each image, in order to reconstruct a more realistic BRDF. For this, a set-up with a light and a separate camera seems more appropriate, with the need for a method to position them with respect to the object studied.

3.3 In-situ and highly portable acquisition method

The symbols used in the following are summarized in Annex V.

We have therefore seen that the current acquisition methods do not really suit our needs. Indeed, the set-ups in the laboratory are not easily transportable for on-site application. Similarly, many portable techniques are often imprecise or require strong priors. The most advanced overcome it but include very sparse data. That is why we have developed below our on-site and portable acquisition method that can be scalable across a range of systems. The goal is therefore to develop a system to measure the BRDF of an object by respecting the following criteria. First, the set-up has to be easy to transport for on-site acquisition. Then, it has to be inexpensive and based on classic photographic acquisition equipment (light spots, cameras, tripod...). But we remind that it should be scalable for the acquisition from small objects to monumental heritage elements. Finally, a relatively fast acquisition is preferred.

Having said that, we can build on ideas presented earlier. We see that in this context of acquisition, we must carry out a model driven approach over a data driven one. Although, we have seen that the crucial problem is the need to have different light and camera directions (so, positions) around the captured object. It means we should not be limited to paraxial illumination at the camera with an attached flashlight. That is the reason why we decide to develop a technique with a camera and a spot light we can move by hand for a better flexibility on a case-by-case basis. Also, [107] and [155] shows good results with a pipeline based on (1) a data processing pipeline in which we define the coarse 3D first, and (2) an iterative process around normals and SVBRDF optimization then (as presented in Figure 3.12). So we will be moving in that direction as well.

In the following, we present our new approach for a SVBRDF on-site acquisition method. We have separated the implementation into 3 different parts:

1. The acquisition process
2. Transforming image data into BRDF measurements
3. Obtaining BRDF models from the BRDF measurements.

3.3.1 Acquisition process

Here we present the acquisition part. In this part, we focus on the process developed to take pictures from various spots with different light positions, and a way to be able to calibrate these elements. For this, we have sought to develop a fairly simple and quick set-up while respecting the constraints listed above, using the equipment below:

- Two fixed-focus cameras
- A calibration checkerboard
- Two 6 cm mirror spheres and their supports
- 2 tripods
- A B10 ProFoto spotlight
- A computer
- A Specim IQ portable hyperspectral camera
- Associated connectors
- A spectralon

These different elements are visible during the measurement in the Figure 3.13.



Figure 3.13: *Set-up acquisition in laboratory. Example of the acquisition of a Balinese mask.*

We describe here the acquisition pipeline. As a reminder, we must multiply the camera points of view and the light positions, as part of a BRDF acquisition. The general idea of this acquisition is therefore first to obtain a 3D model of the scene, then to be able to quickly change the position of the lights and the camera, while quickly estimating their position each time. Note that we independently speak about the scene or the real object when we speak about the studied object. Indeed, both words can fit more if we are speaking about a handheld item or a architectural element. Here is the measurement process, which is also visible in Figure 3.14:

1. We place the calibration checkerboard near the scene to be measured, then two mirror spheres of radius $\delta_{rad,sph}$ on it, centered in $q_{sph,l}$ and $q_{sph,r}$. Finally we also place a camera in front, in a position that allows it to take a picture of the 2 spheres as well as the




Figure 3.14: *Acquisition process. First, we place C_l in front of the checkerboard and the mirror spheres (A). Then we place C_s in front of the real object and a spectralon (B). The following part is iterative. $\mathcal{L}_{a \in [1, m_{\mathcal{L}}]}$ is placed close to the real object, and images are taken by C_s and C_l (C). We change the position of $\mathcal{L}_{a \in [1, m_{\mathcal{L}}]}$ and do it again (D). After $a = m_{\mathcal{L}}$ iterations, we change the position q_s of C_s (E). The $\mathcal{L}_{a \in [1, m_{\mathcal{L}}]}$ position step is done again (F).*

checkerboard. This camera will be called in the following the light camera C_l , at the position q_l (Figure 3.14A). We set its focus manually on the mirror spheres and block it.

2. A complete photogrammetry is carried out including the scene, as well as the checkerboard. This photogrammetry has 2 goals: to create the 3D mesh on which we are going to work, and to place the mirror spheres precisely in the space of the scene. Here we refer to the 3D mesh when we speak about the object from the photogrammetry, and the 3D object is about the whole captured object.
3. One or more hyperspectral images of the scene are produced using the hyperspectral camera for the material segmentation. These images can also be achieved with a classic camera and the segmentation will be done with the RGB images then.
4. We place the second camera (call the scene camera C_s) facing the scene to be measured, on a tripod at the position q_s . A spectralon (a lambertian white reference material) is also placed close to the scene (Figure 3.14B). Then we take a first HDR photo without light in order to have our “reference black”, being in an uncontrolled environment. Like the other camera, we set its focus manually on the scene and block it. It must be blocked until the end, even if the camera is moved.

5. We perform the following steps in a loop in order to multiply the number of light positions $m_{\mathcal{L}}$ and camera positions m_s as much as possible:
 - We place our light spot \mathcal{L} (with the position $q_{\mathcal{L}}$ and the direction vector $\mathcal{D}_{\mathcal{L}}$) which illuminates the scene, and we take a picture of the mirror spheres with the light camera \mathcal{C}_l (Figure 3.14C). Based on a method inspired by *Corsini et al.* [46], the goal here is to place \mathcal{L} in the space of the scene very quickly, thanks to a simple photo.
 - We then take an HDR photo of the scene with the scene camera \mathcal{C}_s .
 - We move the light spot \mathcal{L} and perform the previous two steps again (Figure 3.14D).
 - Once several light positions $q_{\mathcal{L}, a \in [1, m_{\mathcal{L}}]}$ have been obtained, move the scene camera \mathcal{C}_s , and repeat the previous steps, being careful not to change the focus (Figure 3.14E and Figure 3.14F).
6. Once the measurement is finished, we remove the cameras \mathcal{C}_s and \mathcal{C}_l from the tripod while keeping the focus blocked, and we take a picture of the checkerboard several times by moving it, in order to carry out the geometric calibration of the two cameras.

After performing this acquisition pipeline, a lot of data is generated such as:

- One or more hyperspectral images of the scene for the material segmentation. These images can also be RGB photos.
- Photos for the photogrammetry of the scene containing the mirror spheres and the scene itself
- Images of the mirror spheres under different light positions $q_{\mathcal{L}}$ from \mathcal{C}_l
- HDR images of the scene from \mathcal{C}_s
- Images for the calibration of the light camera \mathcal{C}_l and the scene camera \mathcal{C}_s

From this data, it is now necessary to set up a second process, which is generating BRDF measurements from these images.

3.3.2 From image data to BRDF measurements

In this part, we describe the data process (photos, spectral images and mesh) starting from the exploitation of the data acquired on site, until the creation of a file containing for each material an RGB intensity value associated with a couple (ω_i, ω_o) of the 3D object. The exploitation to realize a coherent SVBRDF model is then obtained from these data.

As explained previously, we decide not to measure a total SVBRDF (i.e. each triangle in our model would have an associated BRDF) but a simplified SVBRDF, for which each triangle is associated with a material β , itself associated with a BRDF ρ_{β} .

The explanatory diagram of the complex pipeline that we are setting up is visible in Figure 3.14.

We need to get (ω_i, ω_o) and the material β for each pixel projected on the measured mesh, to generate our BRDF measurements input. Similarly, to obtain a BRDF value, it is necessary to have the precise distance from the camera $\delta_s = \|q_s - \mathbf{x}\|$ and the light $\delta_{\mathcal{L}} = \|q_{\mathcal{L}} - \mathbf{x}\|$ to the point \mathbf{x} on the 3D mesh, as well as a calibration of the light intensity of the spotlight. This is why we create 4 blocks of elements that allow to have one BRDF per material: (1) a segmented mesh, (2) a light intensity calibration, (3) the camera positions q_s , rotations r_s and intrinsic parameters A_s for the determination of ω_o , and (4) the light positions $q_{\mathcal{L}}$ and direction vectors $\mathcal{D}_{\mathcal{L}}$ for the determination of ω_i . It should be noted that the light intensity calibration is not absolute but proportional for reasons of *in-situ* constraints presented previously. Our measure of BRDF is therefore relative.

In order to generate the BRDF samples, it is necessary to add a 5th block: the generation of the 3D mesh by **photogrammetry**. We do not dwell on this part because it is very simple: the

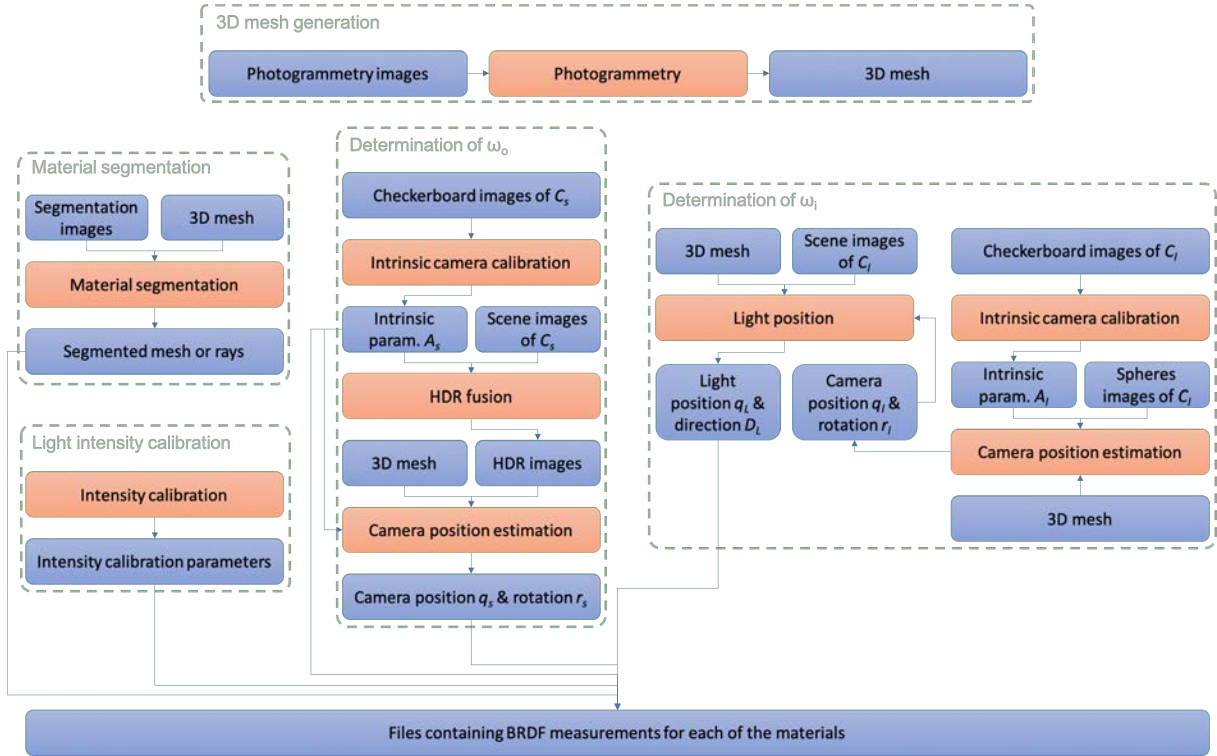


Figure 3.15: Process generating BRDF measurements from images data. The gray dotted squares represent the 4 major axes we developed. The blue boxes are the inputs and outputs generated at each step while the red boxes are the methods used.

photos with a flash taken for the photogrammetry generate a 3D model colored by SfM-MVS approach on the *Metashape* software. The extracted mesh is then used for the other blocks.

Likewise, we do not detail the **HDR fusion** red box as it is simply an HDR fusion based on a Gaussian weight function as described by Robertson *et al.* [182].

In the following subsections, we explain the other steps of the process (red boxes) of Figure 3.14 in more detail. We thus see how we arrange them in order to arrive at BRDF values.

Camera calibration

The camera calibration yet contains two parts: the intrinsic calibration and the extrinsic calibration. The intrinsic calibration makes it possible to determine the internal parameters of the camera, that is the distortion parameters (that we set aside here for simplification purpose but are integrated for distortion corrections on our images), and the camera matrix A containing the optical centers and the focal lengths of the camera. Then, the extrinsic calibration gives the rotation matrix R and the translation vector t , which makes it possible to obtain the position q and the rotation vector r of a camera. Indeed, the rotation matrix R can be obtained from r using the Rodrigues conversion [74]. Likewise, the position q can be obtained from $O = (0, 0, 0)$ the origin in the world coordinate. So we have $O = Rq + t$. Then we obtain q easily with:

$$q = -R^{-1}t \quad (3.14)$$

These parameters are necessary for C_s , but also for C_l , as shown in Figure 3.14. Indeed, they make it possible to obtain the rays at each pixel in the direction of the mirror spheres for C_l , and of the 3D mesh for C_s .

It is then necessary to carry out a calibration of the camera for the BRDF measurements. This calibration ultimately makes it possible to perform ray casting and ray tracing for each pixel of the images taken by the cameras C_l and C_s , based on the pinhole camera model.




Figure 3.16: *Pinhole camera model. Image taken from OpenCV documentation [6].*

Note: Ray casting for a pinhole camera

As explained in [6], the pinhole camera model is a simple model computing for given parameters, the link between a 2D pixel coordinate (u, v) and a 3D position in world coordinate $P_w = (X_w, Y_w, Z_w)$ (see Figure 3.16). In this model, we simplify the camera as a virtual pinhole camera and we connect 2D and 3D position with the equation:

$$sp = A[R|t]P_w \quad (3.15)$$

where A is the camera matrix, R the rotation matrix, t the translation vector, q is a 2D pixel in the image plane, and s a scaling factor. So we can rewrite it :

$$s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} [R|t] \begin{bmatrix} X_w \\ Y_w \\ Z_w \\ 1 \end{bmatrix} \quad (3.16)$$

where the parameters are visible in Figure 3.16. By solving the system, we can then obtain from our parameters the normalized direction vector for each pixel (u, v) defined by $\hat{p}_w(u, v) = (x_y, y_w, z_w)$ [6]:

$$p_w(\hat{u}, v) = \frac{p_w(u, v)}{\|p_w(u, v)\|} \quad \text{with} \quad p_w(u, v) = \begin{bmatrix} \frac{u-c_x}{f_x} & \frac{v-c_y}{f_y} & 1 \end{bmatrix} R \quad (3.17)$$

The ray cast process is done as followed: for each pixel of an image, we can now generate a vector coming from the center of the camera into the world space. This is crucial in our case: we can cast a ray from the camera which is crossing the 3D mesh, and we know exactly what is the correspondence between a pixel RGB value, the direction vector p_w , its intersection on the mesh \mathbf{x} and the normal n at this position. At the same time, we see then that it is necessary to know the parameters of A , as well as R (and t for its position) to obtain this vector. That is a reason why a calibration is necessary.

This section does not contain any new contribution relative to the actual state of the art in Computer Vision, but it is necessary to explain it precisely for the following.

Intrinsic calibration. In this part, we basically use OpenCV [5]. The methodology and the functions used are based on [222] and [36]. The principle consists in taking a large number of photos of the same black and white checkerboard (in our case, 15 x 15 squares, visible in Figure 3.13) from different positions.

First, we detect the corners of the squares on the image using *cv.findChessboardCorners()* function. This openCV function gives as an output the corners as 2D pixel positions. Then, *cv.cornerSubPix()* is used to refine them precisely.

Next, we use the corners found to get the intrinsic parameters of the camera using openCV function *cv.calibrateCamera()*. Based on [222] and [36], it returns the camera matrix A as well as the distortion parameters of the camera, but also r and t , respectively the rotation vectors and the translation vectors of each of the images accordingly to the images taken.

In our case, we use direction distortion parameters with *cv.undistort()* to undistort all images of \mathcal{C}_s and \mathcal{C}_l . This is why we do not mention it any further. Similarly, r and t are not useful here as they are the relative poses to the checkerboard. Finally, we get A which is necessary for the ray casting and the extrinsic parameters calibration.

Camera pose estimation. After obtaining A , we estimate the translation t and the rotation vector r of each camera position in world space in order to get R and q . This is called the pose estimation, or the extrinsic calibration of the camera. There are 2 techniques to obtain these parameters: 2D-3D points matches, known as structure-based pose estimation or 2D-2D points matches (structure-less pose estimation).

Structure-based methods are usually based on a **Perspective-and-Point algorithms** (PnP) [136]. It globally consists in solving the r and t vectors that can minimize the reprojection error from 3D-2D point correspondences. These 3D points are usually placed on the mesh and the correspondence is done manually. The main interest is that we only require the image we want to estimate, and the mesh. But 2D-3D pose estimation methods are constrained in their effectiveness by the accuracy of the 3D point triangulations and the quality of the generated 3D model.

On the other hand, structure-less techniques are based on **Resectioning** [224]. Here, the estimation of the pose for a given image relies on establishing correspondences between the query image and two or more images within the reconstruction, achieved through 2D-2D matching. Even if they are more precise, these pose estimation techniques tend to exhibit higher computational complexity and frequently yield numerous potential solutions. This problem is in addition to the need for several images around the mesh.

In our case, we choose to focus on a PnP estimation. This choice is made possible by better integration into the global pipeline: it is easier for us to obtain the extrinsic parameters from A , an image and a mesh, rather than several images. Then, in the case of resectioning, it is necessary to have a starting image precisely located around the mesh. That is to say that it is necessary either to use the PnP on a first image (which therefore have an imperfect position), or to start from the images captured for the photogrammetry on *Metashape*. This then means placing points on a large number of images. However, the placement of the points being done manually, it is much longer to place the correspondence points by hand on our photogrammetry photographs than simply on the mesh and the studied image. In addition, tests were carried out on *Metashape* and the positions did not suit very well in comparison to PnP. We therefore apply a PnP method using the OpenCV functions of pose estimation [7].

Camera \mathcal{C}_s . We separate the pose estimation of \mathcal{C}_s and \mathcal{C}_l into two distinct cases. To begin, we look at the first one. Several PnP algorithms implemented in OpenCV are tested: PnP Iterative [136], EPnP [126], SQPnP [199], and RANSAC method using PnP Iterative [136]. For each of the methods, we place 2D and 3D points manually on the mesh and some images.

We then compute r_s and t_s for each of them. Next, we reproject the 3D points on the images and compute the square error between 2D points and their reprojections. The reprojection of 3D points on the image is visible in Figure 3.17 left while the reprojection error for each method is visible in Table 3.1 According to our tests, the PnP Iterative method with or without RANSAC gives the best results.

Table 3.1: Reprojection error for each of the PnP method for the Balinese mask acquisition. Input parameters used are the parameters by default.

	Iterative	EPnP	SQPnP	RANSAC (on Iterative)
Error	14.9	220.4	109.7	14.9

We originally place the position of \mathcal{C}_s roughly in a scene containing the 3D mesh on *Blender*. Then, *Blender* positions and rotation are converted into $r_{s,orig}$ and $t_{s,orig}$ (see Annex IV). It is indeed necessary to have initial values, as Iterative PnP is based on a Levenberg-Marquardt optimization, which then requires an initial solution [64]. We finally use *cv.solvePnP()* for the PnP Iterative method implementation and so we get r_s and t_s , then R_s and q_s and visualize the position of the camera according to the mesh (see Figure 3.17).




Figure 3.17: Illustration of the pose estimation for \mathcal{C}_s . **Left:** Photograph taken during the acquisition. White dots are the points selected by hand on the image while red ones are the reprojections of 3D corresponding points. **Center:** View through the virtual camera on *Blender* having the same intrinsic and extrinsic parameters as \mathcal{C}_s (for the conversion from *OpenCV* to *Blender*, see Annex IV). **Right:** External view of the camera and the mesh on *Blender*.

Camera \mathcal{C}_l . In the case of \mathcal{C}_l , we also execute a PnP pose estimate using Iterative algorithm. But in this situation, we use the checkerboard for a finer estimation. At first, we place manually 5 to 10 points on some corners of the checkerboard in the 3D mesh. From this, we create an analytical checkerboard in the scene. As we know the coordinates of the 3D points, as well as the real dimensions of the squares, we can execute a rigid transformation of an analytical checkerboard to position it precisely (see Figure 3.18 left). We can then extract most of the 3D corners (we choose to only extract those visible on the image of \mathcal{C}_s) from this analytical checkerboard, as well as their position on the image thanks to the function *cv.findChessboardCorners()*. We finally have a large number of 2D-3D correspondences, which we use to find r_l and t_l , then obtain R_l and q_l with PnP Iterative method as previously. As before, the position can be checked by reprojecting the 3D points on the image (see Figure 3.18 right). Note that this part is quite close to the process used for the intrinsic parameters but in our case, we already know A_l so we focus on PnP pose estimation instead of *cv.calibrateCamera()*.

At this stage, we hence have the parameters that lead to the determination of ω_o , as well as those at the base of the pipeline for the determination of ω_i for the light calibration.




Figure 3.18: Pose estimation for C_s . **Left:** 3D graph of the rigid transformation of an analytical checkerboard (in blue). The red dots are the position placed by hand, and their labels are their coordinates on the checkerboard. **Right:** Photograph from C_l of the checkerboard and the mirror spheres. White dots are the 2D points on the image while red ones are the reprojections of 3D corresponding points.

Light calibration

Here we focus on the light calibration. The goals can be separated in two parts, which can be compared to the camera calibration: the intrinsic and the extrinsic parameters of a spotlight. The first part is about centering on the determination of a light model corresponding to our situation while the second one concentrate on the pose estimation of a light fitting to the model.

Radiometric calibration of light. We use a Profoto B10 which is adjustable in color and intensity (visible in Figure 3.13, on the bottom left). This is why we do not seek to carry out an absolute radiometric calibration of the light intensity emitted. Indeed, we perform a relative BRDF measurement, as explained above. Instead, we observe the relative angular intensity variation of the source, and try to see if a model would suit our situation better than the point light model.

Point light model. Acquisitions methods coming from the state-of-the-art do not usually care about the light profile. Many of them use a simple flashlight that is modeled by a point light model. This model consists of assuming that light is emitted from a point in space, uniformly in all directions. Globally, coming from eq. 3.1, the BRDF can be written:

$$\rho(\omega_i, \omega_o) = e \frac{dL_o(\omega_o)}{dE_i(\omega_i)} = e \frac{dL_o(\omega_o)}{\mathcal{F}_i} \frac{\delta_{\mathcal{L}}^2}{\cos \theta_i}, \quad (3.18)$$

where e is a scaling term to take into account that we our measures are relative and not absolute ones, $dL_o(\omega_o)$ is the reflected luminance we get from the acquired image \mathcal{I}_s . $dE_i(\omega_i)$ is the incident irradiance, \mathcal{F}_i the flux/power of the light source. We introduce the view factor $V_{\text{ponct}} = \frac{\delta_{\mathcal{L}}^2}{\cos \theta_i}$ between the source and the surface (i.e. the proportion of the radiation which leaves the light source that strikes the surface measured [44]). Finally we can write:

$$\rho(\omega_i, \omega_o) = e \frac{dL_o(\omega_o)}{\mathcal{F}_i} V_{\text{ponct}}, \quad (3.19)$$

Disc light model. But in our case, the source is relatively large: it has a diffuser disc with a diameter $\delta_{rad,\mathcal{L}} = 10$ cm. It is therefore more judicious to model the source by a disc light model. It consists in modeling the light as an extended source having the shape of a disc, emitting in a diffuse way. The light is then much more directed, because the apparent surface of the disc is much larger in front than on the side. Similarly, the disc is oriented: only one side illuminates. In this case, the view factor can be then replaced by the formula given by Hollands [100] and corrected in [8]:

$$V_{disc} = \frac{1}{2} \left((\cos \theta_k - H \cos \theta_j) + \frac{HZ \cos \theta_j}{P} - \frac{1 + H^2 - R^2}{P} \cos \theta_k \right) \quad (3.20)$$

with $\begin{cases} H = \frac{h}{a} & Z = 1 + H^2 + R^2 \\ R = \frac{r}{a} & P = (Z^2 - 4R^2)^{1/2} \end{cases}$

where the parameters given here are visible in Figure 3.19 (for a reason of the multiplicity of variables, variables used in eq. 3.20 and Figure 3.19 are only used here).




Figure 3.19: View factor from differential element at coordinate system origin tilted at arbitrary angle to a disc bisected by the y - z plane. Plane of element does not intersect disc. Image taken from [8].

And then the BRDF can be written:

$$\rho(\omega_i, \omega_o) = e \frac{dL_o(\omega_o)}{\mathcal{F}_i} V_{disc} \quad (3.21)$$

Further, we see that we therefore need 3 parameters to define a disc source. So we call the radius $\delta_{rad,\mathcal{L}}$, the position of the source defined at the center of the disc $q_{\mathcal{L}}$ and the direction vector of the disc $\mathcal{D}_{\mathcal{L}}$.

Validation. To demonstrate our model is more accurate than a point light model, we perform a simple laboratory test. We place the spotlight in front of an almost lambertian wall. We then place the camera behind the spot in front of the wall too, and we take a picture of the wall (see Figure 3.20 left). Then, we calibrate our camera as seen previously.

Afterwards, a simple simulation with raytracing is done with a lambertian wall, a camera and a light. The camera and the spotlight are located at the same positions as in the real configuration, and we use the camera matrix obtained after calibration for the camera model. We adjust the value of each pixel using the eq. 3.21 with the view factors of a punctual and a disc light model V_{punct} and V_{disc} . We compare them in Figure 3.20 right.

We can see that simulation is much closer to real data with a disc model. Moreover, if we notice a gap on the edges, it is also because the wall is not perfectly Lambertian. The real data




Figure 3.20: Validation process for the disc source model. **Left:** Picture of the process in lab. **Center:** Image taken by the camera (bottom), transformed into a intensity image instead of RGB (top). **Right:** Graph comparing the intensity along the horizontal axis of the real image, and the simulations with a disc and a punctual light. The 3 curves are rescaled at their top in order to be compared (relative BRDF measurement).

curve is therefore slightly less intense on the edges than the simulations. But despite this, our results show that it is more physically realistic to use the BRDF expression based on V_{disc} .

In the following, we find the values of BRDF by applying eq. 3.21 to the flux values extracted from each pixel of the HDR images taken by \mathcal{C}_s , called \mathcal{I}_s . This BRDF is nevertheless proportional as we do not know precisely \mathcal{F}_i , so we add a proportional term a to the eq. 3.21. The radius of the disc being known, we then need to determine $q_{\mathcal{L}}$ and $\mathcal{D}_{\mathcal{L}}$ for a complete light characterization.

Light pose from light probes. Here, we develop a new fast way to estimate the light pose from parameters and images of \mathcal{C}_l , and the 3D mesh. The process can be divided as (1) the position estimation of the mirror spheres, (2), the parameters optimization of the light source using Inverse Rendering. For this second part, we use a disc source model in the following, but we also show that this method works for a point sphere light model or a sphere light model as well.

Mirror spheres positions. At first, it is necessary to obtain with precision the center $q_{sph,l}$ and $q_{sph,r}$ of the two mirror spheres in 3D space. However, on the mesh, the spheres are very poorly represented because photogrammetry has trouble with reflective surfaces. Even by selecting a point by hand, we would not be at the center of the sphere.

We then rely on the work of Lucat [137]. Below the process is explained for one sphere but it is applied to both. First, we perform a circle detection based on Hough Circle Transform OpenCV functions to detect the sphere outlines (see Figure 3.21 left) . But the shape projected of a sphere onto the pixels is actually an ellipse and not a circle. This is why we select 6 points distributed all around the circle in a uniform way, which we move manually on the edges of the contours of the sphere. Then, we fit an ellipse which corresponds to the edge of our sphere on the image (see Figure 3.21 right). From it, we can obtain an accurate position of the center of the sphere $q_{sph,x}$ by ray tracing, applying the method from [137].

To explain quickly, it is a matter of casting rays from \mathcal{C}_l on the contours and the center of

the detected ellipse. As we know precisely the diameter of the sphere, we can then see how far the gap between the edge of the ellipse and its center corresponds to $\delta_{rad, sph}$. This is then the distance to the sphere, and its center passes through the ray casted along the center of the ellipse. Thus, one can very easily obtain a precise position approximated of the centers of each sphere. Figure 3.21 bottom underlines the efficiency of the method by showing how the real sphere is superimposed (the pixels detected in red are those within the ellipse) with a perfect analytical sphere at the position $q_{sph, x}$ obtained and of diameter $\delta_{rad, sph}$ (green pixels). Indeed, the vast majority of visible pixels are yellow and correspond to the overlay.




Figure 3.21: *Mirror sphere pose estimation. Left: Sphere or the Hough Circle Transform for circle detection. Middle: Ellipse fitting. In red are the 6 points manually placed. In green the ellipse fitted. Right: Validation process: inside the ellipse representing the real sphere on the photo (red pixels) and projected analytical sphere (green pixels) are overlaying well for left and right spheres.*

Inverse rendering for a point source estimation. Here we show how we get the position of a light source from our images and our parameters of \mathcal{C}_l and mirror spheres. For this, we carry out a simple inverse rendering, that is we seek to minimize the difference between a result that we obtain on a fictitious scene, and the studied image. In a first time, we show the simple case of a point source, then we present the more complex case of the extended spherical source, to finally tackle our case of the disc source.

First, a reference image is created from the studied image. We place two regions of interest (ROI) around the reflections of the lights on the spheres, then we apply a threshold manually in order to keep the reflections at 1, and at 0 the rest of the image. This reference image is called \mathcal{I}_{ref} .

Second, we create a fictitious scene, in which we have our camera \mathcal{C}_l and its previously obtained parameters, and two perfect spheres in $q_{sph, l}$ and $q_{sph, r}$, of radius $\delta_{rad, sph}$. We then roughly place one ponctual light in world space, most coherently with reality, at a position $q_{\mathcal{L}, orig}$. Moreover, in this case, we define the surface of the sphere as a quasi-mirror surface, reflecting light with a GGX whose α is very close to 0.

From this scene, we generate an image \mathcal{I}_{scene} similar to \mathcal{I}_{ref} . As for \mathcal{I}_{ref} , we apply a threshold so as to keep only the specular reflection in the spheres at 1, and the rest of the image at 0. We therefore have an image obtained which is a function of the position of the point light $q_{\mathcal{L}}$, and of α (the spheres and the camera being fixed). We can finally define an error err_1 to be minimized such that

$$err_1(q_{\mathcal{L}}, \alpha) = \|\mathcal{I}_{ref} - \mathcal{I}_{scene}(q_{\mathcal{L}}, \alpha)\| \quad (3.22)$$

A nonlinear minimization of the error is then carried out using Sequential Least Squares Programming (SLSQP). The optimization of $q_{\mathcal{L}}$ from $q_{\mathcal{L}, orig}$ is used to find the precise position

of the light in coherence with the image capture while α allows above all to adjust to a imperfect model.

Inverse rendering for a sphere source estimation. However, the need to adjust our spheres with an α from GGX shows that our model is not realistic enough. Moreover, we know that our source is not punctual. This is why we now replace the point source with a spherical light of radius $\delta_{rad,\mathcal{L}}$ at position $q_{\mathcal{L}}$. In this case, we assume that our mirror spheres are now perfect mirrors. So we therefore have the error to minimize:

$$err_2(q_{\mathcal{L}}, \delta_{rad,\mathcal{L}}) = \|\mathcal{I}_{ref} - \mathcal{I}_{scene}(q_{\mathcal{L}}, \delta_{rad,\mathcal{L}})\| \quad (3.23)$$

This case is a good transitional case between disc and point sources because it represents an extended source, and more realistically takes into account the shape of the specular reflection on \mathcal{I}_{ref} . Indeed, now the size of the reflection shape is not only a question of the light radius, but also of its distance. It is therefore more taken into account. But this is a simpler case of the disc one: a spherical source remains spherical regardless of the viewing direction.

Inverse rendering for a disc source estimation. But in our case, we do have an extended source, but it is a disc. The disc is a more complex source to optimize, because it includes one more parameter: its direction $\mathcal{D}_{\mathcal{L}}$:

$$err_3(q_{\mathcal{L}}, \delta_{rad,\mathcal{L}}, \mathcal{D}_{\mathcal{L}}) = \|\mathcal{I}_{ref} - \mathcal{I}_{scene}(q_{\mathcal{L}}, \delta_{rad,\mathcal{L}}, \mathcal{D}_{\mathcal{L}})\| \quad (3.24)$$

However, $\mathcal{D}_{\mathcal{L}}$ is a unit vector. It is therefore necessary to constrain it during the SLSQP minimization such that:

$$\begin{cases} -1 \leq \mathcal{D}_{\mathcal{L}}(i) \leq 1 & \text{with } i \in (x, y, z) \\ \|\mathcal{D}_{\mathcal{L}}\| = 1 \end{cases}$$

A second problem is the visible shape of the reflection on the mirror sphere. Several directions for the same position can give the same shape. There are therefore several possible local minima. This is why it is necessary to be careful to manually define the initial parameters $\mathcal{D}_{\mathcal{L},orig}$ and $q_{\mathcal{L},orig}$. In our case, we define them by placing the disc source on the scene in *Blender*.

We can note, however, that we can measure $\delta_{rad,\mathcal{L}}$ by hand on our spotlight. It is possible to keep it as a parameter of the optimization as a way to check if the method works, or to fix it and have only two vectors to optimize. Apart from these particularities, the minimization is identical to the previous cases (see Figure 3.22).

After these steps, we therefore obtain all the parameters resulting from the light calibration. The intrinsic calibration allows us to model our light as being a disc source, while the new position calibration is used to obtain a position, a direction, and a size to this disc. Thanks to this parameters, we can obtain a ω_i direction for each rays coming from \mathcal{C}_s , as well as a BRDF value from the flux received on the pixel.

Material segmentation

The last part to achieve before having SVBRDF measurements is to operate the material measurement. In the following part we present a segmentation based on RGB calibrated images, and a segmentation based on hyperspectral images obtained from the *Specim IQ* camera used in Chapter 2.

Segmentation from RGB images. In this section, we want to obtain the material of each ray of \mathcal{C}_s . For this, the segmentation is done in two steps: (1) an image segmentation by material, (2) the passage of material information from the image to the BRDF (via the mesh or not).




Figure 3.22: *Schema of the inverse rendering pipeline for a disc source. **Top line:** Generation of the \mathcal{I}_{ref} image from the real image. **Bottom line:** Generation of the \mathcal{I}_{scene} image from a synthesized scene with parameters we have measured before. Here, we show the mirror spheres in gray in order to get a better understanding of the image. **Right:** The comparison between the two reflection shapes on the mirror sphere surface is finally done for optimization. The difference between the two images is computed and minimized.*

Image segmentation. First, we segment an image taken for each position q_s in ambient lighting, as wide and diffuse as possible (see Figure 3.23 left). The purpose of this image segmentation is to separate the different materials that are assumed to have different colors. To do this, we transform our image from RGB to xyY and then we keep only the chromaticity xy. Indeed, xy contain the chromaticity (i.e., the color) of the material while Y contains its luminance. However, within the framework of a BRDF, we can simplify to our samples that, roughly, the same material has the same color, and varies especially in intensity. This is not completely true but this assumption remains consistent for a segmentation of relatively different materials.

Next we segment our xy image with K-Means [140]. For better results, we indicate a number of clusters approximately 1.5 times to 2 times higher, then we manually merge the clusters of the same materials. This allows to correctly prevent fusions of one material while another will be separated into two clusters.

However, the results have limitations. Some areas appear more complex to be correctly labeled: for example overexposed specular reflections or shadow areas too black to have a correct xy. It is also necessary to pay attention to metamerism if the materials begin to become more complex.

Results of the segmentation is illustrated in Figure 3.23 center. In the following, we call this segmented image the material map M_{mat} .

From segmented image to segmented BRDF. Images are taken by \mathcal{C}_s at known poses, so ray casting seems to be easy to realize. From there, there are then two complementary ways to segment the SVBRDF: the indirect projection on the mesh, or the direct method on the image.

The indirect method consists of projecting the material map M_{mat} onto the mesh. The intersected triangle is then associated with a value β of material connected to the corresponding cast pixel. Then, each ray cast from \mathcal{C}_s for the BRDF measurement is assigned to a corresponding




Figure 3.23: *Left:* Image used for the segmentation. Example with the Balinese mask. *Center:* material map M_{mat} . One colour is corresponding to one material. *Right:* Segmented mesh with the previous image, in the case of indirect method use. As previously, each colour is relative to a material. As we can see, results are globally coherent but some little offset appears due to the pose uncertainty.

β relative to the intersected triangle. The main advantage of this method is that it requires far fewer photos. It is not necessary to segment an image for all positions of \mathcal{C}_s . In addition, we can also cross different positions to fill the “holes” with information due to specular reflections, shadows and occlusions. At the end, we also obtain a fully segmented mesh, which can be reused later to make the 3D object associated with a complete SVBRDF (see Figure 3.23 right). But it remains drawbacks. The projection is done on the images in which pose estimation was made by PnP. It is therefore not perfect, and can present a little offset. But above all, it depends enormously on the geometrical resolution of the mesh which must be in adequacy with the resolution of the image. On a too dense mesh, a large number of triangles will not be intersected and must be assigned to a material by filling, while on a poorly resolved mesh, triangles overlapping several materials can exist. The ray tracing that follows can therefore include a lot of error in the assignment of material to each BRDF measurements.

That is the reason why we focus on, the direct method. It consists of performing an image segmentation for each point of view of \mathcal{C}_s then directly allocating a material β per pixel. This pixel (u, v) will have β for material for each ray launched for the measurement of BRDF. The method is direct because the material information is directly transmitted from the material map M_{mat} to the measurement image: there is no passage through the mesh. This is one of its main drawback, because it can only be used to segment the measurements and not the 3D object. In addition, it requires performing a segmentation for each point of view, which can take longer than the indirect method. But it is much more precise: the material error during the measurement ray tracing of \mathcal{C}_s is much lower. That is why we use it in our pipeline.

However, as in the previous case, the main limitation lies in the segmentation of the material map. As we demonstrated in Annex III, RGB values are not enough to properly segment materials, and a minimum of 10-15 spectral values would be needed.

Segmentation from hyperspectral images For the previously named reason, we also show that we can segment our BRDF by adding HSI. The objective is to take advantage of the high spectral resolution to overcome metamerism. We use the *Specim IQ* camera for this, and we integrate the hyperspectral images into our segmentation pipeline.

First, we perform the image segmentation on the hyperspectral image. We are in much simpler cases than in Chapter 2, so we can just perform a segmentation by K-Means by taking as data for each pixel the array containing all the spectral values after normalization of the spectrum. The advantage is that the segmentation is unambiguous. Complex areas due to metamerism are segmented correctly (for example, the difference between yellow or golden areas is easier). The problem with the *Specim IQ* remains its low spatial resolution. We therefore have a less precise segmentation at this level.




Figure 3.24: Homography to get a distorted segmented image which can be used for the material segmentation. **Top left:** Image used for the segmentation. Example with the Balinese mask. Points are the correspondence points placed manually. **Top right:** Image from C_s with known parameters. The same points are placed. **Bottom left:** Homography matrix computation from the corresponding points of the segmented image (in red) and the C_s one (in white). **Bottom right:** The segmented image is plotted here on the C_s image with an transparency of 30%. This distorted image is the material map M_{mat}

The second main problem occurs when going from a segmented image to a segmented BRDF with the direct method. Indeed, we do not know the intrinsic and extrinsic parameters of the spectral camera. In addition, the pinhole camera model is not necessarily suitable for a pushbroom camera. It would hence be very complicated to calibrate the Specim IQ, taking into account its low spatial resolution and its acquisition time, for a result that may be incorrect.

To solve this problem, we perform a **homography** between the RGB image from C_s (with known parameters), and the segmented hyperspectral image. A homography is basically an isomorphism of projective spaces, which is used to make a transformation between two planes, that is two images. To apply it, we compute a homography matrix, which is a 3×3 matrix but with 8 degrees of freedom as it is estimated up to a scale. This matrix allows the transformation of the position of a pixel (u, v) to the position (u', v') .

So we place 10-20 match points manually on the two images, then we calculate the homography matrix from the segmented image to the image of C_s . From this matrix, we distort the image. This final transformed image can then be used for the step of passing material information from the image to the BRDF as a material map M_{mat} . Homography method is illustrated in Figure 3.24.

At the end of the segmentation, we therefore have a value $\rho_\beta(\omega_i, \omega_o)$ for each ray launched from each position of C_s intersecting the mesh. We consequently converted the image data into BRDF measurements.

3.3.3 From BRDF measurements to BRDF model

We now have some raw BRDF measurements to deal with. The objective of this section is to show how these measurements are processed to obtain a realistic 3D object. For this, we optimize both our BRDFs via a model fitting approach, as well as the distribution of normals on the surface of the mesh. So, the developed pipeline is an iterative process, in order to enhance both the geometry and the appearance simultaneously. It is presented in Figure 3.25.

After the step presented previously allowing to pass from data images to a set of BRDF values $\rho_\beta(\omega_i, \omega_o)$, a very large amount of data is generated. But this data is not necessarily representative (many (ω_i, ω_o) pairs are similar), and contains a lot of noise. So a first step is about the **BRDF decimation and noise filtering**. Then, for each material, we apply a **BRDF model fitting** followed by a **normal optimization** on the normal maps. The normals being modified, this leads to a modification of (ω_i, ω_o) according to Figure 3.4. We then return to the first step “From image data to BRDF” in which we will just modify the values of its pair. Afterwards the different steps are carried out again. The loop is ultimately calculated a second time before ending the process. To finish, the last step presents the **generation of the realistic 3D object**, containing the UV maps for the normals and the BRDF parameters.

These different steps are presented in the following sections.

Measurements processing

First, the BRDF data is processed in order to be visualized and modelled. This step can be split in two parts: the data decimation and the noise filtering, presented below.

Data decimation. As seen previously, we have a very large amount of data (in the case of the Balinese mask example, we have more than 15 million measurements). However, this data is very redundant, and requires significant decimation. But this decimation must be done without eliminating the interesting values. For this, we then start by performing an angular segmentation (K-Means) on $(\theta_i, \theta_o, \Delta\phi)$, on our data. The goal is therefore to have n clusters of equivalent sizes representing well the whole of the angular distribution that we have on our measurement. Then, we randomly select 1 measurement out of 500 within these clusters which we keep for the rest. This allows us to greatly reduce the amount of data, while remaining representative.

Noise filtering. Then, the BRDF values are very noisy and have to be filtered. It can be explained by the imprecision on the segmentation of material, the presence of occlusion, or some inaccurate normals. It is therefore necessary to filter our data with an averaging type low-pass filter. For that, we allocate for each BRDF measures the average value of the x closest measurements (of the same material) in the $(\theta_i, \theta_o, \Delta\phi)$ space, to which we add a weight. We take x chosen upstream, and the weight used here corresponds to the inverse of the distance to the filtered measurement in this space. The larger x is, the stronger the filtering will be. In this case, the diffuse aspect will accordingly be more precise, but there is a risk of losing the specular peak.




Figure 3.25: Process generating a realistic 3D object with SVBRDF models and corresponding normal maps from the BRDF measurements.

Measurements visualization. The processed data can thus be observed. For this, we move into Rusinkiewicz space and first propose the visualizations visible in Figure 3.26.




Figure 3.26: BRDF visualizations of the golden material of the Balinese mask.

The graph on the left simply offers a 3D view of the BRDFs in the $(\theta_h, \theta_d, \Phi_d)$ space that can be manipulated in order to see what is the representativeness of the BRDF in the light and viewer hemispheres. On the right, the figures are **BRDF slices** as visible in Figure 3.3. We have represented them in the form of grids in which the measured BRDFs are averaged in each square. The 4 grids present the number of measurements, the color, the chromaticity (or tint), and the luminance for each square. Thanks to this, we can see what are the variations of intensity and color in the diffuse or specular parts, as well as the number of values measured in each zone. But the slices remain difficult to interpret and in particular do not allow an easy comparison with a fitted model.

So we can also represent the BRDFs in the form of a **polar plot**. This view, used Figure 3.2, displays a cross-sectional view of the material. An incident ray at a value θ_i is plotted, and the BRDF is displayed as a lobe: for each angular direction, the larger the lobe, the greater the reflection intensity. Figure 3.27 thus presents polar plots for several light incidences. In our case, to display the BRDFs, we linearly interpolate our measurements, and give a value of zero to the areas beyond our limit angles. That's why we have to see it only as a truncated polar plot.




Figure 3.27: Polar plots of the golden material of the Balinese mask with $\theta_i = 1^\circ$, $\theta_i = 30^\circ$ and $\theta_i = 50^\circ$.

The advantage of this visualization is that it is more telling than the previous one. Moreover, it is much easier to superimpose the measured BRDF and a fitted BRDF model for comparison than in the previous case. But it has the drawback of only showing the lobe in the case $\theta_h = 0$. The information is therefore interesting and complementary to visualization in the form of a slice, but remains incomplete.




Figure 3.28: Normals fitting. Material maps $M_{app,k,l}$ are made from $M_{mat,k}$, $M_{norm,k}$ and the BRDF models fitted in (a). Then, the material maps $M_{app,k,l}$ and the HDR images taken $\mathcal{I}_{s,k,l}$ are compared in $M_{error,k}$ in (b). This error map is used as the error for the optimization of $M_{norm,k}$ (c).

Fitting the appearance

As we see on the slices, our measurements are sparse. This is why we adopt a model fitting approach, based on GGX microfacets model presented in section 3.1.2.

Fitting of a Microfacets Model. In our case, we thus have as parameters to optimize k_d the RGB triplets of the lambertian part, η the Fresnel refractive index (also a RGB triplet in our case as we define F varying for R, G and B channels according to Cook-Torrance results [45]) and α the roughness. In order to determine these parameters, we implement a cross-fitting method on our data. First, we determine an original value of k_d by taking for each RGB channel the top of the histogram of BRDF values. Yet we define a first error by:

$$err_4(k_d, \alpha, \eta) = \sum_o \sum_i (g_{fit}(\omega_i, \omega_o) - g_{meas}(\omega_i, \omega_o))^2 \cdot \sin \theta_o \quad (3.25)$$

$$\text{with } \begin{cases} g_{fit}(\omega_i, \omega_o) = -\log(1 + \rho_{GGX, [k_d, \alpha, \eta]}(\omega_i, \omega_o) \cdot \cos \theta_i) \\ g_{meas}(\omega_i, \omega_o) = -\log(1 + \rho_{meas}(\omega_i, \omega_o) \cdot \cos \theta_i) \end{cases}$$

based on Low *et al.* [135], with $\rho_{GGX, [k_d, \alpha, \eta]}$ the GGX model we want to optimize and ρ_{meas} our BRDF measurements. Low *et al.* show that a log-based error metric performs better than an error directly on the BRDF because it captures more of the wide-angle scattering. The result on the specular can be a little bite depreciated but the overall result is improved.

After that, we optimize $Err_{k_d, \alpha, \eta}$ to find our parameters. We first perform a nonlinear minimization (*L-BFGS-B* method) on α and η keeping the value of $k_{d,init}$ obtained previously, and with $\alpha_{init} = 0.4$ and $\eta_{init} = (4, 4, 4)$. Next, we recalculate k_d with the new α and η with a linear least squares optimization. As long as the error decreases, we continue this iterative alternation between the *L-BFGS-B* optimization on α and η , and the least squares for k_d . Once the error starts to rise, we stop and keep our obtained parameters.

We finally obtain the parameters k_d , α and η corresponding to our fitted BRDF, for each material.

Fitting the Normals. At this stage, we actually get a GGX model fitted for each material. Moreover, we can project our materials correctly on the mesh thanks to the material maps $M_{mat,k}$ that we have for each pose k . We thus generate for each k a normal map $M_{norm,k}$ thanks to the parameters of corresponding \mathcal{C}_s . It is important to note that $M_{norm,k}$, $M_{mat,k}$ and the HDR images $\mathcal{I}_{s,k}$ taken by $\mathcal{C}_{s,k}$ are perfectly aligned. By ray casting, we can therefore deduce an appearance map $M_{app,k,l}$ by applying the formula obtained from eq. 3.21 for each of the light poses measured l from the pose k :

$$\mathcal{F}_o = \mathcal{F}_i \frac{\rho_{GGX,[k_d,\alpha,\eta]}(\omega_i, \omega_o)}{V_{disc}} \quad (3.26)$$

in which the flux \mathcal{F}_o is proportional to an intensity value, that is a pixel value. We finally get $m_{\mathcal{L}}$ appearance maps $M_{app,k,l}$, which can be compared to the images $\mathcal{I}_{s,k,l}$ taken. To compare them, we look at the Symmetric mean absolute percentage error (SMAPE) map which can be detailed for each pixel (u, v) as:

$$M_{error,k}(u, v) = \frac{\sum_l |M_{app,k,l}(u, v) - \mathcal{I}_{s,k,l}(u, v)|}{\sum_l |M_{app,k,l}(u, v)| + |\mathcal{I}_{s,k,l}(u, v)|} \quad (3.27)$$

which can be actually rewrite for each pixel (u, v) related to a pair (ω_i, ω_o) as:

$$M_{error,k}(u, v) = \frac{\sum_l |a\mathcal{F}_i \frac{\rho_{GGX,[k_d,\alpha,\eta]}(\omega_i, \omega_o)}{V_{disc}} - \mathcal{I}_{s,k,l}(u, v)|}{\sum_l |a\mathcal{F}_i \frac{\rho_{GGX,[k_d,\alpha,\eta]}(\omega_i, \omega_o)}{V_{disc}}| + |\mathcal{I}_{s,k,l}(u, v)|} \quad (3.28)$$

where a is the fixed proportional coefficient. This form of SMAPE with sums in the numerator and in the denominator makes it possible to suppress as much as possible the outliers due to cases where the image is black, such as in the occlusion zones.

The error map $M_{error,k}$ is finally minimized by taking the normal map $M_{norm,k}$ as minimization parameter. Indeed, by modifying the normal n of each pixel, we then modify the values of (ω_i, ω_o) . We can therefore optimize the appearance obtained, and get a better normal map containing the microdetails on the surface of the mesh.

The optimization process is based on a SLSQP nonlinear minimization, and is done pixel by pixel. It can consequently take a long time (for an object like the Balinese mask, the minimization lasts 20 hours on *Python* after parallelization on the CPU). The whole process is illustrated in Figure 3.28.

As we take several origins of different lights, we can tend towards a coherent result. By then multiplying the poses k of \mathcal{C}_s , we can globally reconstruct the normal maps on the surface of the mesh.

Crossed Optimization between Normal and BRDF. At this point, we actually have a new normal map $M_{norm,k}$ for each pose k . The measurements of ρ_β obtained at each pixel of the images $\mathcal{I}_{s,k,l}$ no longer correspond to the corresponding pairs (ω_i, ω_o) measured in section 3.3.2. Indeed, changing the normals n modifies (ω_i, ω_o) . Moreover, according to eq. 3.21, the value of ρ_β itself is modified by the change of (ω_i, ω_o) induced.

We must therefore resume the data by recalculating (ω_i, ω_o) and ρ_β for each pixel of $\mathcal{I}_{s,k,l}$ at each pose k and light l . This computation is achieved from the new n from $M_{norm,k}$. So, we resume at the first stage on the Figure 3.25.

Next, the different steps are achieved again iteratively. Due to the very high computation time of the normals fitting step, we limit ourselves to 2 or 3 iterations. The results tend to converge very quickly, so the gain does not seem significant enough in balance over time.




Figure 3.29: *Synthetic scene created on Blender.*

3D object generation. We finally have normal maps, material maps, as well as fitted BRDF models. We create from materials maps and BRDF models, a map for each BRDF parameter: M_α , M_η and M_{k_d} . From there, we simply reproject them on our 3D mesh to obtain a realistic 3D object. We implement it in *Blender* using a script in *Open Shading Language* (see Annex IV).

In the context of a more complex object, we project the materials maps onto the mesh so as to obtain a segmented mesh, as we saw the indirect method in Section 3.3.2. From this we generate create UV coordinates with a Trivial Per-Triangle texture parameterization. Finally, new UV maps are created: the normal maps as well as the M_α , M_η and M_{k_d} maps are generated. In case of inconsistency for the same triangle, we use the map with the largest dot product $\omega_o \cdot n$. As previously, results can be seen on *Blender*.

3.4 Results and evaluation

To assess the reliability of our method, we develop some evaluations. The validation of this chapter is actually more complex than in the previous chapters. We therefore neither get a measurement as simple as a thickness (Chapter 1), nor test color charts to compare (Chapter 2). So, the best way to validate the process is to develop relevant metrics for the different contributions we made. Initially, we create synthetic simple scenes on *Python* and *Blender* to validate and evaluate the different steps and the final results. Then, we use some metrics to compare the results obtained on real objects with the HDR images taken.

3.4.1 Evaluation on synthetic scenes

It is very difficult to compare each step individually on real acquisitions. For example, our SVBRDF is based on microfacets model parameters obtained from a fitting. In this case, its error on a real sample is hard to quantify. Similarly, it is difficult to find the pose of a light, because we are not looking for the location of the spotlight but of its equivalent model which can be slightly different.

This is why we create a simple equivalent and synthetic scene: we generate the scene on Blender as in Figure 3.29. This scene is relatively simple: we create two perfect spheres on a checkerboard, and a plane wall wich is a very basic case for the BRDF acquisition. We keep the size of the spheres, as well as the distance between them in order to have the most realistic scene. Next, we create \mathcal{C}_s and \mathcal{C}_l cameras, and some disc spotlights placed in the scene in such a way as to keep our order of magnitude distances. We then define the material of the object according to a GGX distribution with the following parameters: $k_d = (0.73, 0.54, 0.02)$, $\alpha = 0.380$ and $\eta = (4.18, 4.12, 4.06)$.

Finally, we export the light poses, the mesh of the wall, and all the camera parameters (see Annex IV). In the following, each part is evaluated individually, so we can estimate the uncertainty of each step of the process alone.

Light pose

At first, we evaluate the error on the pose of the lights. So we generate images in Python from our \mathcal{C}_l camera parameters, our spheres, and the actual \mathcal{L} light parameters. This image generation is done as in Section 3.3.2, successively for point, spherical and disc light sources.

Sphere position estimation. Here we evaluate the mirror sphere position estimation. We use the images computed previously, and apply the method developed in Section 3.3.2. We get the positions $q_{sph,l}$ and $q_{sph,r}$ and compute the distance with the real position: we find an error on the left of 1.97 mm and on the right of 2.26 mm, that is less than 3% of the diameter of the spheres.

Light pose estimation. Then, we take the real positions of the spheres, and images are generated for 4 light poses as visible in Figure 3.29. Next, we apply the inverse rendering algorithm using the synthetically generated images here as reference images (see Figure 3.30).




Figure 3.30: *Synthetic light pose. **Left:** Reference image based on the thresholding of lights on the synthetic image. **Center:** Superposition between the thresholding (in yellow) and the halo obtained after optimization. **Right:** New position obtained compared to the real one, in 3D.*

From the light pose estimation, we find a light position (and a direction for the disc) for each light in the scene. The average error is then calculated and presented in Table 3.2.

Table 3.2: *Error obtained for the evaluation of the light poses with a synthetic scene. % error on position is computed with respect to the distance between the real light and the center between the spheres.*

	Total error	% error
Point light source		
Position after inverse rendering (in cm)	0.75	0.52
Spherical light source		
Position after inverse rendering (in cm)	0.95	0.66
Radius after inverse rendering (in cm)	0.25	0.83
Disc light source		
Position after inverse rendering (in cm)	0.97	0.67
Direction after inverse rendering (in °)	2.6	/
Radius after inverse rendering (in cm)	0.42	1.41

Taking into account that our light source is systematically at more than one meter from the sphere mirrors, the results observed here in an ideal case are excellent. Of course, it must be nuanced by taking into account that our camera position and our intrinsic parameters are perfect. But the results presented here make it possible to validate the method in a synthetic framework.

BRDF fitting

Next, we evaluate the BRDF measurements fitting. We generate the images taken by \mathcal{C}_s from 3 different camera poses (visible in Figure 3.29) with the chosen BRDF parameters, as we do for $M_{app,k}$ in the Section 3.3.3 (see Figure 3.31). These images can be called $\mathcal{I}_{s,k,l}$ since they are our synthetic images used for BRDF measurements.




Figure 3.31: 4 synthetic images obtained on Python.

We apply the pipeline presented in Section 3.3.2 and we get a set of BRDF measurements ρ_β . From this set, we process our data as in Section 3.3.3 and we compute the microfacets model fitting described in Section 3.3.3. We get BRDF fitted parameters and we can render $M_{app,k,l}$ images.

For the evaluation, we compare two parts: the BRDF parameters themselves, and the rendering (i.e. the comparison between $\mathcal{I}_{s,k,l}$ and $M_{app,k,l}$), which is evaluated with a SMAPE for each of the RGB channel as:

$$SMAPE_k = \frac{1}{m_{l,k}} \sum_l \sum_k \frac{\sum_{(u,v)} |M_{app,k,l}(u,v) - \mathcal{I}_{s,k,l}(u,v)|}{\sum_{(u,v)} |M_{app,k,l}(u,v)| + |\mathcal{I}_{s,k,l}(u,v)|} \quad (3.29)$$

with $m_{l,k}$ the total number of images. We then also compare the results for 3 poses and for only one pose perpendicular to the plane object. Obtained results are visible in the Table 3.3.

Table 3.3: New BRDF parameters and error on the reprojection obtained for the evaluation of the BRDF fitting with a synthetic scene.

	BRDF parameters			Reprojection (%)		
	k_d	α	η	R	G	B
Reference	(0.73, 0.54, 0.02)	0.380	(4.18, 4.12, 4.06)	-	-	-
3 poses fitting	(0.74, 0.53, 0.04)	0.373	(4.00, 4.03, 3.82)	0.1	0.8	0.9
1 pose fitting	(0.77, 0.56, 0.07)	0.350	(3.59, 3.62, 3.33)	0.5	0.7	2.2

If the results are better with 3 poses of cameras, they are nevertheless very good with a single pose. We notice that overall, we find a slightly undervalued roughness and index of refraction in the specular, compensated by a slightly bigger k_d . These results though present an error of less than 1% overall on the reprojected images. We can also see on Figure 3.32 that the reprojections $M_{app,k,l}$ are close to $\mathcal{I}_{s,k,l}$.

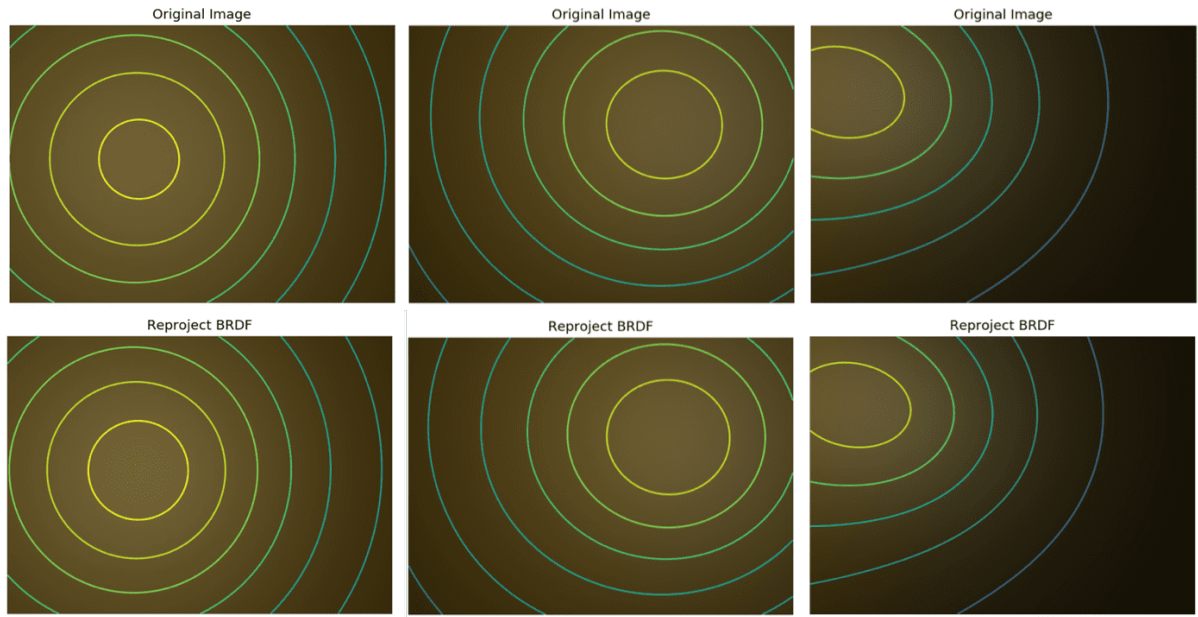


Figure 3.32: *Examples of BRDF reprojections for different k and l with synthetic data. The concentric circles are the intensity level lines for better visual comparison.*

Normals fitting

We then produce a synthetic image in order to see if we can find the normals correctly and evaluate the accuracy. To do this, we create a normal map of 30×30 pixels in which the normals are randomly generated with a variation of less than 80° in θ and random in ϕ . We then obtain a new noisy normal map.

Then, four images are generated with known BRDF parameters (k_d, α, η) from 4 given light poses and one camera pose. So we have input $M_{app,l}$ data for the optimization. In this part, the uncertainty is only on the value of the normals. It allows to estimate our results only by




Figure 3.33: *Mean difference between the normal obtained and the real one in degree as a function of α . Note that the error between initial value and real data is 23° .*

varying their parameters.

We can therefore apply our minimization algorithm pixel by pixel. In order to quantify the result, we calculate the mean error between the normals obtained and the real normals, in degrees. Moreover, we vary α in order to see its influence on the final error. Results are visible in the Figure 3.33.

The average error before optimization being 23° , we see that the optimization makes it possible to significantly improve the results obtained. Furthermore, we see that the roughness of the material has a rather important role in the final precision. We observe that the minimization works very well when we are on a diffuse part of the material. We therefore have very good results as we approach $\alpha = 1$, or when α is close to zero. In this specific case, the lobe will be on such a thin part that the majority of the viewer directions are found in the diffuse too.

A limitation that we note here is that a non negligible part of the measured materials has a roughness between 0.2 and 0.4. They are thus located in the most complicated area to measure, even if the normals fitting makes it possible to find more consistent values.

Evaluation of the iterative process

Finally, we evaluate the complete BRDF + normals fitting iterative process. So we generate the images $\mathcal{I}_{s,k,l}$ with the previously chosen BRDF parameters, while adding a random normal distribution (with an angle θ_{max} of 45°). As we have seen that roughness has an influence on normals fitting, we perform up to 3 iterations for several values of α . The results are visible in Table 3.4.

As previously, we notice that the specular part tends to be slightly less intense (η is lower globally), and compensated by a more intense k_d . For a very small α , the specular peak is hard to detect, and α will be bigger. Otherwise, we have a globally smaller α , except when we tend to 1, for which the result tends to assume that the material to be perfectly diffuse. This result is not surprising because Low *et al.* [135] demonstrates that our error metrics in eq. 3.25 tends to underestimate the specular peak a bit to better represent the rest.

Then we see that not having a perfect BRDF somewhat limits the normals fitting, even if the result obtained is much better afterwards and is inflated by some aberrant results.

Finally we notice by observing the error of reprojection that the results are better after two iterations, and then decrease. This confirms that only 2 iterations are enough.

3.4.2 Results on real objects

We carry out an acquisition in a controlled environment in order to test the global acquisition method (see Figure 3.13). The acquisition is made on the Balinese mask visible in the previous illustrations. For the acquisition, two Canon EOS 5D IV cameras are used, with 50mm (scene camera) and 35mm (light camera) lenses. The light spot used is the Profoto B10 that we calibrated previously. The photogrammetry of the scene has been realized with a Sony α 6000 camera. The photogrammetry of the mask has also been made with this camera, and contains 55 photos. For the acquisition, we have $m_s = 2$ camera scene poses, and a total of $m_L = 5$ light poses. We use a reduced number of views and lights because our object already presents a wide variety of normal directions.

Some examples of the BRDF measurements made on the mask are exposed in Figure 3.34. The slices show that the measurements relate to a fairly large area of the pair (θ_h, θ_d) . We notice that overall, the tint remains homogeneous, and that we have mostly a variation on the intensity on these materials.

Figure 3.35 illustrates the normals fitting process with a zoom on the right part of the normal map before and after normals optimization. We can see that it is a shading normal map

Table 3.4: *New BRDF parameters, error on the normals, and reprojection error obtained for the evaluation of the BRDF/normal fitting with a synthetic scene for various α at each step of the iteration pipeline.*

	GGX Parameters			Normals (°)	Reprojection (%)		
	k_d	α	η		R	G	B
Reference	(0.73, 0.54, 0.02)	0.010	(4.18, 4.12, 4.06)	-	-	-	-
BRDF fitting 1st iter.	(0.71, 0.53, 0.07)	0.086	(1.29, 1.31, 1.31)	13.3	12.3	15.1	59.0
Normals fitting 1st iter.	-	-	-	7.2	11.1	13.7	50.0
BRDF fitting 2nd iter.	(0.72, 0.54, 0.07)	0.056	(1.34, 1.36, 1.36)	-	11.0	13.1	45.7
Normals fitting 2nd iter.	-	-	-	7.5	10.6	12.1	42.5
BRDF fitting 3rd iter.	(0.73, 0.55, 0.07)	0.047	(1.34, 1.35, 1.35)	-	10.9	12.4	42.1
Normals fitting 3rd iter.	-	-	-	8.6	10.6	12.1	40.7
Reference	(0.73, 0.54, 0.02)	0.200	(4.18, 4.12, 4.06)	-	-	-	-
BRDF fitting 1st iter.	(0.81, 0.63, 0.20)	0.268	(2.96, 2.95, 2.69)	13.3	19.4	23.0	42.8
Normals fitting 1st iter.	-	-	-	11.3	11.1	13.3	26.4
BRDF fitting 2nd iter.	(0.85, 0.63, 0.10)	0.231	(3.42, 3.59, 3.78)	-	11.1	12.3	20.2
Normals fitting 2nd iter.	-	-	-	11.6	9.6	10.6	17.8
BRDF fitting 3rd iter.	(0.86, 0.64, 0.10)	0.237	(3.89, 4.13, 4.40)	-	10.2	11.2	17.9
Normals fitting 3rd iter.	-	-	-	12.5	9.5	10.6	17.6
Reference	(0.73, 0.54, 0.02)	0.380	(4.18, 4.12, 4.06)	-	-	-	-
BRDF fitting 1st iter.	(0.81, 0.64, 0.18)	0.369	(2.53, 2.43, 2.21)	13.3	10.8	12.9	29.4
Normals fitting 1st iter.	-	-	-	8.9	8.4	9.3	20.7
BRDF fitting 2nd iter.	(0.89, 0.68, 0.14)	0.340	(2.15, 2.33, 2.66)	-	9.1	10.0	19.5
Normals fitting 2nd iter.	-	-	-	9.1	8.4	9.3	18.8
BRDF fitting 3rd iter.	(0.95, 0.70, 0.15)	0.319	(2.04, 2.22, 2.54)	-	8.7	9.5	18.9
Normals fitting 3rd iter.	-	-	-	9.6	8.6	9.4	18.9
Reference	(0.73, 0.54, 0.02)	0.600	(4.18, 4.12, 4.06)	-	-	-	-
BRDF fitting 1st iter.	(0.88, 0.65, 0.18)	0.402	(1.00, 1.68, 1.48)	13.3	6.8	8.1	21.7
Normals fitting 1st iter.	-	-	-	8.4	7.7	8.7	17.9
BRDF fitting 2nd iter.	(0.89, 0.71, 0.14)	0.508	(1.00, 1.00, 2.06)	-	7.8	8.0	18.0
Normals fitting 2nd iter.	-	-	-	8.1	7.7	7.8	17.9
BRDF fitting 3rd iter.	(0.90, 0.72, 0.14)	0.523	(1.00, 1.00, 2.08)	-	8.0	8.1	17.9
Normals fitting 3rd iter. r	-	-	-	8.2	7.9	8.0	17.9
Reference	(0.73, 0.54, 0.02)	0.800	(4.18, 4.12, 4.06)	-	-	-	-
BRDF fitting 1st iter.	(0.80, 0.61, 0.14)	0.554	(1.60, 1.68, 1.30)	13.3	6.5	7.1	20.8
Normals fitting 1st iter.	-	-	-	8.0	8.5	8.4	18.3
BRDF fitting 2nd iter.	(0.76, 0.57, 0.14)	1.000	(2.90, 3.10, 1.63)	-	7.6	7.6	18.5
Normals fitting 2nd iter.	-	-	-	8.2	8.0	7.9	18.3
BRDF fitting 3rd iter.	(0.78, 0.59, 0.13)	1.000	(2.00, 2.53, 1.90)	-	7.9	7.8	18.4
Normals fitting 3rd iter.	-	-	-	8.7	8.0	7.9	18.3

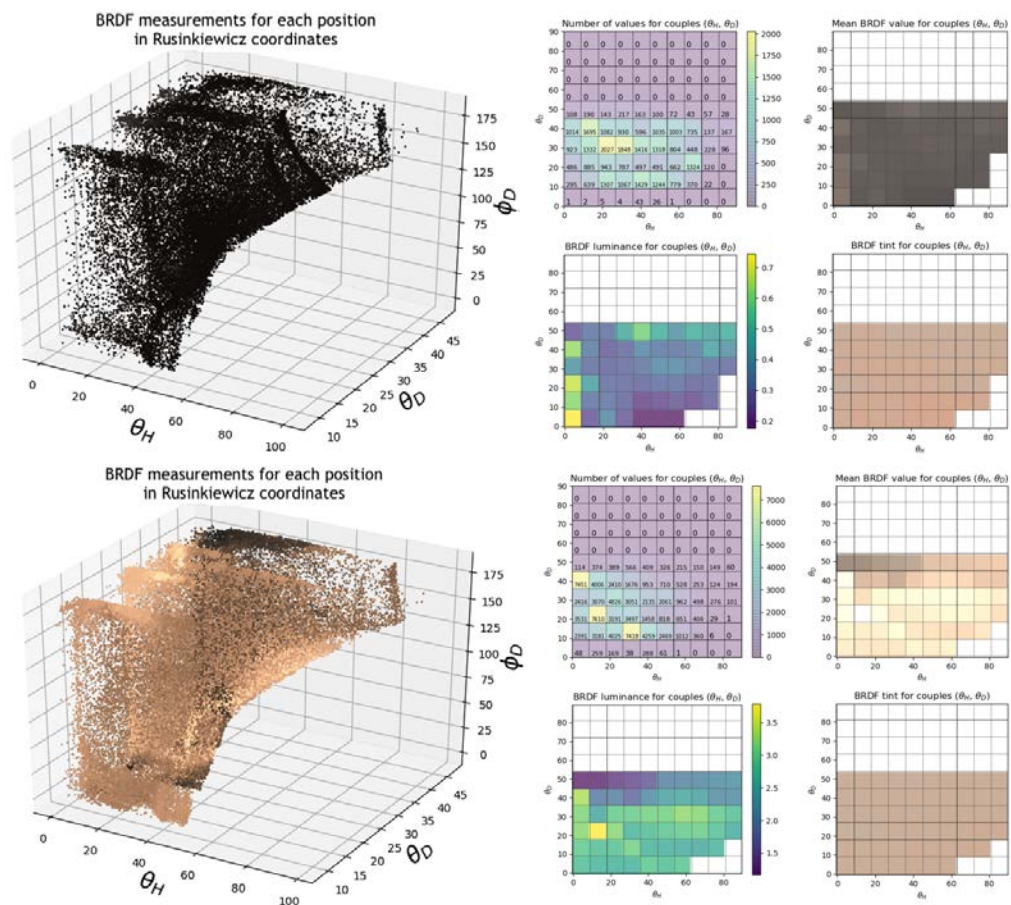


Figure 3.34: Examples of BRDF measurements made on the Balinese mask. First line refers to the black material and second line refers to the white one. See *Measurement visualizations* in Section 3.3.3 for more explanation.



Figure 3.35: Zoom on part of the normap map before (left) and after optimization (right).

rather than a 3D normal map, as we have variation that are not physically realist. But the micro-details of the white parts and the center of the eye are visible. Also, the golden crown shape is much better drawn after optimization.

Figure 3.36 shows some reprojection results. The reprojections are the $M_{app,k,l}$ generated after 2 iterations of BRDF and normals fitting. We observe a little variation in the specular part of the black material (for example in the specular peak in the eyes of the mask). In addition, the 3D object shows less similarity in grazing light. This can be explained by the overall difficulties in correctly acquiring the BRDFs in raking light (for $\theta_i \geq 60^\circ$). But the microstructures looks realistic on the diffuse (white) and shiny (golden) parts with an incidence closer to normal.



Figure 3.36: *Examples of comparison between re-simulated images $M_{app,k,l}$ (top) and real images $\mathcal{I}_{s,k,l}$ (bottom) for different light poses, in the case of the Balinese mask*

We calculate the SMAPE for each average channel between the real images and the re-projection ones, and we obtain (in %): **(19.4, 20.6, 21.9)**. But these results must be nuanced: they take into account both the errors on the BRDF and the normals, but also on the segmentation of the materials which can greatly increase their value. In addition, our image has significant dark areas. And areas where the base value is close to zero tend to greatly increase the SMAPE.

Figure 3.37 finally presents some renderings obtained on the mask after the 3D object generation with different angles.



Figure 3.37: *Examples of renderings of the Balinese mask.*

3.5 Application to Saint-Maurice residence

The previous results show acquisitions carried out in the laboratory. But the interest of our method is that it can be deployed on site, on larger scale objects. We thus use it in-situ in the case of the acquisition of the Saint-Maurice residence. Captures of the walls and a ceiling have been made in order to represent more faithfully the effects caused by the mosaics or the gildings (see Figure 3.38).

As well as for the Balinese mask, we employ a Profoto B10 as a light source. The cameras used were a Canon EOS 5D IV mounted with a 50mm lens for \mathcal{C}_l , and a Nikon D850 with a 60mm lens for \mathcal{C}_s . The acquired part of the wall measures around 1m x 1.60m, and the ceiling is around 4 x 6m at 4m high. Results are presented in Table 3.5.




Figure 3.38: *Left:* Capture of the wall. *Right:* Capture of the ceiling.

In the case of the wall, we can see that we get the different aspects of the stones used for the mosaic. Some stones are polished while other present a matte surface. So a 3D model with a global diffuse albedo cannot represent it well. But the SVBRDF acquisition underlines faithfully the shininess of mother-of-pearl and polished stone. Moreover, the normals fitting generates a precise micro-structure of the pieces of mosaic. It draws the borders of each pieces within the segmentation, and reveals the inhomogeneity of some stones (for example the red ones).

However, the acquisition is not perfect: light with normal incidence does not generate a specular peak as intense as in the real case. This problem is caused again by the choice of our BRDF fitting error metrics (see [135]). Note that the rendering and the restitution shows different colors. It is caused by the slightly yellow light used during the capture.

The ceiling originally presents the same issues as the wall. Some parts are recovered with gilding paintings. Renderings and reprojection presented in Table 3.5 shows that the golden parts are correctly captured. Globally, the reprojections look similar to the real ceiling. After all, the results remains encouraging for an in-situ acquisition.













In summary, we have presented a novel solution to enhance a 3D model acquired by photogrammetry, by adding an on site SVBRDF capture and a normal optimization for the micro details representation. Current portable methods are limited to simple materials or SVBRDF, or only measure the SVBRDF for particular (ω_i, ω_o) couples. Our method overcomes them by proposing a new approach to determine the light position, with a more physically realistic light model for the measurements. Thanks to this, we can compute new light directions fastly and create a real complete set of SVBRDF measures for a more realistic fitting. Moreover, we propose a global pipeline which is working on site, at different scale of objects.

As part of the Saint-Maurice residence study, we developed an acquisition pipeline that

enhance drastically the 3D model by incorporating the reflectance information within a dense model. This improvement opens new perspectives for more realistic representations of the architectural heritage thanks to the acquisition.

Moreover, in this thesis, it allows to show that the reflectance as the 3rd mentioned aspect of a 3D model can also be acquired on site. This acquisition can be achieved using a more physically based pipeline of measurements instead of a simple visualization method.

Table 3.5: *Example of on site results get in Saint-Maurice residence, Cairo.*

Real Images	Reprojected images	Renderings
Wall		
		
		
Ceiling		
		
		
		

Conclusion

During the study of Saint-Maurice residence, we developed different aspects of on site and non-invasive acquisition of visual information. The objective was to enrich the visual data of a classic 3D model including a RGB colored mesh. In this context, we have created methods to capture and process information on the shape, the color and the reflectance of material.

At first, we were interested in the acquisition of macro details of the shape of the objects, that is a capture of their 3D profile. In this context, we realized that an acquisition under the order of $100\ \mu m$ of resolution becomes very challenging because of the very narrow depth-of-field at such a scale. We have therefore used this limitation at our advantage by exploring Depth from Focus approaches [169] classical methods, and successfully achieved reconstruction at a resolution of around $10\ \mu m$. In addition, we used the limits on transparent materials of the Depth from Focus to measure in a fast and non-invasive way the thickness of ceramic glazes.

After having validated our method on various samples of ceramics, we used it to measure the glazes of the Iznik ceramics of Saint-Maurice residence. Preliminary study on samples from various time and origin has shown that depending on their origin, their age, and their uses, the ceramics nevertheless have large differences in glaze thickness. But the specific results on Iznik ceramics from Saint-Maurice residence then showed that the variations within the same tile did not make it possible to make difference between them for the study of their similarities. Despite this point, we have seen that the late Iznik ceramics of Saint-Maurice residence and the Izniks of Tunis studied have the same glaze thickness. This indicates that they seem to have been made by common manufacturing processes.

Secondly, we wanted to deepen the study of the Izniks by focusing on the used pigments. Unlike glaze thickness or drawing patterns, a pigment is prepared to manufacture a large set of tiles. It is therefore expected to vary less from one tile to another when they have a common origin and its pigments can be identified and studied by their reflectance spectrum [47, 79, 114]. This study therefore served as a pretext for the development of new processes and algorithms for processing hyperspectral data, particularly data from hyperspectral imaging. The main problem in the hyperspectral images data processing we faced is the enormous amount of information to be processed. So the first step therefore generally consists of a dimensionality reduction, that is to say that the number of values contained by a pixel is drastically reduced. It is in this context that we have developed our method which consists in using a database of reference spectra. From this base, we computed a Pearson correlation coefficient value with the studied spectrum for each reference spectrum. This allows a very strong reduction of data while preserving as much as possible the relevance of the information contained in the spectrum.

We used it on Aubusson tapestries [54] to see if we were able to identify red and yellow pigments for validation purpose. Yet we took hyperspectral images of the Iznik tiles in Cairo, using the portable *Specim IQ* camera. We applied the Pearson correlation method to the ceramic tiles to get a specific array of values to each of them. This allowed us to look at their similarity more precisely for making hypotheses in agreement with historians on common or different origins depending on the rarity of the patterns.

Finally, we turned our attention to another part of the residence: the restitution of the walls and ceilings. The palace actually presents mosaics and quite significant variations of materials that can be very matte or shiny. So a classic restitution in the form of a 3D model is not enough to represent the appearance as faithfully as possible. This is why we have endeavored to capture the properties of reflection of light in the palace (i.e., its SVBRDF). However, the current methods of portable acquisition of SVBRDF are not sufficiently precise in our case. Either they only work for very simple geometries or BRDFs, or they are limited in terms of directions.

This explains why we have come up with a new way to capture appearance. This process acquires BRDF properties for various (ω_i, ω_o) directions, then uses a model fitting approach for both BRDF and normals, to finally reconstruct the 3D object with new normals and BRDF properties instead of simple diffuse RGB colors. The problem that arises is then the validation of our acquisition pipeline. The BRDF being a complex property, we favored a synthetic validation only, by recreating a computer scene, and processing all the data obtained from there. Each stage of the pipeline could be tested and evaluated individually. Subsequently we tested it on an object in the lab to see if the results were visually consistent. In the end, we made the acquisition on a larger scale, on a ceiling and a section of wall of Saint-Maurice residence in Cairo.

We finally developed 3 axes of acquisition of visual information that are at first sight quite different. These 3 parts actually overlap quite well because they represent the whole of a global pipeline for a more physically and visually realistic approach possible to the capture of an object. First, we have the shape of the digitized object. Second, we have its diffuse color reflected in spectral format, which makes it possible to overcome metamerism and to represent it faithfully for any lighting. Third, we have its BRDF properties that faithfully represents how light is reflected, in particular to depict the specular aspect which is added to the color of the albedo mentioned just before.

We have also seen through our results that this data set can be used to answer (or participate in the validation of answers) to scientific questions on Cultural Heritage. This is the case here for measurement and analysis methods on ceramics, but also for better visualization of restitutions from the past and present. A line of research is then to push in the on-site measurement direction, for these visual properties. This axis is a real challenge because in-situ measurement is much more complex than in the laboratory. However, in the context of heritage, these methods are crucial because it is not always possible to transport samples to the lab, even more if the measurement is carried out for monumental heritage, that is on a large scale.

Future Work

The challenges of on-site acquisition we have identified, and the implementation of new techniques to address them, raise questions for future work. These possibilities revolve in my eyes mainly around 3 ideas: the improvement of current acquisition techniques, the improvement of access to a wider audience for their use, and their development to help answering more historical questions for Cultural Heritage.

Improving the acquisition technique is a necessary step for the other two points. For example, our SVBRDF acquisition method could be enhanced. To begin with, it would be interesting to make it more automatic because it currently requires a lot of human intervention, in particular on the segmentation of materials. This segmentation step could be greatly improved too. A direction on this subject could be to seek for a segmentation of the materials by their complete BRDF and not by their color alone. Indeed, the use of colors segmentation consists in segmenting by the diffuse albedo. This can cause great difficulties on very specular materials such as metals. A segmentation based on the BRDF measurements, for example using spherical harmonics (see [215] as a starting point), could make it possible to override this.

It may also be interesting to develop new devices for the acquisition of hyperspectral images. Indeed, we were limited by our hyperspectral camera because of its acquisition time and its low spatial resolution. By sacrificing some spectral resolution, which was much higher than our real needs, we could create a faster, better spatially resolved set-up. One idea could be to develop a system close to tunable filters devices by using filters placed on the flashlight. The spatial resolution would be maintained and the noise problems inherent to the filters placed in front of the sensor or the lens would be reduced.

Then comes the access to a wider audience. Reusing codes or algorithms is an important element for the reproducibility of results. This issue becomes even more challenging in the multi-disciplinary context of this thesis. Indeed, the data processing methods put in place can be used by historians or archaeologists, in addition to computer scientists. However, these are areas where people are not trained in computer science. So, in order to help them get to grips with the Depth from Focus algorithm and the Pearson correlation coefficient method, it will then be important to transfer it in the form of add-ons or software instead of Matlab or python scripts.

The third point is the development of the use of these methods by historians on a larger scale to provide them as a tool. The study of Iznik ceramics is a very good example. Initially, the acquisition of the spectra of all the ceramics tiles of Saint-Maurice residence would be a good start to further validate the obtained results. Then, it would be interesting to extend this analysis of the spectra to targeted samples of Iznik ceramic tiles present elsewhere in Cairo in order to constitute a much more consistent database for the researches about modern Cairo.


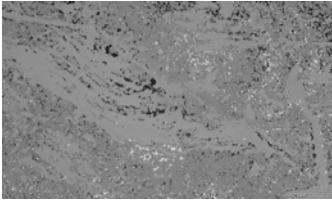
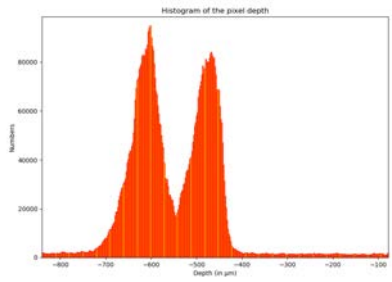

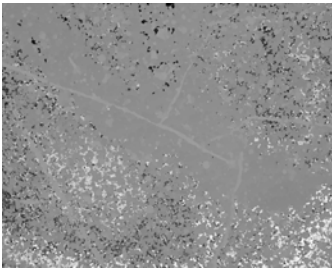
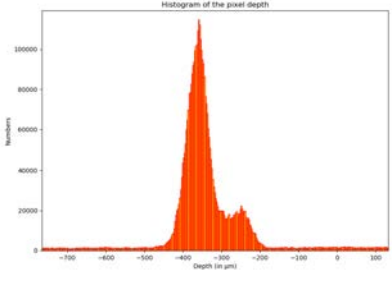

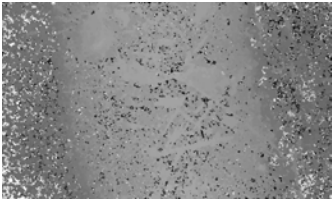
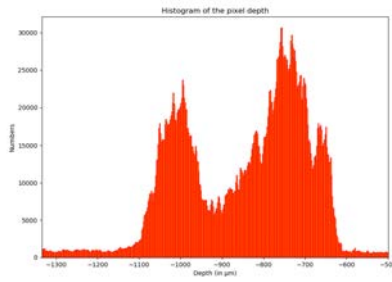
Similarly, measuring the thickness of ceramic glazes via Depth from Focus on Iznik ceramics of various origins and periods could be tested to see if a significant variation can be observed. For example, one could examine and compare Golden Age tiles (16th century) with a perfectly white glaze, and later tiles (17th and 18th century) with colored glazes. A variation in thickness could support the fact that the glazing technique itself has been lost.


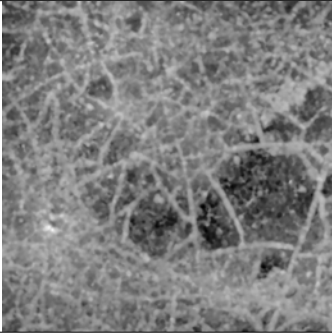
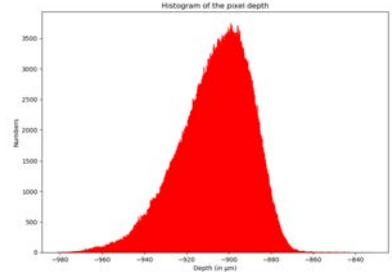

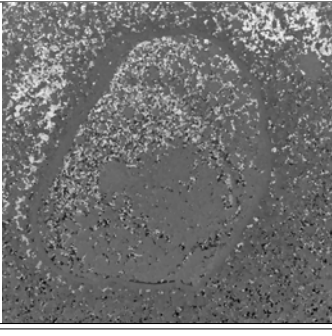
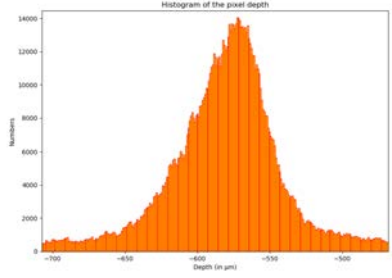

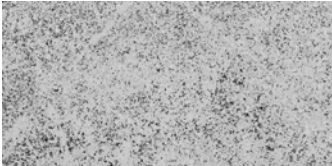
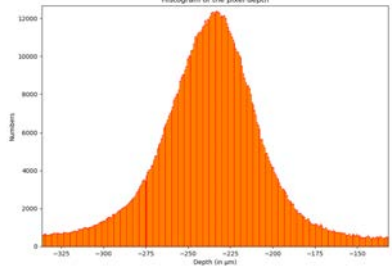
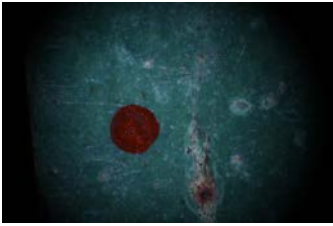
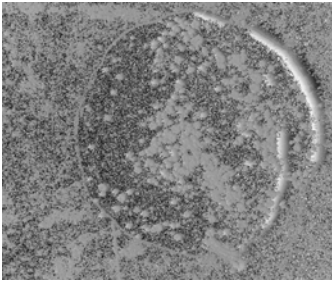
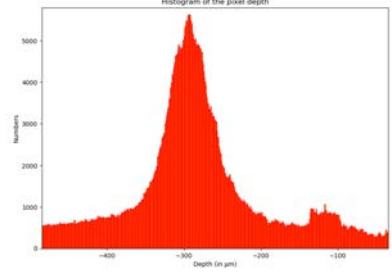
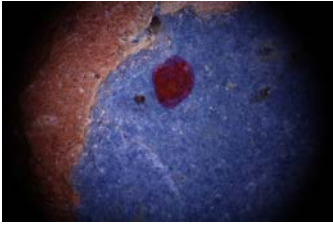
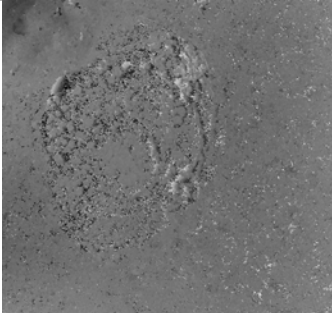
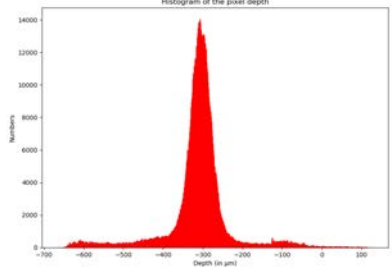
Annexes

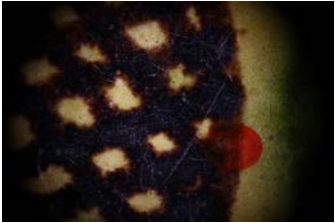
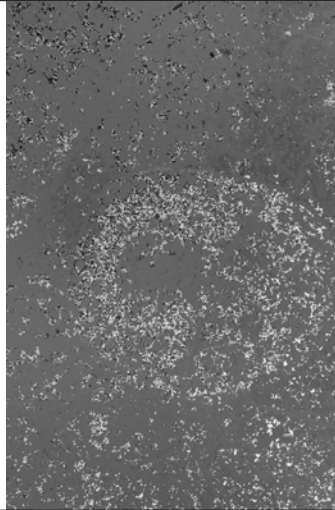
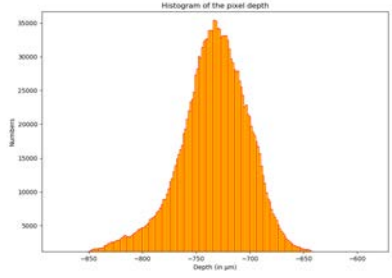
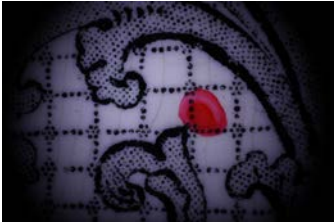
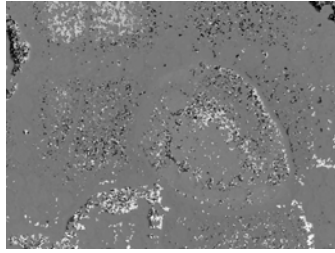
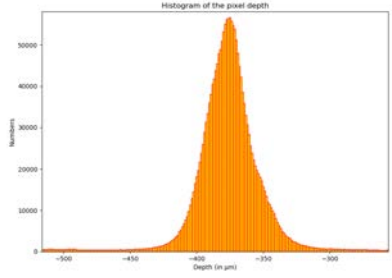

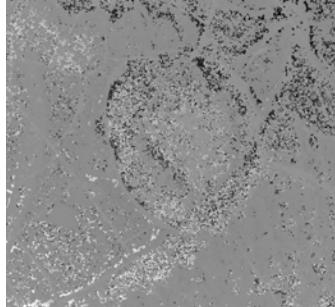
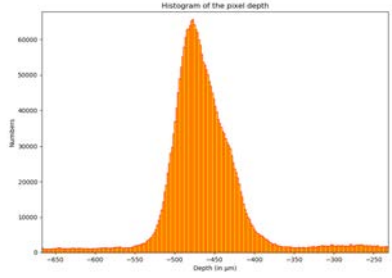
I Results of Depth from focus for ceramics glaze thickness

On lab measurements

Below is the visual data for each of the tiles we measured in lab for the validation of the glaze thickness measurement method. From left to right : all-in-focus image of the tile, corresponding depth map, histogram of the depth map. As we can see for some cases, despite seeing clearly the scratches on the depth map, the histogram does not present two clear peaks. In this case, an interactive visualization is required as explained in Chapter 1.


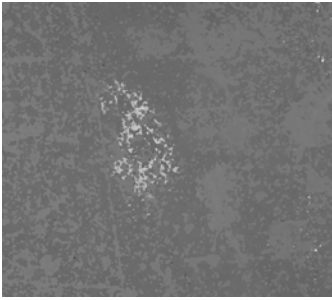
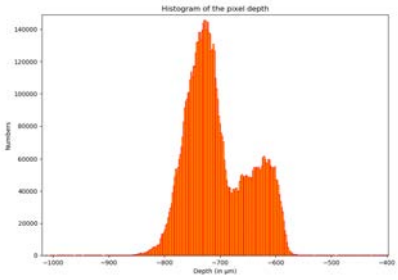
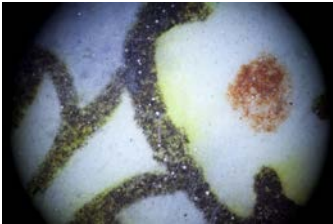

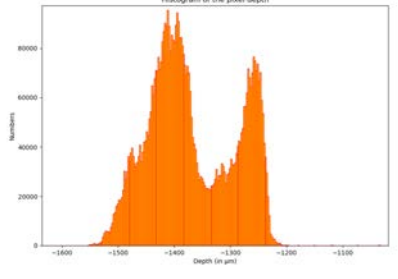

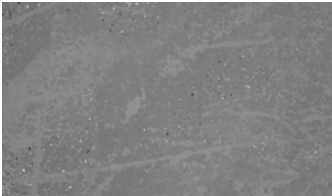
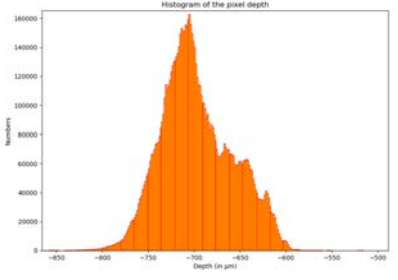
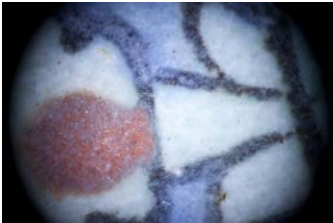
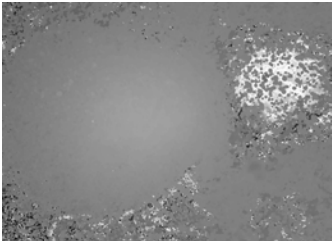
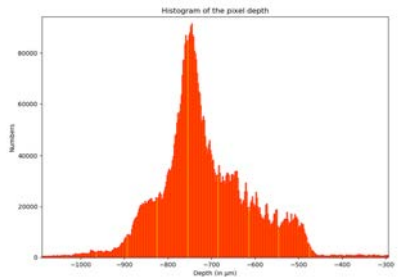
All-in-focus	Depth map	Histogram
BDX 6502 		
BDX 6493 		
BDX 6501 (glaze to red) 		

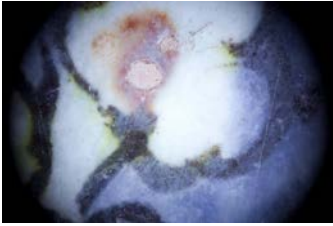
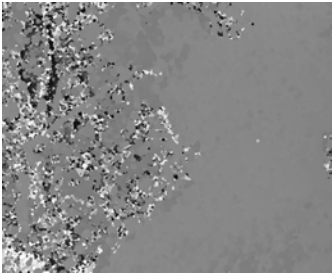
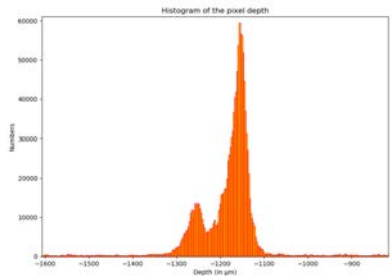
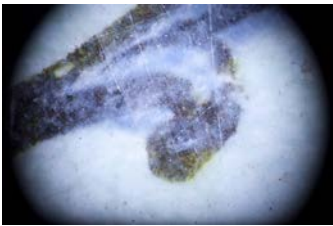
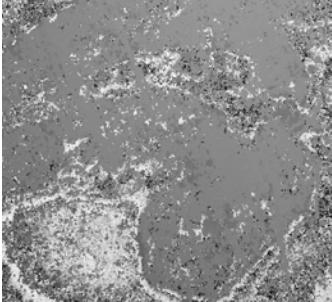
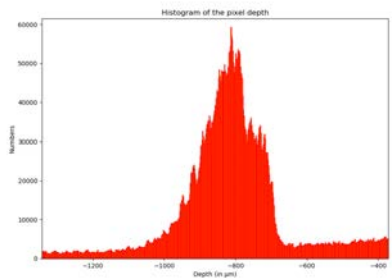

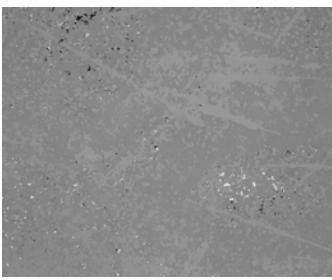
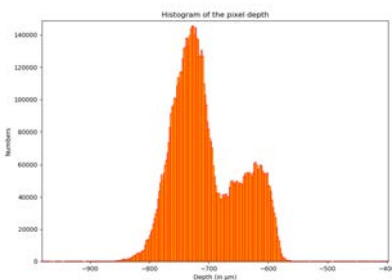

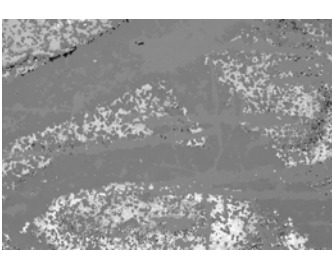
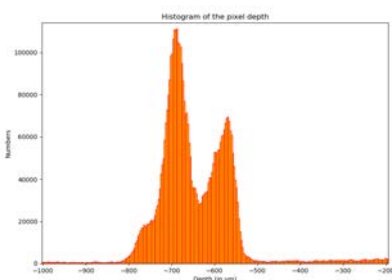

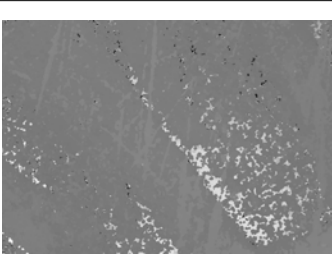
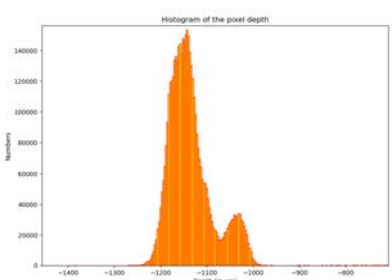
All-in-focus	Depth map	Histogram
BDX 16 621 		
BDX 2591 (white part) 		
BDX 2591 (yellow part) 		
BDX 6522 		
BDX 6524 		


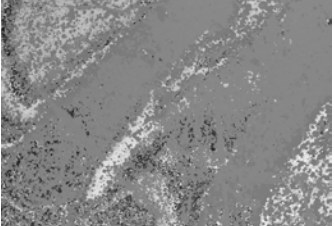
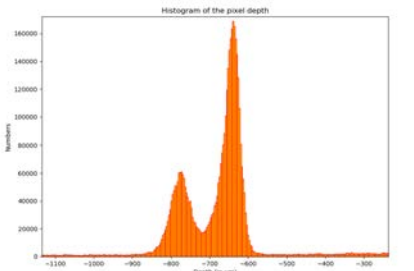

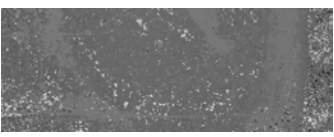
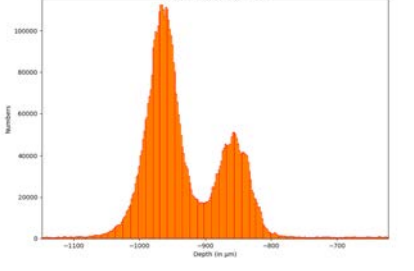

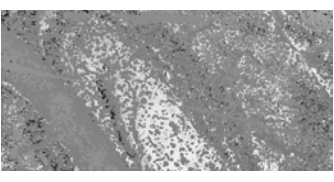
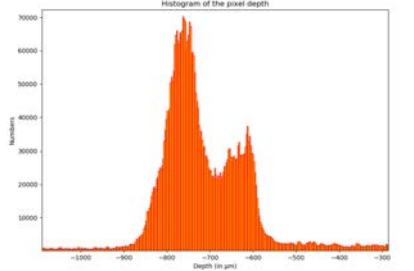

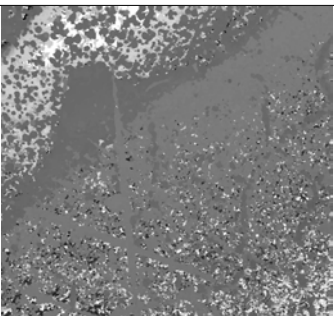
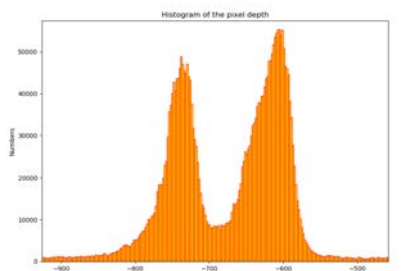
All-in-focus	Depth map	Histogram
<p data-bbox="194 230 304 253">BDX 5502</p> 		
<p data-bbox="194 775 316 797">BDX 21066</p> 		
<p data-bbox="194 1122 316 1144">BDX 22625</p> 		

Saint-Maurice residence measurements

Below is the visual data for each of the tiles we measured the glaze thickness in Saint-Maurice residence. From left to right : all-in-focus image of the tile, corresponding depth map, histogram of the depth map.

All-in-focus	Depth map	Histogram
80 P2 		
81 P2 		
87 P2 		
89 P2 		

All-in-focus	Depth map	Histogram
90 P2 		
91 P2 		
92 P2 		
1 C2 (A) 		
1 C2 (B) 		

All-in-focus	Depth map	Histogram
<p>1 C2 (C)</p> 		 <p>Histogram of the pixel depth</p>
<p>4 C2 (A)</p> 		 <p>Histogram of the pixel depth</p>
<p>4 C2 (B)</p> 		 <p>Histogram of the pixel depth</p>
<p>4 C2 (C)</p> 		 <p>Histogram of the pixel depth</p>

II Hyperspectral resolution resampling and impacts on our results

As it was shown in Chapter 2, the use of hyperspectral imaging (HSI) systems is a considerable advantage compared to classical RGB cameras. Because of metamerism, two different spectra can exhibit the same RGB color under a given illuminant. It is for this reason that we used HSI as main part of the study of pigments on the Iznik ceramics of Saint-Maurice residence.

In this context, as we have previously seen, we used a push-broom type camera: the Specim IQ camera. Its main interest is that it has excellent spectral resolution (204 spectral bands). However, the use of the Specim IQ portable camera has major flaws that a classic RGB camera does not have [192]. First, its high spectral resolution requires a sacrifice of its spatial resolution. Indeed, we obtain frames of 512×512 pixels on our obtained data-cube. Within the framework of the study of Iznik ceramics, we cannot therefore make an image containing several tiles at the same time. Yet, the second major drawback of the push-broom camera is the extremely long acquisition time. As it scans for each pixel line on the data-cube, the acquisition time is multiplied by 512. It can then last more than 5 min in the context of on-site acquisition with lighting constraints. It is for this reason in particular that we were not able to acquire all the tiles even though we spent several days on site.

The idea of developing or using another camera technology then emerged from these issues. Some technologies allow to take only one acquisition (like snapshot), or a dozen (like tunable filter). The case of tunable filter cameras and their derivatives even seems to be the best compromise for us. We obtain effectively our data-cube with a transportable set-up, and with a high spatial resolution over a fast acquisition time. The sacrifice is then made on the spectral resolution. But in our case, the spectral resolution seems to be able to be reduced. Indeed, we are not looking to obtain a large number of precise bands, but rather to roughly recognize the type of pigment used, then to make comparisons of similarities between the pigments.

In this case, it is essential to know how many spectral bands are enough to keep a reliable and robust study. This is why we start from our data, and make a progressive resampling of our spectra until we see what the limits of spectral resolution are.

Resampling method

We thus resample our spectral values to observe when the quality of the segmentation deteriorates. In order to make this resampling, we start by applying a cubic interpolation [113] on our spectrum for having a non-discretized function. Then, we sample at regular intervals according to the number of values chosen upstream.

Results are visible below in Figure 39.

Impact on the pigments segmentation

With the resampling computed above, we will therefore try to see how low we can go without too much loss of harmful information. For different resampling values, we apply the pipeline of the PCCR method seen in chapter 2 and we observe the error with respect to the non-resampled data.

Here we see more precisely the pipeline. First, we start from the representative spectra of each pigment, for each tile. We perform an n -value resampling, as seen previously on the representative spectra. However, it should be noted that the 12 characteristic spectra have 120 values, and we do not wish to resample them. So, from the resampling obtained, we perform an interpolation on the new n values (see the blue curve in Figure 39). From this new interpolation, we sample at 120 values. The PCCR method is applied to these new representative spectra.

In a first time, we get a new database of correlation values. By simply calculating the least-squares distance between the original database and the new one, we obtain a distance value.




Figure 39: *Resampled spectrum. **Cyan:** starting spectrum. **Red:** values after resampling (from 204 to 10 values here). **Blue:** Interpolation on the new values.*

We finally apply this method to numerous resamplings between $n = 120$ and $n = 4$ values. The distance curve is plotted in Figure 40.




Figure 40: *Plot of the distance of the correlations as a function of the number of n sampling values. **Yellow:** The limit of resampling.*

We can see that the curve presents an elbow around 9 to 10 values of spectral resolution. It therefore seems that for the generation of our database by PCCR, it is not necessary to have an excellent spectral resolution.

However, here, the distance measurement is made only on the correlation coefficients and not on the sensitivity of results obtained by K-means within the framework of a tiles segmentation. We therefore seek to compute the distance between the subgroups obtained with the first set of data and those obtained after resampling. However, the subgroups obtained for each K-means after resampling are not necessarily the same, and some are closer than others. So we measure for each tile the euclidean distance between the center of its subgroup obtained at the start, and after resampling. Then we made the difference. Thus, if the subgroups are identical, the center is the same. Moreover, if two subgroups are extremely close, the distance is not too large, and neither is the error. By summing the whole, we therefore obtain a general value of distance. As before, it is applied to a large number of resamplings, and the curve visible in Figure 41 is drawn.




Figure 41: Plot of the distance of the segmentation as a function of the number of n sampling values. **Yellow:** The limit of resampling.

As for the distance on the correlation coefficients, we observe that the elbow of the curve is located at a spectral resolution of 9 values. Thus, it seems that at least 9-10 values of spectral resolution are enough to obtain a result as precise as with the Specim IQ camera. In the following, we limit ourselves to the observation of the error on the correlation coefficients because we see that the error on the segmentation generally follows that on the correlations, and that we finally did not focus on the segmentation in the Chapter 2.

Problem of the narrow band resampling, and wide band resampling

The results found are very interesting but raise a problem if we work with a spectral camera allowing this resampling to be obtained. Indeed, in order to obtain points so close to the spectrum, it is necessary to have very fine bands for each wavelength measured, which means a large light intensity, and therefore either noise, or a time significant exposure (which goes against our initial objective). So, we try to measure our results on wider bands, while keeping the peaks.

For this, we take a band shape following a Gaussian type distribution. We want the integral over the visible spectrum to be 1. In addition, we take place on a uniform distribution of the centers of our curves. We will then only manipulate the σ parameter of our Gaussians, defined by:

$$g(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2} \frac{(x - \mu)^2}{\sigma^2}\right) \quad (30)$$

We vary σ according to the sampling step, with $\sigma = k.p$, p being the sampling step and with several k . We then obtain the graph in Figure 42.

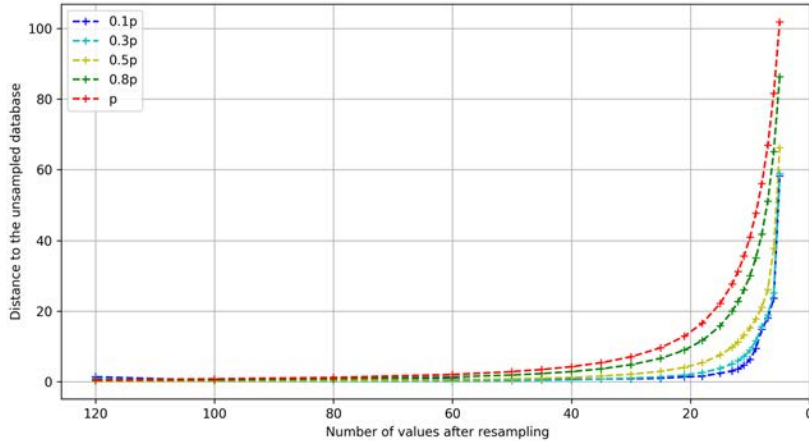


Figure 42: Plot of the distance of the correlations as a function of the number of n sampling values, for different σ

Thus we see that the modification of σ has an important influence on the accuracy of our PCCR. Indeed, a large σ will cause the elbow to bend towards $n = 18$ values for resampling.

Adaptative filtering for a better resampling

As we see, results are better with thinner bands while we want to develop some wider ones. An intuition would be to enhance the $\sigma_{f, f \in \text{filters}}$ of each filter individually in order to find a compromise and overcome it. So, we seek to optimize σ_f as well as possible for each filter individually. For this, we perform a nonlinear regression on the σ_f of the Gaussians, seeking to minimize using a least squares distance between the original non-resampled curve and the cubic interpolated curve obtained from the resampling.

We then perform a nonlinear regression for different resampling steps, taking as the original value $\sigma_f = 0.5p$. Indeed, this value seems to be a good intermediate value which makes it possible to keep interesting filter widths without losing too much information. The minimization here is only carried out on the distance applied to the correlations (and not on the segmentation error) because of the computation time. Figure 43 shows the curve comparing the error on the correlations.

We can see that we then obtain a lower error, and above all, the elbow of the curve begins at about $n = 9$ resampling values.

We can then look at the σ_f of the Gaussian curves obtained with 9 values. We notice that least squares minimization process consists in making the width of the Gaussians tend towards 0, so as to obtain thin bands. This confirms what we see in Figure 43: the optimized curve is in fact the error curve obtained for thin bands. Even so, adding limits for minimization is not better: results are simply fitting the lower bands of the limits. Relying on non-optimized spectra of a pre-established width therefore seems more efficient.

Conclusion and discussions

To sum up, the low spatial resolution and the long acquisition time due to the push-broom technology system suggest that a tunable filter-like camera would fit more to the Iznik ceramics study. However, this type of camera has to sacrifice the spectral resolution instead of the spatial one. That is why we observe how low could it be. The first results with a narrow bands resampling model on our data demonstrate that a resolution of only 9 values would be enough for fair results. But narrow bands filters would not be great for acquisition for a question of intensity received. It would increase the acquisition time, and that would streer against our requirements. In this case, a wide bands resampling model is more relatable and shows a minimum up to 18

values (i.e. 18 filters) required. We could finally try to adapt the width of each filter, but the results remains inconstant: a pre-established width is enough.

N.B.: These results also finally show that we are very far from 3 sufficient values. This confirms our affirmations on metamerism, and shows that an RGB acquisition of the tiles would clearly not suffice for their study.



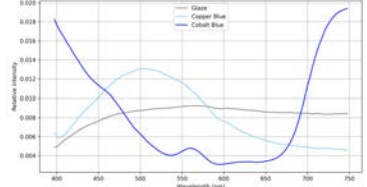

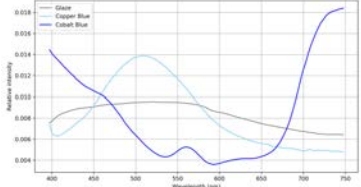

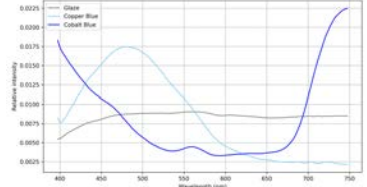

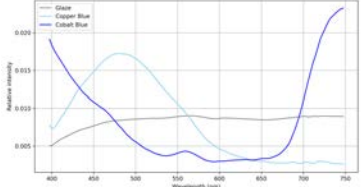

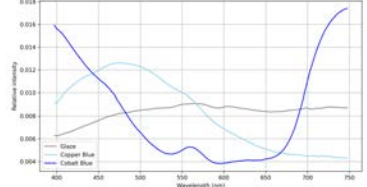

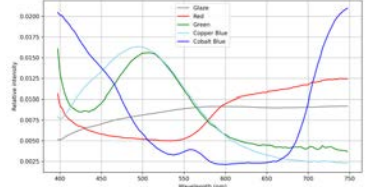

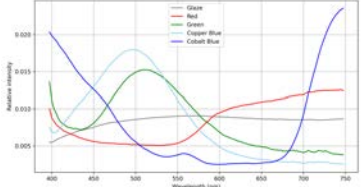

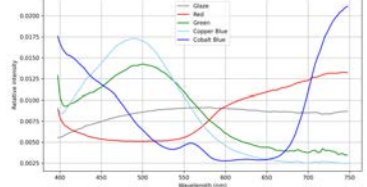

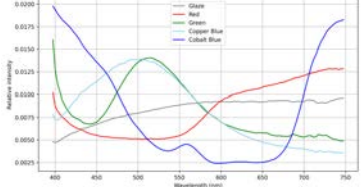

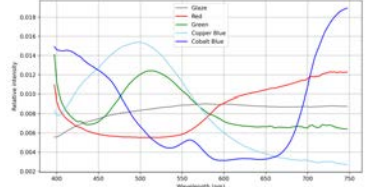

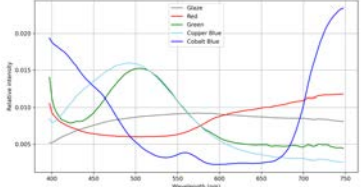



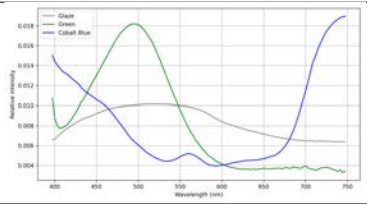

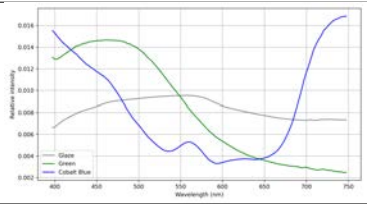

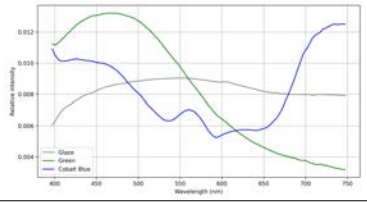
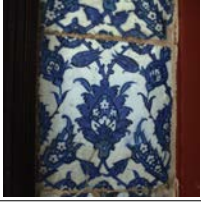
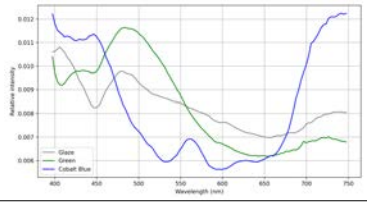

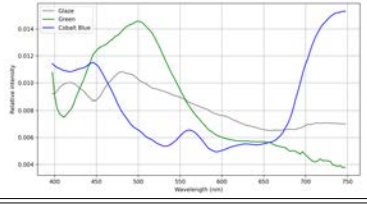

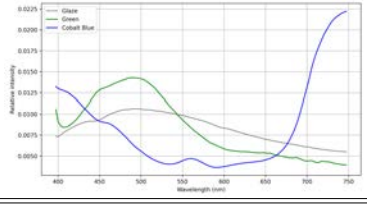
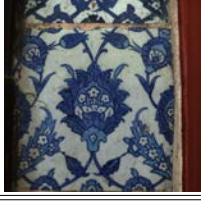
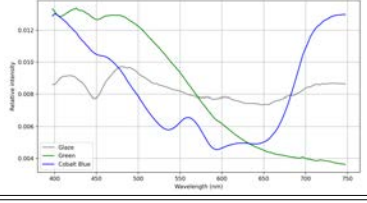

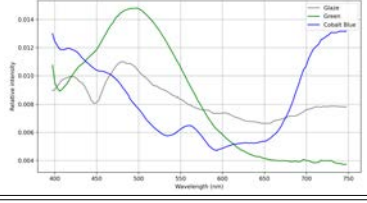
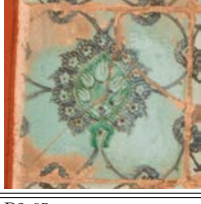
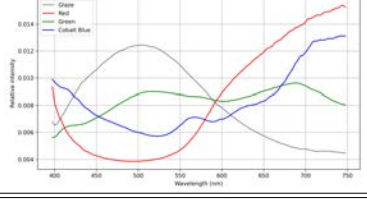

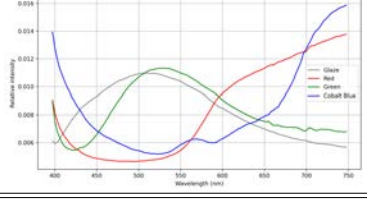

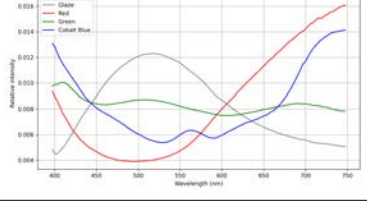

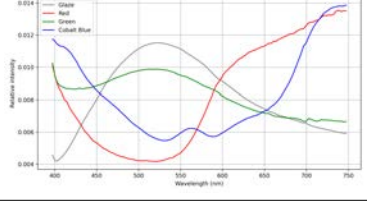

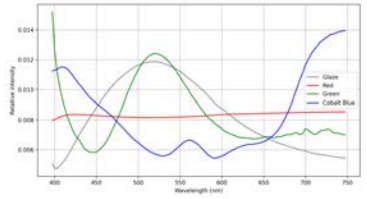

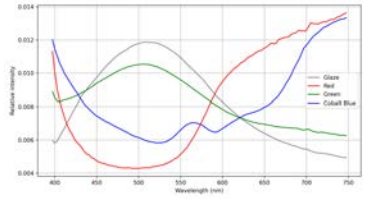
Figure 43: Plot of the distance of the correlations as a function of the number of n sampling values, before and after fitting


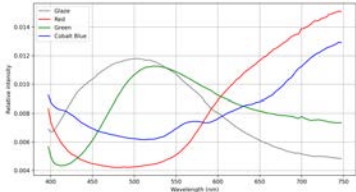

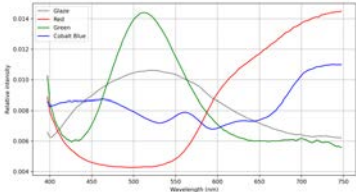

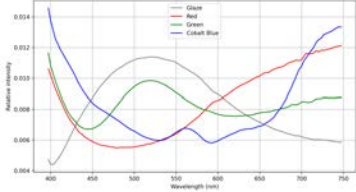

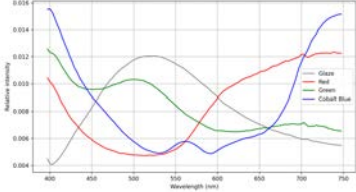

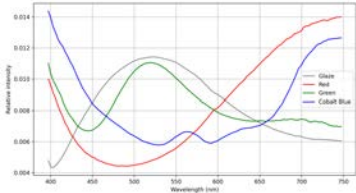

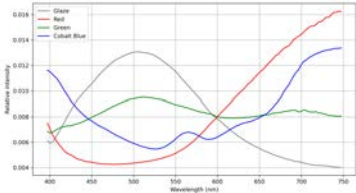

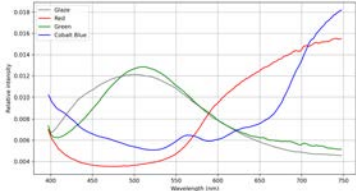
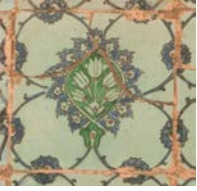
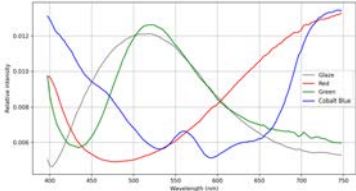

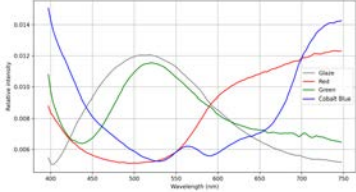

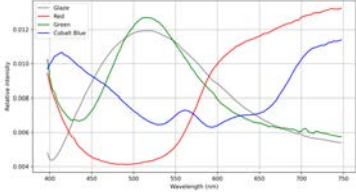

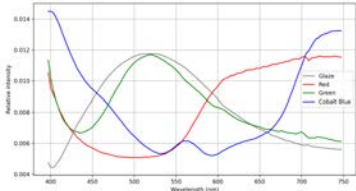

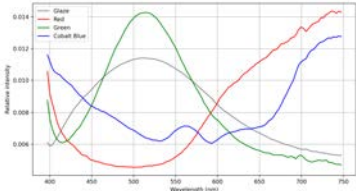

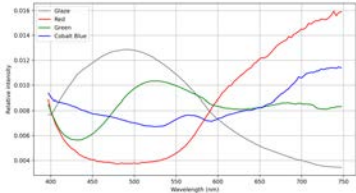

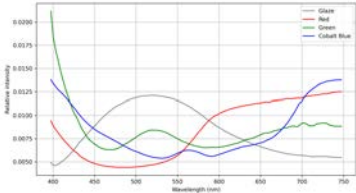
III Pigments studies of the Iznik ceramics tiles


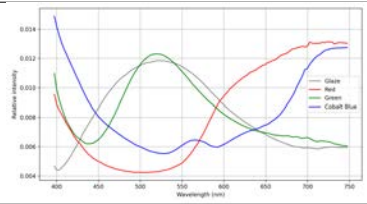

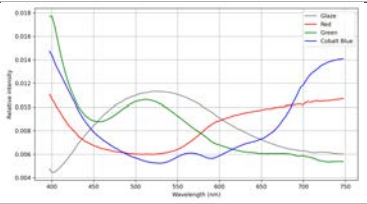

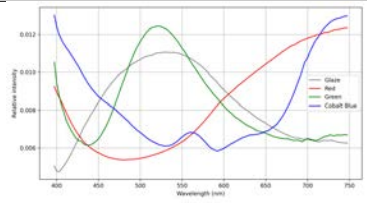
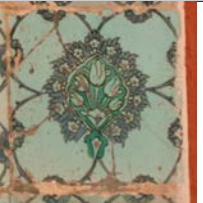
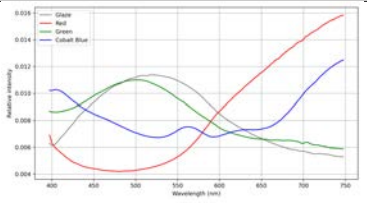

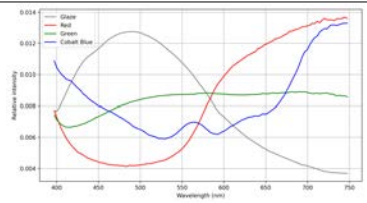

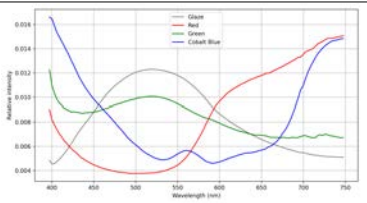

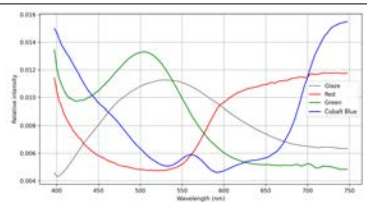
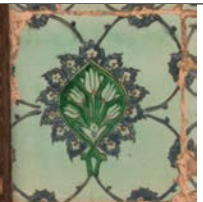
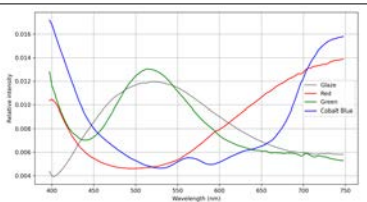

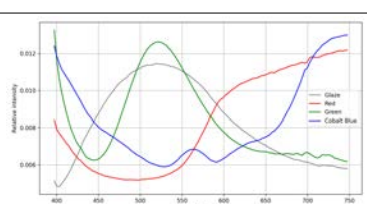

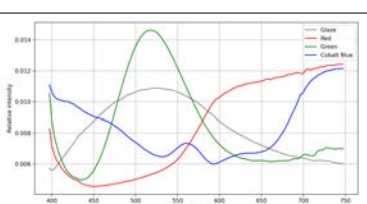

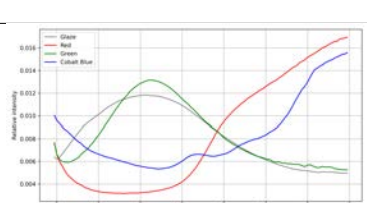

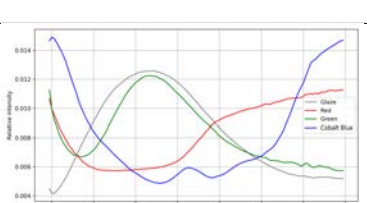

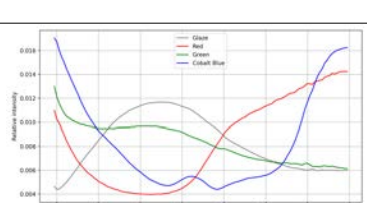

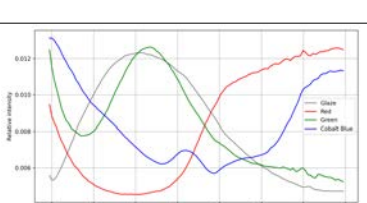
Tile presentation


Below are presented all the tiles acquired through the hyperspectral camera in Saint-Maurice residence. For each of the tile mentioned, we can find an RGB image on the left, and its representative spectra on the right.

RGB image	Representative spectra	RGB image	Representative spectra
			
			
			
			
			
			

RGB image	Representative spectra	RGB image	Representative spectra
			
			
			
			
			
			
			


RGB image	Representative spectra	RGB image	Representative spectra
P2-69 		P2-70 	
P2-71 		P2-72 	
P2-73 		P2-74 	
P2-75 		P2-76 	
P2-77 		P2-78 	
P2-79 		P2-80 	
P2-81 		P2-82 	

RGB image	Representative spectra	RGB image	Representative spectra
P2-83 		P2-84 	
P2-85 		P2-86 	
P2-87 		P2-88 	
P2-89 		P2-90 	
P2-91 		P2-92 	
P2-93 		P2-94 	
P2-95 		P2-96 	




Visualizations

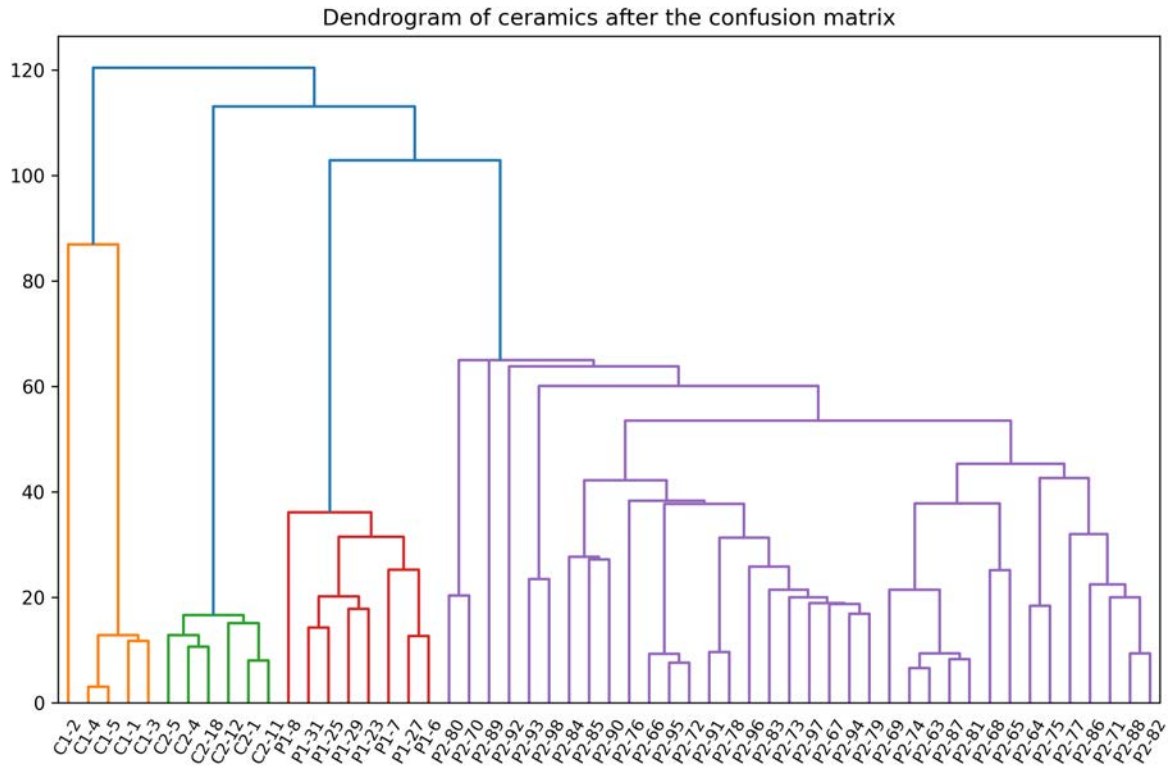
The following figures presents the visualizations of PCCR method including P2 tiles.




PCA graphs for all the Iznik tiles. Each of the group has its own color for better visualization.



Confusion matrix for all the Iznik tiles.



Dendrogram from the confusion matrix.



Dendrogram from agglomerative clustering.

IV Personal Blender/Python interface manual

Blender is a very powerful tool for 3D object manipulation. There are many aspects which can be used for our SVBRD acquisition process, especially for the visualization of a scene: rendering, placement of lights, cameras, etc. Moreover, its free open source aspect makes it great for researches, but also for sharing it.

In our case, *Blender* can be very convenient. First, we find its interest because it allows us to create our set-ups on entirely fictitious models, before using it on real hardware. So it can be very useful for our acquisition tests and validation. Indeed we manipulate it for example for the placement of lights, but also for the validation of our fitting methods on the BRDFs. In this sway, it is therefore important to find how to export cameras, objects, from *Blender* to *OpenCV*.

Similarly, *Blender* is also used to visualize acquisition results. For example, after having placed them in our scene with *OpenCV*, we can place cameras in Blender to see if it corresponds well. To do this, we therefore must be able to export a camera model from *OpenCV* to *Blender*.

However, we see that *Blender* 3D conventions may not match with *OpenCV*. So there are some axes on which we looked to ensure a good link between *Blender* and *OpenCV* on Python:

- Importing and exporting objects in *.obj* on *Blender*,
- Getting a *Blender* camera model to *OpenCV* and vice versa,
- Exporting our results from the SVBRDF acquisition to a *Blender* visualization.

Import and export *.obj* files on *Blender*

It is possible to import or export wavefront *.obj* object on Blender. However, we notice that they do not fit well in the 3D reference. When importing, we must specify:

Transform >

- *Forward > Y Forward*
- *Up > Z Up*

Likewise, when exporting a **Blender** scene, it is important to respect the same order, otherwise the 3D mesh will not have the desired meaning and will be unusable on OpenCV. Moreover, we should select the 3D object to export, and then to check the "Selected only" box.

Export intrinsic camera from Blender to *OpenCV* and back

The *OpenCV* and *Blender* cameras model used are both a pinhole camera model, but they present great differences, whether by the parameters or the references used. First, we will see the different parameters on each side, then how to switch from one to the other.

Parameters on *OpenCV*. With *OpenCV*, the parameters are as follows (we will use the same names as used as input or output in the *OpenCV* documentation):

- The matrix of intrinsic parameters *mtx*
- The rotation vector *rvec*, which is the same type of *rotvec* defined by *Scipy*
- The position vector *pos* (here you can either use the translation vector *tvec* or the position vector. In our case, we prefer to speak of position vector *pos* which is more intuitive).
- The *dist* distortion vector
- The size of our image (*h, w*)

We can represent the intrinsic matrix *mtx* by:

$$mtx = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \quad (31)$$

where f_x and f_y are the focal lengths along each ax of our camera model, and c_x and c_y the optical centers.

Parameters on Blender. Settings vary a bit on *Blender*. They are :

- The xyz position of the camera
- Euler angles orientation rot
- The focal length of the lens (in mm)
- The shifts $shift_x$ and $shift_y$
- The $pixe_aspect_x$ (which remains at 1) and $pixel_aspect_y$

Blender to OpenCV. First, the simple aspect is that the position of the camera remains the same on each side. Then, the rotation is not simple because he two camera models does not follow the same reference (see Figure 44). In order to solve this problem, we use the Python code 1. Finally, we apply the code presented in [3] on the *Blender* scripting interface to obtain the matrix mtx .

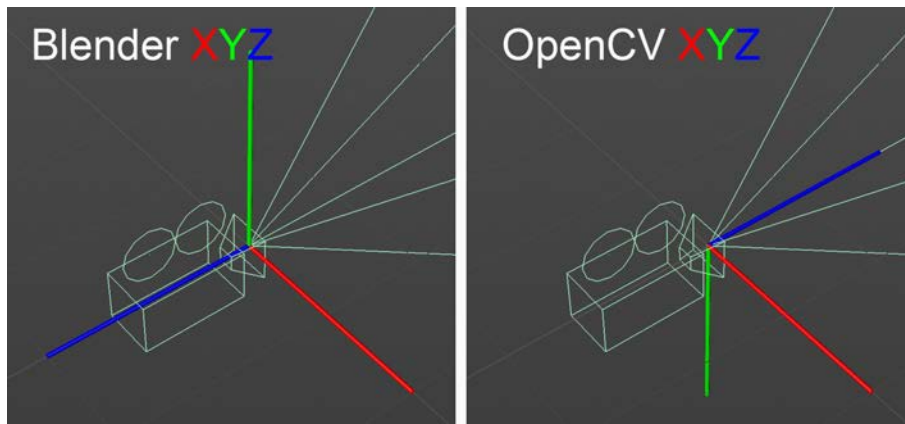


Figure 44: Orientation of the camera reference in each case.

Algorithm 2 Rotation conversion from *Blender* to *OpenCV*

```
list_ = ...                                ▷ Euler angles orientation vector from Blender
rot = np.reshape(list_, 3).astype(np.float)
R1 = Rot.from_euler('xyz',rot).as_matrix()
blender2opencv = np.array([0, 0, 90])*np.pi/180
B2CV = Rot.from_euler('xyz',blender2opencv).as_matrix()
rvec = -Rot.from_matrix(np.dot(R1, B2CV)).as_rotvec()
```

Be careful: in *Blender*, and in *OpenCV*, f_x and f_y are reversed (c_x and c_y accordingly).

OpenCV to Blender. Similarly, from *OpenCV* to *Blender*, we keep the position but we do the opposite path on the rotation with the code 2.

Then, the conversions from *OpenCV* to *Blender* of shifts $shift_x$ and $shift_y$, $pixe_aspect_x$ and $pixe_aspect_y$ are also described in [3]. For the focal length, we prefer to use the length given in the EXIF of the real photo. Finally, we assume that the distortions are not taken into account: they are zero on *Blender*.

Algorithm 3 Rotation conversion from *OpenCV* to *Blender*

```
R1 = Rot.from_rotvec(-rvec/180*np.pi).as_matrix()
blender2opencv = np.array([0, 0, 90])*np.pi/180
B2CV = Rot.from_euler('xyz',blender2opencv).as_matrix()
rot = Rot.from_matrix(np.dot(R1, np.linalg.inv(B2CV))).as_euler('xyz')
```

Using *Blender* for P-n-P Pose computation with *OpenCV*

Similarly, one can use *Blender* to help with the PnP pose computation by placing 3D points on the mesh, and exporting them to OpenCV. For this, we set up an easy and fast method to place points by hand on the object, on click. While remaining in *Layout* mode, we click where we want to place our point, by doing *Shift+Right click*, or by selecting the cursor in the menu on the left. Then we click on the zone to be selected, and we click at the top on *Add>Empty>Plain Axes*.

We then have a positioned point called *Empty*. We repeat for each point, then we run the code 3 in *Blender* script reader.

Algorithm 4 3D points export

```
import bpy
from mathutils import Matrix, Vector
import numpy as np

function EXPORT_POSITION(cam_name)
    list_pos = []
    for ob in bpy.context.selected_objects do
        list_pos = []
        list_pos.append(ob.location[:])
    end for
    nP = np.matrix(list_pos)
    path = bpy.path.abspath("//")
    filename = cam_name + ".txt"
    file = path + filename
    np.savetxt(file, nP)
end function

function MAIN
    export_position('Collection')
end function

if __name__ == "__main__" then
    main()
end if
```

We then obtain a *.txt* file, which contains all the 3D points that we can use on Python for the PnP Pose estimation.

Exporting SVBRDF parameters results to *Blender*

Here we describe how we export the results obtained in Chapter 3 on Blender in order to get a 3D object linked to an SVBRDF. For this we need to follow two steps. First, we have to rewrap a UV mapping in order to project the maps we made in Chapter 3 on the mesh from the poses




Figure 45: *Top:* Screenshot of Blender during UV mapping edition. *Bottom:* Screenshot of Blender when adding a new shader using OSL.

of the cameras. Next, we create a shader that allows us to use its maps to render the object correctly.

Creating new UV coordinates with projecting mapping

Usually, we use a mesh which is textured. It means that a UV mapping has already been done. The problem we face is that we want to project our new UV maps containing our SBRDF parameters from the camera positions. This UV mapping is called projecting mapping.

To convert our UV mapping into a projection mapping, we place ourselves in the *UV Editor* workplace. Two windows open: the *UV Editor* and the *3D Viewport*. In the *3D viewport*, we switch to *Edit Mode*, we select the 3D object and we indicate *UV>Project from View*. Next, we can import a photo from the camera pose into the UV Editor to see if it works well.

Adding a new shader with *Open Shading Language* and nodes

To apply our BRDF model on *Blender*, we want to be able to import our own BRDF model using shaders. For this, we use *Open Shader Language*. This is a language close to C++ that we can manipulate in the *Text Editor*. We start by authorizing it by choosing *Render Engine > Cycles* in *Render Properties*, and then checking *Open Shading Language*. So, we open a window with the *Text Editor* and one with the *Shader Editor*. We then write our shader in the Text Editor, which we save in *.osl* format. In the *Shader editor*, we can yet add a node *Add > Script* where we choose our shader. It is afterward connected to other nodes, such as the colors, the normal, etc.

V List of symbols in Chapter 3

Table 6: *List of symbols.*

Symbol	Description
ρ	BRDF function
ω_i	Light angular vector
ω_o	View angular vector
$(\theta_o, \phi_o), (\theta_i, \phi_i)$	Polar and azimuthal angles
\mathcal{F}_o	Flux received by the camera
\mathcal{F}_i	Incident flux
Ω_{ω_i}	Solide angle from the light source around ω_i
β	Material indentation
(u, v)	Studied pixel on the image
\mathbf{x}	Studied point on the mesh
n	Studied normal in \mathbf{x} defined by (θ, ϕ)
$\delta_{rad, sph}$	Radius of the mirror sphere
$q_{sph, l}, q_{sph, r}$	Position vector of the left (resp. right) mirror sphere
$\mathcal{C}_l, \mathcal{C}_s$	Camera light (resp. scene)
q_l, q_s	Position vector of the camera light (resp. scene)
t_l, t_s	Translation vector of the camera light (resp. scene)
r_l, r_s	Rotation vector of the camera light (resp. scene)
R_l, R_s	Rotation matrix of the camera light (resp. scene)
A_l, A_s	Intrinsic matrix of the camera light (resp. scene)
f_x, f_y	Focal distances of the camera
c_x, c_y	Central point coordinates of the camera
δ_s	Distance between \mathbf{x} and the camera scene
m_s	Number of camera scene positions
\mathcal{L}	Lightspot
$q_{\mathcal{L}}$	Lightspot position vector
$\delta_{\mathcal{L}}$	Distance between \mathbf{x} and the lightspot
$\mathcal{D}_{\mathcal{L}}$	Lightspot angular vector
$m_{\mathcal{L}}$	Number of light positions
\mathcal{I}_s	HDR images taken by \mathcal{C}_s
M_{norm}	Normal map
M_{norm}	Segmentation map
M_{app}	Material map (rendered image)
M_{error}	Error map between M_{app} and \mathcal{I}_s
\mathcal{I}_{ref}	Reference image for disc source estimation
\mathcal{I}_{scene}	Synthetic scene image for disc source estimation
α	Roughness parameters for GGX
η	Fresnel refractive index
k_d	RGB triplets of the lambertian part

Bibliography

- [1] Archéogrid, Le CAIRE Villa St-Maurice. <https://www.archeogrid.fr/project/5369>.
- [2] ENVI-SHW software. [https://www.nv5geospatialsoftware.com/docs/SpectralHou
rglassWizard.html#1](https://www.nv5geospatialsoftware.com/docs/SpectralHou
rglassWizard.html#1).
- [3] From Blender to OpenCV Camera and back. [https://www.rojtberg.net/1601/from-b
lender-to-opencv-camera-and-back/](https://www.rojtberg.net/1601/from-b
lender-to-opencv-camera-and-back/).
- [4] Hugin sourceforge. <https://hugin.sourceforge.io/releases/2022.0.0/en.shtml>.
- [5] OpenCV Documentation, Camera Calibration. [https://docs.opencv.org/4.x/dc/dbb
/tutorial_py_calibration.html](https://docs.opencv.org/4.x/dc/dbb
/tutorial_py_calibration.html).
- [6] OpenCV Documentation, Camera Calibration and 3D Reconstruction. [https://docs.o
pencv.org/3.4/d9/d0c/group__calib3d.html](https://docs.o
pencv.org/3.4/d9/d0c/group__calib3d.html).
- [7] OpenCV Documentation, Perspective-n-Point (PnP) pose computation. [https://docs.o
pencv.org/4.x/d5/d1f/calib3d_solvePnP.html](https://docs.o
pencv.org/4.x/d5/d1f/calib3d_solvePnP.html).
- [8] The Thermal radiation, Factors From Differential Elements to Finite Areas. [http://ww
w.thermalradiation.net/sectionb/B-17.html](http://ww
w.thermalradiation.net/sectionb/B-17.html).
- [9] Miika Aittala, Tim Weyrich, Jaakko Lehtinen, et al. Two-shot SVBRDF capture for stationary materials. *ACM Trans. Graph.*, 34(4):110–1, 2015.
- [10] Arash Akbarinia and Karl R Gegenfurtner. Color metamerism and the structure of illuminant space. *JOSA A*, 35(4):B231–B238, 2018.
- [11] Hanaa A Al-Gaoudi and Roberta Iannaccone. Multiband imaging techniques incorporated into the study of dyed ancient Egyptian textile fragments. *International Journal of Conservation Science*, 12(3):893–906, 2021.
- [12] Giacomo Aletti, Alessandro Benfenati, and Giovanni Naldi. A semi-supervised reduced-space method for hyperspectral imaging segmentation. *Journal of Imaging*, 7(12):267, 2021.
- [13] Neil Alldrin, Todd Zickler, and David Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [14] Ayed Ben Amara, Maïa Cuin, Max Schvoerer, and Paolo Peduto. Voyage en majolique-Texture, composition et couleurs d’une céramique glaçurée trouvée à Ravello (XVIe siècle, Italie), 2005.
- [15] Ayed Ben Amara, Max Schvoerer, Maïa Cuin, and Mohamed Baji. An investigation of the ceramic technology of a late Iznik ceramic production (XVIIth century AD). In *34th international symposium on Archaeometry*, pages 393–398. Institución Fernando el Católico, 2006.

- [16] Nurhan Atasoy, Julian Raby, and Yanni Petsopoulos. *Iznik: the pottery of Ottoman Turkey*. (No Title), 1989.
- [17] Nebahat Avcioğlu and Mercedes Volait. “Jeux de miroir”: Architecture of Istanbul and Cairo from Empire to Modernism. In *A Companion to Islamic Art and Architecture*, chapter 43, pages 1122–1149. John Wiley & Sons, Ltd, 2017.
- [18] Seung-Hwan Baek, Tizian Zeltner, Hyunjin Ku, Inseung Hwang, Xin Tong, Wenzel Jakob, and Min H Kim. Image-based acquisition and modeling of polarimetric reflectance. *ACM Trans. Graph.*, 39(4):139, 2020.
- [19] Vincent Baillet, Pascal Mora, Corentin Cou, Sarah Tournon-Valiente, Mercedes Volait, Xavier Granier, Romain Pacanowski, and Gaël Guennebaud. 3D for Studying Reuse in 19th Century Cairo: the Case of Saint-Maurice Residence. In *GCH 2021-Eurographics Workshop on Graphics and Cultural Heritage*, 2021.
- [20] S. Gonizzi Barsanti, F. Remondino, B. Jiménez Fenández-Palacios, and D. Visintini. Critical Factors and Guidelines for 3D Surveying and Modelling in Cultural Heritage. *Int. J. of Heritage in the Digital Era*, 3(1):141–158, 2014.
- [21] Delphine Bauer, Ayed Ben Amara, and Nadia Cantin. La céramique de poêle du site Berg Armo (XVI s., Sainte-Marie-aux-Mines, France): matériaux et techniques de fabrication. *ArcheoSciences*, 43:275–286, 2019.
- [22] Emmie Beauvoit, Ayed Ben Amara, Nadia Cantin, Quentin Lemasson, Christophe Sireix, Valérie Marache, and Rémy Chapoulie. Technological investigation on ceramic bodies of 19th century French white earthenware from the Bordeaux region. *Journal of Archaeological Science: Reports*, 31:102314, 2020.
- [23] Emmie Beauvoit, Ayed Ben Amara, Nicolas Tessier-Doyen, Camille Frugier, Quentin Lemasson, Brice Moignard, Claire Pacheco, Laurent Pichon, Rémy Chapoulie, and Bernard Gratuze. Chemical and Mechanical Characterisation of White Earthenware Glazes from the Johnston-Vieillard Manufactory (France, 19th Century). *Archaeometry*, 63(5):941–959, 2021.
- [24] Petr Beckmann and Andre Spizzichino. *The scattering of electromagnetic waves from rough surfaces*. Norwood, 1987.
- [25] Jan Behmann, Kelvin Acebron, Dzhaner Emin, Simon Bennertz, Shizue Matsubara, Stefan Thomas, David Bohnenkamp, Matheus T Kuska, Jouni Jussila, Harri Salo, et al. Specim IQ: Evaluation of a new, miniaturized handheld hyperspectral camera and its application for plant phenotyping and disease detection. *Sensors*, 18(2):441, 2018.
- [26] Ayed Ben Amara. Céramiques glaçurées de l’espace méditerranéen (IXe-XVIIe siècles ap. J.-C.): matériaux, techniques et altération. 2002.
- [27] Ayed Ben Amara, Max Schvoerer, Abdelaziz Daoulatli, and Mourad Rammah. “Jaune de Raqqada” et autres couleurs de céramiques glaçurées Aghlabides de Tunisie (IXe-Xe siècles). *ArchéoSciences, revue d’Archéométrie*, 25(1):179–186, 2001.
- [28] Moshe Ben-Ezra, Jiaping Wang, Bennett Wilburn, Xiaoyang Li, and Le Ma. An LED-only BRDF measurement device. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [29] Pascal-François Bertrand. *Aubusson, tapisseries des Lumières: splendeurs de la Manufacture royale, fournisseur de l’Europe au XVIIIe siècle*. 2013.

- [30] William S Bickel and Wilbur M Bailey. Stokes vectors, Mueller matrices, and polarized scattered light. American Journal of Physics, 53(5):468–478, 1985.
- [31] Bastien Billiot, Frédéric Cointault, Ludovic Journaux, Jean-Claude Simon, and Pierre Gouton. 3D image acquisition system based on shape from focus technique. Sensors, 13(4):5040–5053, 2013.
- [32] Francois Blais. Review of 20 years of range sensor development. Journal of electronic imaging, 13(1):231–243, 2004.
- [33] Ilya Blayvas, Ron Kimmel, and Ehud Rivlin. Role of optics in the accuracy of depth-from-defocus systems. JOSA A, 24(4):967–972, 2007.
- [34] Wolfgang Boehler, Guido Heinz, and Andreas Marbs. The potential of non-contact close range laser scanners for cultural heritage recording. International archives of photogrammetry remote sensing and spatial information sciences, 34(5/C7):430–436, 2002.
- [35] Nirmal K Bose and Nilesh A Ahuja. Superresolution and noise filtering using moving least squares. IEEE Transactions on Image Processing, 15(8):2239–2248, 2006.
- [36] J-Y Bouguet. Camera calibration toolbox for matlab. http://www.vision.caltech.edu/bouguetj/calib_doc/index.html, 2004.
- [37] Brent Burley and Walt Disney Animation Studios. Physically-based shading at disney. In Acm Siggraph, volume 2012, pages 1–7. vol. 2012, 2012.
- [38] Willy Anugrah Cahyadi, Yeon Ho Chung, Zabih Ghassemlooy, and Navid Bani Hassan. Optical camera communications: Principles, modulations, potential and challenges. Electronics, 9(9):1339, 2020.
- [39] Tadeusz Caliński and Jerzy Harabasz. A dendrite method for cluster analysis. Communications in Statistics-theory and Methods, 3(1):1–27, 1974.
- [40] Stefano Campana and Maurizio Forte. From space to place. BAR International Series, 1568, 2006.
- [41] Jeanne Capelle. Reflectance Transformation Imaging (RTI) et épigraphie. 2017.
- [42] Guojun Chen, Yue Dong, Pieter Peers, Jiawan Zhang, and Xin Tong. Reflectance scanning: Estimating shading frame and BRDF with generalized linear light sources. ACM Transactions on Graphics (TOG), 33(4):1–11, 2014.
- [43] Catia Clementi, Gloria Basconi, Roberto Pellegrino, and Aldo Romani. Carthamus tinctorius L.: A photophysical study of the main coloured species for artwork diagnostic purposes. Dyes and Pigments, 103:127–137, 2014.
- [44] Michael F Cohen and John R Wallace. Radiosity and realistic image synthesis. Morgan Kaufmann, 1993.
- [45] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. ACM Transactions on Graphics (ToG), 1(1):7–24, 1982.
- [46] Massimiliano Corsini, Marco Callieri, and Paolo Cignoni. Stereo light probe. Computer Graphics Forum, 27(2):291–300, 2008.
- [47] Antonino Cosentino. FORS spectral database of historical pigments in different binders. E Conserv. J, 2:54–65, 2014.

- [48] Corentin Cou, Ayed Ben Amara, and Xavier Granier. Non-invasive on-site method for thickness measurement of transparent ceramic glazes based on depth from focus. *Archaeometry*, 2022.
- [49] Corentin Cou and Gaël Guennebaud. Depth from Focus using Windowed Linear Least Squares Regressions. *The Visual Computer*, pages 1–10, 2023.
- [50] Costanza Cucci, John K Delaney, and Marcello Picollo. Reflectance hyperspectral imaging for investigation of works of art: old master paintings and illuminated manuscripts. *Accounts of chemical research*, 49(10):2070–2079, 2016.
- [51] Costanza Cucci, Marcello Picollo, Leandro Chiarantini, Gianni Uda, Lorenzo Fiori, Bruno De Nigris, and Massimo Osanna. Remote-sensing hyperspectral imaging for applications in archaeological areas: Non-invasive investigations on wall paintings and on mural inscriptions in the Pompeii site. *Microchemical Journal*, 158:105082, 2020.
- [52] A Daoulatli. La céramique ifriqiyenne du IX^e au XV^e siècle, La céramique médiévale en Méditerranée Occidentale, Col-loques Internationaux CNRS, N 584, Valbonne, 11-14 septembre 1978, éd, 1980.
- [53] David L Davies and Donald W Bouldin. A cluster separation measure. *IEEE transactions on pattern analysis and machine intelligence*, (2):224–227, 1979.
- [54] Hortense de la Codre. Textiles et colorants des tapisseries : développement d’une méthodologie d’analyses sans-contact. Le cas des tapisseries fines dites ”Verdures” d’Aubusson. 2023.
- [55] Hortense De La Codre, Floréal Daniel, Rémy Chapoulie, Laurent Servant, and Aurélie Mounier. Investigating the materials used in eighteenth-century tapestries from the three French Royal Manufactories: inputs of hyperspectral approaches. *The European Physical Journal Plus*, 136(11):1193, 2021.
- [56] Hortense de La Codre, Charlotte Marembert, Pauline Claisse, Floréal Daniel, Rémy Chapoulie, Laurent Servant, and Aurélie Mounier. Non-invasive characterization of yellow dyes in tapestries of the 18th century: Influence of composition on degradation. *Color Research & Application*, 46(3):613–622, 2021.
- [57] John K Delaney, Kathryn A Dooley, Annelies Van Loon, and Abbie Vandivere. Mapping the pigment distribution of Vermeer’s Girl with a Pearl Earring. *Heritage Science*, 8:1–16, 2020.
- [58] John K Delaney, Paola Ricciardi, Lisha Glinsman, Michael Palmer, and Julia Burke. Use of near infrared reflectance imaging spectroscopy to map wool and silk fibres in historic tapestries. *Analytical Methods*, 8(44):7886–7890, 2016.
- [59] Francesco Dell’Endice, Jens Nieke, Benjamin Koetz, Michael E Schaepman, and Klaus Itten. Improving radiometry of imaging spectrometers by using programmable spectral regions of interest. *ISPRS Journal of Photogrammetry and Remote Sensing*, 64(6):632–639, 2009.
- [60] Xi Deng, Fujun Luan, Bruce Walter, Kavita Bala, and Steve Marschner. Reconstructing Translucent Objects using Differentiable Rendering. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022.
- [61] Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis, and Adrien Bousseau. Flexible SVBRDF capture with a multi-image deep network. In *Computer graphics forum*, volume 38, pages 1–13. Wiley Online Library, 2019.

- [62] Yue Dong. Deep appearance modeling: A survey. Visual Informatics, 3(2):59, 2019.
- [63] Yue Dong, Jiaping Wang, Xin Tong, John Snyder, Yanxiang Lan, Moshe Ben-Ezra, and Baining Guo. Manifold bootstrapping for SVBRDF capture. ACM Transactions on Graphics (TOG), 29(4):1–10, 2010.
- [64] Ethan Eade. Gauss-newton/levenberg-marquardt optimization. Tech. Rep., 2013.
- [65] Lennart Eriksson, Tamara Byrne, E Johansson, Johan Trygg, and C Vikström. Multi-and megavariate data analysis basic principles and applications, volume 1. Umetrics Academy, 2013.
- [66] M. Farjas, J. Gregorio Rejas, T. Mostaza, and J. Zancajo. Deepening in the 3D Modelling: Multisource Analysis of a Polychrome Ceramic Vessel Through the Integration of Thermal and Hyperspectral Information. In Revive the Past: CAA 2012, pages 116–124.
- [67] Christian Fischer and Ioanna Kakoulli. Multispectral and hyperspectral imaging technologies in conservation: current research and potential applications. Studies in Conservation, 51(sup1):3–16, 2006.
- [68] LD Fiske, Aggelos K Katsaggelos, MCG Aalders, M Alfeld, M Walton, and Oliver Cossairt. A data fusion method for the delayering of x-ray fluorescence images of painted works of art. In 2021 IEEE International Conference on Image Processing (ICIP), pages 3458–3462. IEEE, 2021.
- [69] Finbarr B. Flood and Gülru Necipoğlu. Modernity, Empire, Colony, and Nation (1700–1950). In A Companion to Islamic Art and Architecture, volume 2, chapter 7, pages 1051–1053. John Wiley & Sons, Ltd, 2017.
- [70] Sing Choong Foo. A gonireflectometer for measuring the bidirectional reflectance of material for use in illumination computation. 1997.
- [71] Jeffrey Froh. Archaeological ceramics studied by scanning electron microscopy. Hyperfine Interactions, 154:159–176, 2004.
- [72] F Gabrieli, KA Dooley, M Facini, and JK Delaney. Near-UV to mid-IR reflectance imaging spectroscopy of paintings on the macroscale. Science advances, 5(8):eaaw7794, 2019.
- [73] Mario Ramírez Galán and Ronda Sandifer Bard. Studies in Archaeometry: Proceedings of the Archaeometry Symposium at NORM 2019, June 16–19, Portland, Oregon, Portland State University. Dedicated to the Rev. H. Richard Rutherford, CSC, Ph. D. 2020.
- [74] Guillermo Gallego and Anthony Yezzi. A compact formula for the derivative of a 3-D rotation in exponential coordinates. Journal of Mathematical Imaging and Vision, 51:378–384, 2015.
- [75] Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. Linear light source reflectometry. ACM Transactions on Graphics (TOG), 22(3):749–758, 2003.
- [76] Nahum Gat. Imaging spectroscopy using tunable filters: a review. Wavelet Applications VII, 4056:50–64, 2000.
- [77] Paul Geladi and Hans F Grahn. Multivariate image analysis. Encyclopedia of Analytical Chemistry: Applications, Theory and Instrumentation, 2006.
- [78] Abhijeet Ghosh, Wolfgang Heidrich, Shruthi Achutha, and Matthew O’Toole. A basis illumination approach to BRDF measurement. International journal of computer vision, 90:183–197, 2010.

- [79] Kremer Pigmente GmbH and Co. KG. 990001 – 990018 color charts <http://www.kremerpigmente.com/en/>. 2013.
- [80] Alexander FH Goetz, Gregg Vane, Jerry E Solomon, and Barrett N Rock. Imaging spectrometry for earth remote sensing. *science*, 228(4704):1147–1153, 1985.
- [81] Richard M Goody and Yuk Ling Yung. *Atmospheric radiation: theoretical basis*. Oxford university press, 1995.
- [82] Sai Siva Gorthi and Pramod Rastogi. Fringe projection techniques: whither we are? *Optics and lasers in engineering*, 48(ARTICLE):133–140, 2010.
- [83] Aoife A Gowen, Colm P O’Donnell, Patrick J Cullen, Gérard Downey, and Jesus M Frias. Hyperspectral imaging—an emerging process analytical tool for food quality and safety control. *Trends in food science & technology*, 18(12):590–598, 2007.
- [84] Leo Grady. Random walks for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 28(11):1768–1783, 2006.
- [85] Hans Grahm and Paul Geladi. *Techniques and applications of hyperspectral image analysis*. John Wiley & Sons, 2007.
- [86] Xavier Granier, Laurent Bergerot, Mehdi Chayani, Bruno Dutailly, Pascal Mora, Jean-Louis Kerouanton, François Daniel, Jean-Baptiste Barreau, Jean-François Bernard, Hervé Bohbot, Philippe Fleury, Olivier Marlet, and Dickinson Richard. 3D lexicon for Human and Social Sciences, April 2021.
- [87] Emanuela Grifoni, Beatrice Campanella, Stefano Legnaioli, Giulia Lorenzetti, Luciano Marras, Stefano Pagnotta, Vincenzo Palleschi, Francesco Poggialini, Emanuele Salerno, and Anna Tonazzini. A new Infrared True-Color approach for visible-infrared multispectral image analysis. *Journal on Computing and Cultural Heritage (JOCCH)*, 12(2):1–11, 2019.
- [88] Paul Grossmann. Depth from focus. *Pattern recognition letters*, 5(1):63–69, 1987.
- [89] D. Guarnera, G.C. Guarnera, A. Ghosh, C. Denk, and M. Glencross. BRDF Representation and Acquisition. *Computer Graphics Forum*, 35(2):625–650, 2016.
- [90] Darya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. BRDF representation and acquisition. In *Computer Graphics Forum*, volume 35, pages 625–650. Wiley Online Library, 2016.
- [91] Gabriele Guidi, Grazia Tucci, Jean-Angelo Beraldin, S Ciofi, D Ostuni, F Costantini, and S El-Hakim. Multiscale archaeological survey based on the integration of 3D scanning and photogrammetry. In *International Workshop on Scanning for Cultural Heritage Recording (CIPA WG/6)*, pages 58–64, 2002.
- [92] Mohit Gupta, Amit Agrawal, Ashok Veeraraghavan, and Srinivasa G Narasimhan. Structured light 3D scanning in the presence of global illumination. In *CVPR 2011*, pages 713–720. IEEE, 2011.
- [93] Nathan Hagen and Michael W Kudenov. Review of snapshot spectral imaging technologies. *Optical Engineering*, 52(9):090901–090901, 2013.
- [94] Tom SF Haines. *Integrating Shape-from-Shading & Stereopsis*. PhD thesis, University of York, 2009.
- [95] M Halioua, RS Krishnamurthy, H Liu, and FP Chiang. Projection moiré with moving gratings for automated 3-D topography. *Applied Optics*, 22(6):850–855, 1983.

- [96] Sara J Hamilton and Robert A Lodder. Hyperspectral imaging technology for pharmaceutical analysis. In Biomedical Nanotechnology Architectures and Applications, volume 4626, pages 136–147. SPIE, 2002.
- [97] Jefferson Y Han and Ken Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. In ACM SIGGRAPH 2003 Papers, pages 741–748. 2003.
- [98] Yanling Han, Yi Gao, Yun Zhang, Jing Wang, and Shuhu Yang. Hyperspectral sea ice image classification based on the spectral-spatial-joint feature with deep learning. Remote Sensing, 11(18):2170, 2019.
- [99] Caner Hazirbas, Sebastian Georg Soyer, Maximilian Christian Staab, Laura Leal-Taixé, and Daniel Cremers. Deep depth from focus. In Asian Conference on Computer Vision, pages 525–541. Springer, 2018.
- [100] KGT Hollands. On the superposition rule for configuration factors. 1995.
- [101] S.A. Hovanessian. Introduction to Sensor Systems. Artech House Communication and. Artech House, 1988.
- [102] Jack J Hsia and Joseph C Richmond. Bidirectional Reflectometry. Part I: A High Resolution Laser Bidirectional Reflectometer With Results on Several Optical Coatings. Journal of Research of the National Bureau of Standards. Section A, Physics and Chemistry, 80(2):189, 1976.
- [103] Yan Hu, Qian Chen, Shijie Feng, Tianyang Tao, Hui Li, and Chao Zuo. Real-time microscopic 3D shape measurement based on optimized pulse-width-modulation binary fringe projection. Measurement Science and Technology, 28(7):075010, 2017.
- [104] Zhuo Hui, Kalyan Sunkavalli, Joon-Young Lee, Sunil Hadap, Jian Wang, and Aswin C Sankaranarayanan. Reflectance capture using univariate sampling of BRDFs. In Proceedings of the IEEE International Conference on Computer Vision, pages 5362–5370, 2017.
- [105] Matthias B Hullin, Johannes Hanika, Boris Ajdin, Hans-Peter Seidel, Jan Kautz, and Hendrik PA Lensch. Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. In ACM SIGGRAPH 2010 papers, pages 1–7. 2010.
- [106] Robert William Gainer Hunt and Michael R Pointer. Measuring colour. John Wiley & Sons, 2011.
- [107] Inseung Hwang, Daniel S Jeon, Adolfo Munoz, Diego Gutierrez, Xin Tong, and Min H Kim. Sparse ellipsometry: portable acquisition of polarimetric SVBRDF and shape with unstructured flash photography. ACM Transactions on Graphics (TOG), 41(4):1–14, 2022.
- [108] Hossein Javidnia and Peter Corcoran. Application of preconditioned alternating direction method of multipliers in depth from focal stack. Journal of Electronic Imaging, 27(2):023019, 2018.
- [109] William R Johnson, Daniel W Wilson, Wolfgang Fink, Mark Humayun, and Greg Bearman. Snapshot hyperspectral imaging in ophthalmology. Journal of biomedical optics, 12(1):014036–014036, 2007.
- [110] Sheethu Jose and Mundlapudi Lakshmi pathi Reddy. Lanthanum–strontium copper silicates as intense blue inorganic pigments with high near-infrared reflectance. Dyes and Pigments, 98(3):540–546, 2013.

- [111] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In Proceedings of the fourth Eurographics symposium on Geometry processing, volume 7, page 0, 2006.
- [112] John T Kent. Information gain and a general measure of correlation. Biometrika, 70(1):163–173, 1983.
- [113] Robert Keys. Cubic convolution interpolation for digital image processing. IEEE transactions on acoustics, speech, and signal processing, 29(6):1153–1160, 1981.
- [114] Sureeporn Khampaeng, Pichayada Katemake, and Chawan Koopipat. Optimizing multi-coloured LEDs for identifying pigments. In Optics for Arts, Architecture, and Archaeology VIII, volume 11784, pages 178–184. SPIE, 2021.
- [115] Tania Kleynhans, David W Messinger, and John K Delaney. Towards automatic classification of diffuse reflectance image cubes from paintings collected with hyperspectral cameras. Microchemical Journal, 157:104934, 2020.
- [116] Tania Kleynhans, Catherine M Schmidt Patterson, Kathryn A Dooley, David W Messinger, and John K Delaney. An alternative approach to mapping pigments in paintings with hyperspectral reflectance image cubes using artificial intelligence. Heritage Science, 8(1):1–16, 2020.
- [117] R Kohavi and F Provost. Confusion matrix. Machine learning, 30(2-3):271–274, 1998.
- [118] Fred A Kruse, AB Lefkoff, JW Boardman, KB Heidebrecht, AT Shapiro, PJ Barloon, and AFH Goetz. The spectral image processing system (SIPS)—interactive visualization and analysis of imaging spectrometer data. Remote sensing of environment, 44(2-3):145–163, 1993.
- [119] Paul Kubelka. Ein Beitrag zur Optik der Farbanstriche (Contribution to the optic of paint). Zeitschrift fur technische Physik, 12:593–601, 1931.
- [120] J. Kulkarni and M.S.R. Chetty. Depth estimation from defocused images: A survey. Journal of Engineering and Applied Sciences, (3):331–335, January 2019.
- [121] Sophia Lahlil, Jiming Xu, and Weidong Li. Influence of manufacturing parameters on the crackling process of ancient Chinese glazed ceramics. Journal of Cultural Heritage, 16(4):401–412, 2015.
- [122] Johann Heinrich Lambert. Photometria sive de mensura et gradibus luminis, colorum et umbrae. 1760.
- [123] Jae Young Lee and Rae-Hong Park. Complex-valued disparity: Unified depth model of depth from stereo, depth from focus, and depth from defocus based on the light field gradient. IEEE transactions on pattern analysis and machine intelligence, 2019.
- [124] Stefano Legnaioli, Giulia Lorenzetti, Gildo H Cavalcanti, Emanuela Grifoni, Luciano Marras, Anna Tonazzini, Emanuele Salerno, Pasquino Pallecchi, Gianna Giachi, and Vincenzo Palleschi. Recovery of archaeological wall paintings using novel multispectral imaging approaches. Heritage Science, 1:1–9, 2013.
- [125] Yoanna Leon, Ph Sciau, Anne Bouquillon, Laurent Pichon, and Ph de Parseval. PIXE (particle induced X-ray emission): A non-destructive analysis method adapted to the thin decorative coatings of antique ceramics. Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms, 291:45–52, 2012.

- [126] Vincent Lepetit, Francesc Moreno-Noguer, and P Fua. EPnP: Efficient perspective-n-point camera pose estimation. Int. J. Comput. Vis, 81(2):155–166, 2009.
- [127] Anat Levin, Rob Fergus, Frédo Durand, and William T Freeman. Image and depth from a conventional camera with a coded aperture. ACM transactions on graphics (TOG), 26(3):70, 2007.
- [128] Marc Levoy, Kari Pulli, Brian Curless, Szymon Rusinkiewicz, David Koller, Lucas Pereira, Matt Ginzton, Sean Anderson, James Davis, Jeremy Ginsberg, et al. The digital Michelangelo project: 3D scanning of large statues. In Proceedings of the 27th annual conference on Computer graphics and interactive techniques, pages 131–144, 2000.
- [129] Junxuan Li, Hongdong Li, and Yasuyuki Matsushita. Lighting, reflectance and geometry estimation from 360 panoramic stereo. In 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 10586–10595. IEEE, 2021.
- [130] Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. ACM Transactions on Graphics (ToG), 36(4):1–11, 2017.
- [131] Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. Materials for masses: SVBRDF acquisition with a single mobile phone image. In Proceedings of the European conference on computer vision (ECCV), pages 72–87, 2018.
- [132] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. ACM Transactions on Graphics (TOG), 37(6):1–11, 2018.
- [133] Yiming Lin, Pieter Peers, and Abhijeet Ghosh. On-site example-based material appearance acquisition. In Computer Graphics Forum, volume 38, pages 15–25. Wiley Online Library, 2019.
- [134] M Llusar, A Forés, JA Badenes, J Calbo, MA Tena, and Guillermo Monrós. Colour analysis of some cobalt-based blue pigments. Journal of the European Ceramic Society, 21(8):1121–1130, 2001.
- [135] Joakim Löw, Joel Kronander, Anders Ynnerman, and Jonas Unger. BRDF models for accurate and efficient rendering of glossy surfaces. ACM Transactions on Graphics (TOG), 31(1):1–14, 2012.
- [136] Xiao Xin Lu. A review of solutions for perspective-n-point problem in camera pose estimation. In Journal of Physics: Conference Series, volume 1087, page 052009. IOP Publishing, 2018.
- [137] Antoine Lucat. Acquisition opto-numérique de vêtements asiatiques anciens. PhD thesis, Université de Bordeaux, 2020.
- [138] Thomas Luhmann, Stuart Robson, Stephen Kyle, and Jan Boehm. Close-range photogrammetry and 3D imaging. In Close-Range Photogrammetry and 3D Imaging. de Gruyter, 2019.
- [139] Lei Ma, Yu Liu, Xueliang Zhang, Yuanxin Ye, Gaofei Yin, and Brian Alan Johnson. Deep learning in remote sensing applications: A meta-analysis and review. ISPRS journal of photogrammetry and remote sensing, 152:166–177, 2019.

- [140] J MacQueen. Classification and analysis of multivariate observations. In 5th Berkeley Symp. Math. Statist. Probability, pages 281–297. University of California Los Angeles LA USA, 1967.
- [141] Muhammad Tariq Mahmood, Abdul Majid, and Tae-Sun Choi. Optimal depth estimation by combining focus measures using genetic programming. Information Sciences, 181(7):1249–1263, 2011.
- [142] Tom Malzbender, Dan Gelb, and Hans Wolters. Polynomial texture maps. In Proceedings of the 28th annual conference on Computer graphics and interactive techniques, pages 519–528, 2001.
- [143] Mohamed Beji Ben Mami. La mosquée M’Hammed Bey: un exemple de la présence architecturale et artistique ottomane dans la médina de Tunis. 1998.
- [144] Stephen R Marschner, Stephen H Westin, Eric PF Lafortune, Kenneth E Torrance, and Donald P Greenberg. Image-based BRDF measurement including human skin. In Rendering Techniques’ 99: Proceedings of the Eurographics Workshop in Granada, Spain, June 21–23, 1999 10, pages 131–144. Springer, 1999.
- [145] Nicola Masini and Francesco Soldovieri. Sensing the Past. Springer, 2017.
- [146] Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. Efficient isotropic BRDF measurement. In Rendering Techniques, pages 241–248. Citeseer, 2003.
- [147] Judit Molera, Mario Vendrell-Saz, and Josefina Pérez-Arantegui. Chemical and textural characterization of tin glazes in Islamic ceramics from Eastern Spain. Journal of Archaeological Science, 28(3):331–340, 2001.
- [148] PRS Moorey. Ancient Mesopotamian materials and industries. The archaeological evidence. Oxford. 1994.
- [149] Aurélie Mounier and Floréal Daniel. Pigments & dyes in a collection of medieval illuminations (14th–16th century). Color Research & Application, 42(6):807–822, 2017.
- [150] Aurélie Mounier, Hortense de la Codre, Charlotte Marembert, Pauline Claisse, and Floréal Daniel. Rediscover the faded colours of an 18th century tapestry kept in the Cité Internationale de la tapisserie in Aubusson (France). In International Colour Association (AIC), pages 178–185, 2020.
- [151] Aurélie Mounier, Charlotte Denoël, and Floréal Daniel. Hyperspectral imaging on three French medieval illuminations of the XVI century (Treasury of Bordeaux Cathedral, France). In COLOURS2015, 2015.
- [152] Aurélie Mounier, Gwénaëlle Le Bourdon, Christian Aupetit, Sylvain Lazare, Carole Biron, Josefina Pérez-Arantegui, David Almazán, Julene Aramendia, Nagore Prieto-Taboada, S Fdez-Ortiz De Vallejuelo, et al. Red and blue colours on 18th–19th century Japanese woodblock prints: In situ analyses by spectrofluorimetry and complementary non-invasive spectroscopic methods. Microchemical Journal, 140:129–141, 2018.
- [153] Pantazis Mouroulis, David A Thomas, Thomas G Chrien, Valerie Duval, Robert O Green, John J Simmonds, and Arthur H Vaughan. Trade studies in multi/hyperspectral imaging systems final report. 1998.
- [154] Harold Mytum and JR Peterson. The application of reflectance transformation imaging (RTI) in historical archaeology. Historical Archaeology, 52:489–503, 2018.

- [155] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. ACM Transactions on Graphics (TOG), 37(6):1–12, 2018.
- [156] Shree K Nayar, Gurunandan Krishnan, Michael D Grossberg, and Ramesh Raskar. Fast separation of direct and global components of a scene using high frequency illumination. In ACM SIGGRAPH 2006 Papers, pages 935–944. 2006.
- [157] Diego Nehab, Szymon Rusinkiewicz, James Davis, and Ravi Ramamoorthi. Efficiently combining positions and normals for precise 3D geometry. ACM transactions on graphics (TOG), 24(3):536–543, 2005.
- [158] Roy G Newton and Sandra Davison. Conservation of glass. 1989.
- [159] Fred Edwin Nicodemus, Joseph C Richmond, Jack J Hsia, Irving W Ginsberg, and Thomas Limperis. Geometrical considerations and nomenclature for reflectance. Final Report National Bureau of Standards, 1977.
- [160] Jens Nieke, Horst H Schwarzer, Andreas Neumann, and Gerhard Zimmermann. Imaging spaceborne and airborne sensor systems in the beginning of the next century. In Sensors, Systems, and Next-Generation Satellites, volume 3221, pages 581–592. SPIE, 1997.
- [161] Frank Nielsen and Frank Nielsen. Hierarchical clustering. Introduction to HPC with MPI for Data Science, pages 195–211, 2016.
- [162] Gael Obein, Robert Bousquet, and Maria E Nadal. New NIST reference goniospectrometer. In Optical Diagnostics, volume 5880, pages 241–250. SPIE, 2005.
- [163] Alianda Dantas de Oliveira, Vitor Hugo da Silva, Maria Fernanda Pimentel, Glória Maria Vinhas, Celio Pasquini, and Yêda Medeiros Bastos de Almeida. Use of Infrared Spectroscopy and Near Infrared Hyperspectral Images to Evaluate Effects of Different Chemical Agents on PET Bottle Surface. Materials Research, 21, 2018.
- [164] Sarah H Parcak. Satellite remote sensing for archaeology. Routledge, 2009.
- [165] Jaesik Park, Sudipta N Sinha, Yasuyuki Matsushita, Yu-Wing Tai, and In So Kweon. Robust multiview photometric stereo using planar mesh parameterization. IEEE transactions on pattern analysis and machine intelligence, 39(8):1591–1604, 2016.
- [166] András Patay-Horváth. The virtual 3D reconstruction of the east pediment of the temple of Zeus at Olympia an old puzzle of classical archaeology in the light of recent technologies. Digital Applications in Archaeology and Cultural Heritage, 1(1):12–22, 2014.
- [167] Karl Pearson. VII. Note on regression and inheritance in the case of two parents. proceedings of the royal society of London, 58(347-352):240–242, 1895.
- [168] P Peduto. Un giardino-palazzo islamico del secolo XIII: l’artificio di Villa Rufolo a Ravello, in Apollo. Bollettino dei Musei Provinciali del Salernitano, 12:57–72, 1996.
- [169] Alex P Pentland. Depth of scene from depth of field. Technical report, SRI INTERNATIONAL MENLO PARK CA, 1982.
- [170] Said Pertuz, Domenec Puig, and Miguel Angel Garcia. Analysis of focus measure operators for shape-from-focus. Pattern Recognition, 46(5):1415–1432, 2013.
- [171] Said Pertuz, Domenec Puig, Miguel Angel Garcia, and Andrea Fusiello. Generation of All-in-Focus Images by Noise-Robust Selective Fusion of Limited Depth-of-Field Images. IEEE Transaction on Image Processing, 22(3), 2013.

- [172] Greta Peruzzi, Costanza Cucci, Marcello Picollo, Franco Quercioli, and Lorenzo Stefani. Non-invasive identification of dyed textiles by using VIS-NIR FORS and hyperspectral imaging techniques. Cultura e Scienza Del Colore-Color Culture and Science, 13(01):61–69, 2021.
- [173] Emeline Pouyet, Neda Rohani, Aggelos K Katsaggelos, Oliver Cossairt, and Marc Walton. Innovative data reduction and visualization strategy for hyperspectral imaging datasets using t-SNE approach. Pure and Applied Chemistry, 90(3):493–506, 2018.
- [174] José Manuel Prats-Montalbán, Ana de Juan, and A Ferrer. Multivariate image analysis: A review with applications. Chemometrics and Intelligent Laboratory Systems, 107(1):1–23, 2011.
- [175] Roberta Ravanelli, Lorenzo Lastilla, and Silvia Ferrara. A high-resolution photogrammetric workflow based on focus stacking for the 3D modeling of small Aegean inscriptions. Journal of Cultural Heritage, 54:130–145, 2022.
- [176] Sidney Ray. Applied photographic optics. Routledge, 2002.
- [177] Fabio Remondino and Sabry El-Hakim. Image-based 3D modelling: a review. The photogrammetric record, 21(115):269–291, 2006.
- [178] Fabio Remondino, Alberto Guarnieri, and Antonio Vettore. 3D modeling of close-range objects: photogrammetry or laser scanning? In Videometrics VIII, volume 5665, pages 216–225. SPIE, 2005.
- [179] Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo. Pocket reflectometry. ACM Transactions on Graphics (TOG), 30(4):1–10, 2011.
- [180] Ilya Reshetouski, Alkhazur Manakov, Hans-Peter Seidel, and Ivo Ihrke. Three-dimensional kaleidoscopic imaging. In CVPR 2011, pages 353–360. IEEE, 2011.
- [181] Jérémy Riviere, Ilya Reshetouski, Luka Filipi, and Abhijeet Ghosh. Polarization imaging reflectometry in the wild. ACM Transactions on Graphics (TOG), 36(6):1–14, 2017.
- [182] Mark A Robertson, Sean Borman, and Robert L Stevenson. Dynamic range improvement through multiple exposures. In Proceedings 1999 international conference on image processing (Cat. 99CH36348), volume 3, pages 159–163. IEEE, 1999.
- [183] Fabiano Romeiro, Yuriy Vasilyev, and Todd Zickler. Passive reflectometry. In Computer Vision—ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12-18, 2008, Proceedings, Part IV 10, pages 859–872. Springer, 2008.
- [184] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. Journal of computational and applied mathematics, 20:53–65, 1987.
- [185] Martin Rump, Gero Müller, Ralf Sarlette, Dirk Koch, and Reinhard Klein. Photo-realistic rendering of metallic car paint from image-based measurements. In Computer Graphics Forum, volume 27, pages 527–536. Wiley Online Library, 2008.
- [186] Szymon M Rusinkiewicz. A new change of variables for efficient BRDF representation. Rendering techniques, 98:11–22, 1998.
- [187] Parikshit Sakurikar and P. J. Narayanan. Composite Focus Measure for High Quality Depth Maps. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Oct 2017.

- [188] Emanuele Salerno, Anna Tonazzini, Emanuela Grifoni, Giulia Lorenzetti, Stefano Legnaioli, Marco Lezzerini, Luciano Marras, Stefano Pagnotta, and Vincenzo Palleschi. Analysis of multispectral images in cultural heritage and archaeology. J. Appl. Laser Spectrosc., 1(1):22–27, 2014.
- [189] Joaquim Salvi, Sergio Fernandez, Tomislav Pribanic, and Xavier Llado. A state of the art in structured light patterns for surface profilometry. Pattern recognition, 43(8):2666–2680, 2010.
- [190] Joaquim Salvi, Jordi Pages, and Joan Batlle. Pattern codification strategies in structured light systems. Pattern recognition, 37(4):827–849, 2004.
- [191] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14, pages 501–518. Springer, 2016.
- [192] Claudia Sciuto, Federico Cantini, Rémy Chapoulie, Corentin Cou, Hortense de la Codre, Gabriele Gattiglia, Xavier Granier, Aurélie Mounier, Vincenzo Palleschi, Germana Sorrentino, and Simona Raneri. What Lies Beyond Sight? Applications of Ultraportable Hyperspectral Imaging (VIS-NIR) for Archaeological Fieldwork. Journal of Field Archaeology, pages 1–14, October 2022.
- [193] Jessi Stumpf, Chris Tchou, Nathan Yun, Philippe Martinez, Tim Hawkins, Andrew Jones, Brian Emerson, and Paul E Debevec. Digital Reunification of the Parthenon and its Sculptures. In VAST, pages 41–50, 2003.
- [194] Wen-Hao Su and Da-Wen Sun. Fourier transform infrared and Raman and hyperspectral imaging techniques for quality determinations of powdery foods: A review. Comprehensive Reviews in Food Science and Food Safety, 17(1):104–122, 2018.
- [195] Yu Sun, Stefan Duthaler, and Bradley J Nelson. Autofocusing in computer microscopy: selecting the optimal focus algorithm. Microscopy research and technique, 65(3):139–149, 2004.
- [196] Jaeheung Surh, Hae-Gon Jeon, Yunwon Park, Sunghoon Im, Hyowon Ha, and In So Kweon. Noise robust depth from focus using a ring difference filter. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 6328–6337, 2017.
- [197] Tianyang Tao, Qian Chen, Shijie Feng, Yan Hu, Minliang Zhang, and Chao Zuo. High-precision real-time 3D shape measurement based on a quad-camera system. Journal of Optics, 20(1):014009, 2017.
- [198] Petra Tatzer, Markus Wolf, and Thomas Panner. Industrial application for inline material sorting using hyperspectral imaging in the NIR range. Real-Time Imaging, 11(2):99–107, 2005.
- [199] George Terzakis and Manolis Lourakis. A consistently fast and globally optimal solution to the perspective-n-point problem. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16, pages 478–494. Springer, 2020.
- [200] Robert L Thorndike. Who belongs in the family? Psychometrika, 18(4):267–276, 1953.
- [201] KR Thorp and LF Tian. A review on remote sensing of weeds in agriculture. Precision Agriculture, 5(5):477–508, 2004.

- [202] Michael Tite, Andrew Shortland, and Sarah Paynter. The beginnings of vitreous materials in the Near East and Egypt. Accounts of chemical research, 35(8):585–593, 2002.
- [203] Michael S Tite. The production technology of Italian maiolica: a reassessment. Journal of Archaeological Science, 36(10):2065–2080, 2009.
- [204] Michael S Tite, Ian C Freestone, Nigel D Meeks, and Mavis Bimson. The use of scanning electron microscopy in the technological examination of ancient ceramics. In Archaeological ceramics. 1982.
- [205] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In Sixth international conference on computer vision (IEEE Cat. No. 98CH36271), pages 839–846. IEEE, 1998.
- [206] Anna Tonazzini, Luigi Bedini, and Emanuele Salerno. Independent component analysis for document restoration. Document Analysis and Recognition, 7(1):17–27, 2004.
- [207] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. Josa, 57(9):1105–1114, 1967.
- [208] Borom Tunwattanapong, Graham Fyffe, Paul Graham, Jay Busch, Xueming Yu, Abhijeet Ghosh, and Paul Debevec. Acquiring reflectance and shape from continuous spherical harmonic illumination. ACM Transactions on graphics (TOG), 32(4):1–12, 2013.
- [209] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. Journal of machine learning research, 9(11), 2008.
- [210] Mercedes Volait. Fous du Caire. Excentriques, architectes et amateurs d’art en Egypte (1867-1914). page 304, 2009.
- [211] Mercedes Volait. Maisons de France au Caire : le remploi de grands décors mamelouks et ottomans dans une architecture moderne. December 2012.
- [212] Hermann Von Helmholtz. Handbuch der physiologischen Optik, volume 9. Voss, 1867.
- [213] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. In Proceedings of the 18th Eurographics conference on Rendering Techniques, pages 195–206, 2007.
- [214] Chun-Po Wang, Noah Snavely, and Steve Marschner. Estimating dual-scale properties of glossy surfaces from step-edge lighting. In Proceedings of the 2011 SIGGRAPH Asia Conference, pages 1–12, 2011.
- [215] Oliver Wang, Prabath Gunawardane, Steve Scher, and James Davis. Material classification using BRDF slices. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 2805–2811. IEEE, 2009.
- [216] Matthew J Westoby, James Brasington, Niel F Glasser, Michael J Hambrey, and Jennifer M Reynolds. ‘Structure-from-Motion’ photogrammetry: A low-cost, effective tool for geoscience applications. Geomorphology, 179:300–314, 2012.
- [217] Robert J Woodham. Photometric method for determining surface orientation from multiple images. Optical engineering, 19(1):139–144, 1980.
- [218] Günther Wyszecki and Walter Stanley Stiles. Color science: concepts and methods, quantitative data and formulae. 40, 2000.

- [219] Wenjie Ye, Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. Single image surface appearance modeling with self-augmented cnns and inexact supervision. In Computer Graphics Forum, volume 37, pages 201–211. Wiley Online Library, 2018.
- [220] Tim Zaman. Development of a topographical imaging device for the near-planar surfaces of paintings. 2013.
- [221] Renata ŽEMAITYTĖ, Vaida JONAITIENĖ, Rimvydas MILAŠIUS, Sigitas STANYŠ, and Regina ULOZAITĖ. Analysis and identification of fibre constitution of archaeological textiles. chemical analysis, 2:7–10, 2006.
- [222] Zhengyou Zhang. A flexible new technique for camera calibration. IEEE Transactions on pattern analysis and machine intelligence, 22(11):1330–1334, 2000.
- [223] Yonghui Zhao, Roy S Berns, and L Taplin. Image segmentation and pigment mapping in spectral imaging. In Proceedings of the 6th International Congress of Imaging Science, pages 294–297, 2006.
- [224] Enliang Zheng and Changchang Wu. Structure from motion using structure-less resection. In Proceedings of the IEEE International Conference on Computer Vision, pages 2075–2083, 2015.
- [225] Zhiming Zhou, Guojun Chen, Yue Dong, David Wipf, Yong Yu, John Snyder, and Xin Tong. Sparse-as-possible SVBRDF acquisition. ACM Transactions on Graphics (TOG), 35(6):1–12, 2016.