

**UNIVERSITAT
JAUME·I**

Programa de Doctorado en Ciencias

Escuela de Doctorado de la Universitat Jaume I

Multidimensional Computational Imaging using Single-pixel Detectors

**Memoria presentada por Fernando Soldevila Torres para optar al grado de
doctor por la Universitat Jaume I**

Escrita y realizada por:

Fernando Soldevila Torres

Dirigida por:

Enrique Tajahuerce Romera &
Vicente Durán Bosch

Castellón de la Plana, noviembre de 2018

“[...] L’amour des Lettres, qui est un mérite chez nos voisins, n’est encore à la vérité qu’une mode parmi nous, & ne sera peut-être jamais autre chose; mais quelque dangereuse que soit cette mode, qui pour un Mécène éclairé produit cent Amateurs ignorans & orgueilleux, peut-être lui sommes-nous redevables de n’être pas encore tombés dans la barbarie où une foule de circonstances tendent à nous précipiter.

On peut regarder comme une des principales, cet amour du faux bel esprit, qui protège l’ignorance, qui s’en fait honneur, & qui la répandra universellement tôt ou tard. Elle sera le fruit & le terme du mauvais goût; j’ajoute qu’elle en sera le remède. Car tout a des révolutions réglées, & l’obscurité se terminera par un nouveau siècle de lumière. Nous serons plus frappés du grand jour, après avoir été quelque tems dans les ténèbres. Elles seront comme une espèce d’anarchie très-funeste par elle-même, mais quelquefois utile par ses suites. Gardons-nous pourtant de souhaiter une révolution si redoutable; la barbarie dure des siècles, il semble que ce soit notre élément; la raison & le bon goût ne font que passer. [...]”

“[...] El amor a las letras, que es un mérito entre nuestros vecinos, entre nosotros no es aún más que una moda, y acaso no sea nunca otra cosa; pero por muy peligrosa que sea esta moda, que, por un Mecenas inteligente produce cien aficionados ignorantes y orgullosos, quizá le debemos el no haber caído todavía en la barbarie a que tienden a precipitarnos multitud de circunstancias.

Se puede considerar como una de las principales ese amor al falso ingenio que protege a la ignorancia, que presume de él y que la difundirá universalmente más tarde o más temprano. Será el fruto y el término del mal gusto; añadido que será su remedio. Pues todo tiene revoluciones previstas, y la oscuridad terminará en un nuevo siglo de luz. La claridad nos impresionará más después de haber permanecido algún tiempo en las tinieblas. Será como una especie de anarquía muy funesta en sí, pero útil en sus consecuencias. Librémonos, sin embargo, de desear una revolución tan temible; la barbarie dura siglos, y parece que es nuestro elemento; la razón y el buen gusto son pasajeros. [...]”

Jean le Rond d’Alembert. *Prólogo a la Encyclopédie, ou Dictionnaire raisonné des sciences, des arts et des métiers.* (1751–1772).

UNIVERSITAT JAUME I

*Abstract*Escola Superior de Tecnologia i Ciències Experimentals, ESTCE
Physics

Doctor of Philosophy

Multidimensional Computational Imaging using Single-pixel Detectors

by Fernando Soldevila

Humans depend on vision to gather information about their surroundings. However, our sight sense is very limited: we cannot see very distant or small objects, and we can only sense light inside the visible spectra. To tackle these limitations, we have been developing optical sensing tools for more than four centuries now. However, even though nowadays we can even see objects at the nanometric scale, or very distant galaxies, there are still fundamental limitations that physical systems cannot bypass.

In this thesis, I will show you how to obtain images with information about multiple dimensions of light (polarization, phase, wavelength) using novel sensing paradigms based on single-pixel detection and signal processing techniques. Using detectors without spatial resolution makes it possible to easily work in exotic spectral ranges, in low light level scenarios, or to build very compact and efficient multidimensional imaging systems. Moreover, the presence of a fast spatial light modulator in all of these systems allows to implement modern recovery techniques based on algorithmic approaches, such as compressive sensing or matrix completion. In doing so, these computational imaging systems can obtain more information than a traditional system, but in a faster and inexpensive manner.

Agradecimientos

Quiero empezar dando las gracias a mis mentores: Pedro, Jesús y Enrique. De vosotros he aprendido cómo hacer Ciencia. Sigo con mis otros mentores: Vicente y Pere. Si no hubiese sido por vuestra ayuda al principio, no habría llegado hasta aquí. Por último, mis otros, otros mentores: Gladys, Esther, Lluís, Merche, Eva, Ángel, Omel, Yessenia, Carlos, Verónica, Miguel, Susana, Vicent, Raúl, Salva, Humberto, Jorge, Alex, Rocío, Eduardo, Dani, Jennifer, Armin, Marc y Omar. Mejor no pongo qué he aprendido de cada uno de vosotros.

Por supuesto, a mi familia, que siempre ha estado a mi lado cuando lo he necesitado. A mis padres por haberme llevado siempre por el buen camino. A mi hermana por soportarme. A mi tía y mis primos, que han sido una segunda madre y unos hermanos. A mis sobrinos, que me han hecho viajar en el tiempo. A los que no han podido ver el final de esta tesis.

To the people in Paris. You made me feel at home, this thesis is also yours. Special thanks to Sylvain and Hilton, for mentoring me in novel fields.

A la UJI, por pagarme durante estos años por hacer lo que más me gusta en el mundo. Y también por haber hecho posible que investigase.

A Nintendo, Valve (aunque tenga problemas con el número 3) y otros, porque también es necesario desconectar de vez en cuando. A Soundcloud, porque esta tesis la he escrito escuchando a Tycho 24 horas al día. A los creadores de Mendeley y L^AT_EX, ya que sin estas herramientas, estas líneas, si bien posibles, habrían sido mucho más feas. A los que suben Matlab (y otras cosas) a ThePirateBay, que somos pobres y sin ánimo de lucro.

Y ahora, lo que muchos están esperando. Si no hubiese sido por todos vosotros, habría terminado esta tesis mucho antes. Que cada cual se vea reflejado donde quiera. Gracias a copito de nieve; al canario más sabrosón; a los que llaman taxis cuando se quedan sin batería en el móvil; a la costra; al plomero marrón que hace doble *clicaso*; a mi Paulo Coelho particular, si no es por ti lo habría dejado todo a medias; a quien quiso ser mi amigo en la fila del colegio; al que vuelca con panteras negras; a la mujer de hielo; a la espía que vino del frío; a la que esconde pollo en los pasillos; al que me dio un triple medusa; a los bares clandestinos que se esconden en bucles del espacio-tiempo; al judío *McMadden*; al sereno de Balsicas; a los lanzadores de árboles; al que va a odiar estas líneas; al ilocalizable en Portugal; a la racista del cercanías; al GPS humano; a los que se pasan Elche; a los cocreadores del triple uno; al hombre a una nariz pegado; a los kurruki; a los bichos; a los Judas Priest; al que hace triple 50; a la media moto de queso; a casa Hitler, con sus cervezas para embarazadas y los ladrones de jarras; al que fotografía la luna; a los que no robaron ninguna jarra; a las albóndigas a -30 °C; a la que me recuerda a mi hermana; a doña exudado vaginal; a la que dibuja penes en Rusia; a los jovencitos confusos que van a vino; al origen del mundo; al que te encuentras aplanado en el sofá cuando llegas a casa; a los que comen *flaons*; a las jovencitas confusas que van a otras cosas; a la ruta de la tapa; a la gente sin sentido del humor; al cuscús; al conductor colérico; a los proletarios; al *team* retrasito; a *Peroni e Calzone*; a la gente con sentido del humor; a la que me perdió las llaves del coche; al que comió hormigas conmigo; al que ya no es alérgico al pescado; al agua; a los autobuses a Murcia; al doble postre; a la que no quería más, pese a estar todo muy bueno; a los que vinieron de lejos, pero han estado muy cerca; a la que me manda audios de Fito; al que hace triple pared; a los desaparecidos; a la señora mayor; a la longaniza seca con leche, un ibuprofeno y medio litro de agua; al abuelete pedante; a la calma que precede a la tempestad; a la que rompe puertas de coche; a la funcionaria antisistema; a los parásitos; al artista

que programa; al exterminador de mosquitos; a los chili; al que se hace ecografías de los riñones pasando por los testículos; a la que se duerme en el coche; al salto patentado; a los que se bañan; al triple postre; a los que no les salen las prácticas; a Piero; a los ordenadores en el baño y a todos los que me dejó en el tintero debido a mi memoria lamentable.

A todos vosotros, gracias.

Contents

Abstract	iii
Agradecimientos	v
1 Introduction	1
1.1 Looking for answers	1
1.2 Making the right questions: multidimensional imaging	3
1.3 Welcome to the world of tomorrow: computational imaging	5
1.4 A drop in the ocean	7
2 Theoretical fundamentals: the pieces of the puzzle	9
2.1 Single-pixel imaging	9
2.2 Going beyond matrix inversion	13
2.2.1 Compressive sensing	16
2.2.2 Adaptive strategies	19
2.2.3 Matrix completion	25
2.3 Coping with the limitations of single-pixel imaging systems	29
3 Experimental results: putting it together	33
3.1 Spectropolarimetric imaging using a single-pixel camera	33
3.2 Computational Raman microspectroscopy	37
3.3 Phase imaging using a single-pixel camera	42
4 So long, and thanks for all the fish	47
A Taming light: Digital Micromirror Devices	49
B Singular Value Decomposition	51
C List of publications	53
Bibliography	113

List of Abbreviations

AI	Artificial Intelligence
APD	Avalanche PhotoDiode
CARS	Coherent Anti-Stokes Raman Spectroscopy
CCD	Charged Coupled Device
CGI	Computational Ghost Imaging
CMOS	Complementary Metal–Oxide Semiconductor
DMD	Digital Micromirror Device
IR	Infra Red
LCVR	Liquid Crystal Variable Retarder
NIR	Near Infra Red
PMT	PhotoMultiplier Tube
RMS	Raman MicroSpectroscopy
RS	Raster Scanning
SHWS	Shack Hartmann Wavefront Sensor
SLM	Spatial Light Modulator
SNR	Signal-to-Noise Ratio
SoP	State of Polarization
SPAD	Single-Photon Avalanche Diode
SPI	Single-Pixel Imaging
SPC	Single-Pixel Camera
SPD	Single-Pixel Detector

A mi familia.

Chapter 1

Introduction

1.1 Looking for answers

Observation is one of the pillars of the scientific method. Whether for proving the existence of gravitational waves or for understanding the inner procedures of how our brain works, scientists need to use precise and reproducible sensing mechanisms.

The most immediate way to perform a physical observation is to use any of our senses. However, as any thermodynamics student would tell, it does not seem a very smart idea to measure the temperature of an oven by using your sense of touch (it could be a one-time measurement!). Furthermore, it is not very practical to say to your colleagues that water boils at “a high temperature”. Anyone wondering the temperature of a system can use a thermometer to precisely (quantitatively) determine it. The act of sensing can be understood as a question-answer process. Maybe you just want to know which is the mass of some apples. In that case, you “ask” the apples what their weight is (using a scale), and after some physical process, the apples “respond” the value. Even if this may seem simplistic, it is useful to interpret sensing processes as questions, as we will see along this thesis.

As science advances and we develop new questions about nature, observation methods also get more and more sophisticated. For example, balances were used since ancient times to compare the masses between objects (see Fig. 1.1). However, when you want to measure Earth’s mass, you need something more complex than a bunch of sticks and strings.



FIGURE 1.1: **Anubis weighing the heart of Hunefer.** Dated from 1300 BC, this painting depicts the use of a balance to determine if Hunefer’s heart was heavier than a feather. Extracted from *Papyrus of Hunefer*.

Optical sensing is a very good example of how sensing has become more and more elaborated. In the beginning, we used our eyes to watch the world around us.

To share that information, we simply told or drew what we saw. Again, this way of understanding our surroundings is very simplistic. First, our sight is limited in both range and detail. We neither can see very distant nor very small objects. Moreover, every observation based on human perception is inherently flawed, and, even in the exceptional case of seeing an event in the exact way it happened, our descriptions, either artistic or linguistic, cannot represent reality as it truly is.

In order to overcome the limitations of human vision, we started using optical instruments. Simple combinations of lenses were used at the end of the 16th century in the Netherlands to build the first optical telescopes and microscopes. These inventions allowed us to study both the most distant and the smallest objects of the known universe at the time. Whereas telescopes were the main tool used to lay the foundations of modern astronomy, microscopes have been a key device for almost all the development of life sciences.

However, there was still the problem of human representation. In order to produce accurate images, scientists needed to develop keen eyes and steady hands. There is a long list of scientists-artists that were key for the scientific advance of their disciplines (see Fig. 1.2). However, that was a temporary period: after the discovery of photography, everything became easier. Photographic cameras have always been better at depicting reality than humans: they do not have personal preferences, or get tired, after all. Once cameras became sophisticated enough, natural sciences rapidly adopted photography as their standard to communicate and store visual information (Fig. 1.2). Nevertheless, photography (as lots of other techniques just by themselves) cannot provide enough information for cutting edge applications in modern industry and science. During the last four decades, researchers have switched from just counting photons to simultaneously count them and their physical properties.




FIGURE 1.2: **Galapagos finches.** On the left we can see a John Gould depiction of the morphology of the animals in the Galapagos Islands, that were the foundations of the theory of evolution by natural selection. On the right, we see a modern photograph of *Camarhynchus pauper*, taken by using a long distance photographic objective and a digital sensor. Pictures adapted from [1] and [2].

1.2 Making the right questions: multidimensional imaging

Nowadays, when you take a photograph with your camera, the question you are formulating is “*how many photons are coming from every point of the scene I want to study?*”. If we take a look at the digital file of a picture, we can see a numerical value for each pixel of the image. This number is proportional to the amount of photons that arrived to that pixel of the sensor, or, if your sensor is good enough (i.e. expensive enough), the exact number of photons that arrived, excluding some small amount of noise.

Counting photons is good, but measuring other parameters can provide additional useful information (see Fig. 1.3). Imagine a doctor that wants to track your health. Of course, your health records contain your name and address, making it easy to identify and distinguish you from other patients. However, tracking your weight, your blood type and the proportions of its contents, and also a record of your past illnesses, provides an excellent way to identify present and possible future health problems, making doctor’s work both easier and more efficient. These additional parameters (or dimensions) can be used in combination to produce results that none of them could provide individually.

In the case of light, changing “*how many*” for “*when*” provides temporal resolution to your system. Now you can measure both the spatial distribution of photons and their arrival timestamps. With this information you can easily discern depth in your pictures. Each physical property is a good candidate to ask about. By looking at the polarization, the stress distribution of a windscreen can be obtained. Wavelength provides information about the maturity of crops. The direction of light coming from distant stars is perturbed by our sky, and measuring those changes is key to obtain good astronomical pictures. Obtaining the spatial distribution of these parameters in addition to just the intensity distribution of light is what we know as multidimensional imaging.

As of today, we can capture images that represent almost any physical quantity you can name: energy, wavelength, polarization, phase, temperature, or pressure, to name a few. In fact, the list has been growing non-stop since the development of digital photography. This capability has been key to the development of industry and science. Nowadays, all the leading industries use some kind of multidimensional image system in quality control or in their safety systems. To name just a few, polarization imaging is widely used to check the stress and defects of transparent objects, such as glass or plastics [5]. Multispectral imaging is commonly used in the food industry to check both maturity and defects of fruits [6]. Augmented reality systems are starting to be used as an everyday tool in medicine and industry, facilitating the work of operators, and therefore increasing efficiency and reducing costs [7]. Phase imaging systems are key in the development of many optical devices, for example in the design and test of ophthalmic lenses [8]. In the case of science, maybe the best example resides in life sciences. As scientists went deeper in the understanding of the human body, it was clear that providing just images of the amount of light going through a sample was not going to be enough. Information about the polarization and the spectrum of light is currently being used to diagnose biological tissue [9]. Phase and spectral information is utilized to study the working principles of neural networks in emerging fields like Optogenetics [10]. Moreover, real-time three-dimensional images help on visualizing and understanding physicians not only the underlying structures of organs, but also their inner functioning. In short, any task were humans need to make a decision is liable to be enhanced by providing extra information.




FIGURE 1.3: **Multidimensional imaging examples.** Three examples of additional dimensions added to a two-dimensional picture. In the top inset, spectral information from a satellite orbiting earth (picture from NASA Terra satellite). In the middle inset, we can see polarization fringes and colors caused by the present stress in a pair of conventional eyeglasses (winner picture from OSA 2008 Photo Contest, by Osvaldo Buccafusca [3]). In the lower inset we can see a group of cells imaged by using differential interference contrast microscopy, extracted from [4].

Nevertheless, to obtain this information, optical systems are required to be built accordingly. Standard digital sensors, such as charge coupled devices (CCD) or complementary metal-oxide semiconductor (CMOS) just provide information about the irradiance distribution of photons (they can only ask a simple question). For example, in order to measure wavelength or polarization, you need to add additional optical elements, such as diffraction gratings and polarizers. During the last decades, great advances have been made in this direction, and nowadays it is possible to buy commercial systems that allow users to obtain spectral and polarization




FIGURE 1.4: **Computational imaging examples.** On the left, an example of STORM microscopy, where multiple fluorophores are activated and imaged sequentially. After all the acquisitions are done, an algorithm provides the final image (shown on the bottom row) with enhanced spatial resolution. On the right, we can see a three-dimensional image of a mouse embryo obtained with computerized tomography. After illuminating the sample with x-rays from different incidence angles, multiple slices of the sample can be obtained. From that, an algorithm provides a three-dimensional representation of the sample. Images adapted from [11, 12].

images. However, there is a price to pay when obtaining this supplemental information. Adding optical elements makes the system more complex, which usually entails bigger, slower, and more expensive devices. In the case of both multispectral and polarization imaging, available systems tend to suffer from low temporal and spatial resolution. Also, having access to bigger amounts of data is very useful, but it compels electronics and computer systems to being able to process and store all this information in an efficient way.

There is a lot of research being done to improve and cheapen the hardware related to multidimensional imaging systems. Advances in electronics and manufacturing have made possible to buy cost-efficient sensors with millions of pixels under the size of a coin. Also, both raw computing power and electronic storage capacity keep growing every day. Yet the increase in the amount of information has outpaced those improvements. In order to tackle this problem, there are novel approaches that rely not only on the hardware but also on intangible elements. Physics, mathematics, and computer algorithms are merging with the hardware to make novel questions to nature and get information in a way that has never been possible before. We will see how next.

1.3 Welcome to the world of tomorrow: computational imaging

How can you take a picture of something transparent? What about something that is hidden by living tissue or behind a corner? How do you get a three-dimensional image of an object? Photons travelling through a transparent object will not suffer from absorption, and thus optical sensors will not serve for imaging that object. Light rays

will deviate in all directions when travelling through a scattering medium, disabling traditional optical systems ability to provide images. Light coming from different planes in your scene will merge at your detector, making it impossible to discern depth. If you want to tackle those problems using classical optics, you need to perturb your scene in some way. For example, you can stain your cell samples, making them opaque. You can also chemically treat a tissue to make it transparent [13]. Last, you can make very thin slices of your object and image each one at a time. All of these solutions share the fact that they are invasive, and thus not good for lots of applications, such as biological imaging, remote sensing, or cultural heritage imaging.

Computational imaging techniques obtain information in an indirect way. Coming back to the doctor analogy, if you go to the clinic because your head hurts, the traditional imaging way would be to open you and see what is inside (or take a sample). A “computational doctor” would check your temperature, look at your eyes and ears, and then make a diagnosis based on all the information and an underlying scientific model. The fundamental idea in our case is, starting from measurements obtained with an optical system, to combine them using computer algorithms to obtain the desired information. The algorithms need to be based on the existing relation between the information contained in the measurements and the desired information to be obtained (in the same way that pupil dilation is related to brain damage). This mathematical relation is where physics gets involved.

Take, for example, the case of obtaining information about a transparent sample. You can solve that problem by using phase shifting interferometry. In this technique, light coming from the sample under study and a reference beam interferes at the sensor. This interference causes the obtained image to present a characteristic fringe pattern, that depends on both the phase introduced by the sample (which is not known), and the phase of the reference beam (which can be controlled). By changing the phase of the reference beam several times, multiple images with different fringe patterns are acquired, and then a simple algorithm that takes into account the laws of wave interference provides an amplitude and phase image of the scene.

As in the case of hardware design, each experimental condition will determine both the optical system and the algorithm used. Knowing which is the information you can obtain from your system and how to go from that to the things you want to measure is what computational imaging tries to accomplish. At first, computational imaging systems were based on capturing images using traditional optical systems with little variations. For example, in the example of imaging a transparent sample, introducing an additional light beam and some phase retardations. However, nowadays it is possible to control light in ways that were unimaginable some decades ago. Spatial light modulators (SLM) have established as a fundamental tool in optics laboratories all over the world to control both amplitude and phase of a light beam at will. By doing so, scientists stopped to just looking at light beams, and started interacting with them in the same way as musicians play with their instruments. This taming of light has led to an apparent paradox: there are a lot of scenarios where the best way to obtain an image of a sample is not to image it with a conventional camera. This has opened the door to imaging systems where detection is not done with CCD or CMOS sensors, but with more exotic detectors, such as spectrometers, beam polarimeters, or photon counting detectors. In doing so, some of the hardware limitations we saw in the last section can be solved, as we will see in the following pages.

Computational imaging has allowed researchers to understand the imaging process in a completely different way. Once you introduce computer science concepts

onto your sensing mechanisms, new ways of retrieving information never stop to appear. Before, we used to ask our systems “simple” questions, such as *how many photons?*, *when did they arrive?*, or *what is their spectral distribution?* Now, we are starting to make questions that do not even need to be formulated in human language. We can obtain images by illuminating scenes with structured patterns and minimizing mathematical functions [14]. We can capture the shadows projected by our scene to obtain its three-dimensional profile [15]. We can even toy around physical limitations, like the resolution limit of an optical system or the rate at which information can be acquired [16, 17]. Of course, there are some limitations. First and most important, you can come up with almost infinite algorithms to perform a task, but most of them will not be fast enough to be useful. In most practical scenarios, imaging systems need to provide images in real time and with high spatial resolution. Computing power has been increasing non-stop since the invention of computers, but it is not an infinite resource. The design of efficient and fast computational techniques is one of the main research areas of the field. There is also the problem of storage. A biology lab can generate terabytes of images every few minutes. Handling all this information (storing, labelling, and making it possible to access all data to process it) has proved to be a headache for researchers, and the amount of pictures being taken keeps getting bigger every day. Historically, the main solution used to tackle this problem has been to compress your data after acquisition, thus reducing storage needs. However, the novel field of signal processing has changed the traditional sensing paradigms. Now, instead of compressing information after acquisition, you can directly measure a limited amount of data in the first place, and then use computer algorithms to “fill in” the missing parts in a clever way. This reduces storage demands, but also provides faster ways to measure (as only the “relevant” parts of information are typically acquired) and reduces bandwidth requirements.

1.4 A drop in the ocean

During my PhD, I have focused on the study of multidimensional imaging systems in combination with computational imaging techniques. However, given the broad range of possibilities of both fields, to just say that is to say nothing. More precisely, I have worked on the development of systems based on a specific kind of technology: the single-pixel camera (SPC). Single-pixel imaging (SPI), also known as computational ghost imaging (CGI), is a computational imaging technique that stands out by obtaining images using detectors without spatial resolution (i.e., using detectors with only one pixel). In order to obtain spatial information, typical single-pixel cameras use SLMs to generate structured patterns that can be used to illuminate your sample. After that, light coming from the scene is collected and measured with the bucket detector. The measured signal, in combination with the knowledge of the patterns that are generated onto the SLM, allows the user to recover the image of the object by solving a simple matrix inversion problem.

This way of obtaining images is beneficial in multiple scenarios. First, single-pixel detectors can be easily built in a wide range of spectral ranges, whereas traditional pixelated sensors outside the visible spectral range have prohibitive costs or are just simply impossible to assemble. This simple fact has led to multiple SPC applications in the infrared [18–22], terahertz [23–27], and X and gamma rays [28–31] regions of the electromagnetic spectrum. Also, there are harsh experimental scenarios that can also benefit from bucket detection. When working with very low amounts of light, SPCs can be built using extremely sensitive detectors, such as

Avalanche Photo-Diodes (APD) or Single-Photon Avalanche Diodes (SPAD) [32–34]. Even though it is possible to build arrays of very sensitive sensors, their costs are much higher than their single-pixel counterparts. SPI systems have also been used to provide images in biological scenarios where light scattering hinders the capability of conventional imaging systems to provide good results [35, 36] or to obtain additional morphological information [37]. Additionally, SPI-based systems have also been proved to be a viable tool to transmit information in a secure way [38, 39]. Last, given the simple nature of SPI detection systems, they have demonstrated to be a good candidate when designing multidimensional imaging systems, as it is easy to build spectral [40, 41], polarization [42, 43], depth [44–46], and phase [47–49] sensitive systems using this sensing paradigm.

During this thesis, I will show several experiments that demonstrate how to obtain spectral [50], polarization [50, 51] and phase [52] information using SPI techniques. However, SPCs still have some drawbacks. The main one is that the patterns are generated onto a SLM in a sequential manner, and thus the capture process time increases. In a traditional camera, images are acquired in one shot. In a SPC, to get an image consisting of N pixels, you need to generate N different patterns onto the SLM. This drawback hinders the capability of SPI systems to work in real time with high spatial resolution.

In order to tackle this problem, the SPI systems used here were based on digital micromirror devices (DMD). DMDs are binary spatial light modulators with refresh rates in the range of tens of kilohertz. Given their speed, their easy polarization characterization, and the fact that they can operate in wide spectral ranges, they have proven themselves to be an excellent choice to work in single-pixel multidimensional applications. Also, the fast control of light they provide has been exploited to further increase the capabilities of our SPCs. Using signal processing techniques, such as compressive sensing (CS), adaptive algorithms, and matrix completion (MC) procedures, it is possible to build systems that surpass the Shannon-Nyquist sampling theorem, increasing the overall acquisition and post-processing speed [53–55], and reducing the information throughput and storage limitations of multi-dimensional systems. These applications will also be explained thorough the text.

Chapter 2

Theoretical fundamentals: the pieces of the puzzle

During this chapter, I am going to show all the fundamental aspects of the techniques I have been using along the thesis. I will start with an introduction to Single-Pixel Imaging, showing both the fundamentals in relation to its hardware (the Digital Micromirror Device and its characteristics). I will continue with the associated signal processing techniques I have used (matrix inversion, compressive sensing, and matrix completion). Finally, I will describe a way to improve the design of single-pixel cameras: the use of a balanced detection scheme to leverage the operation principle of the DMD and increase both the acquisition speed and the SNR of a SPC.

2.1 Single-pixel imaging

The term *megapixel* has recently become common in our daily lives. With the advent of digital consumer photography, manufacturers started advertising their cameras using whatever means they had available to attract people. Bright colors, camera size, ease of use... and sensor size. Even though common people usually do not know what are the benefits or drawbacks of having more or less pixels in your sensor, the common reasoning has been “*the bigger, the better*”. An thus marketing campaigns started using the number of megapixels as a key element to take into account when buying a digital camera. This worked quite well, as can be seen in Fig. 2.1. As search metrics indicate, buyers started worrying about the number of pixels on their future sensors. Even though some time has passed and sensor size has lost some interest, there still exists a trend on manufacturers about creating sensors with higher and higher number of pixels (see Fig. 2.2).

However, when someone asks me the question “*how many megapixels does your camera have?*”, I can answer “ 10^{-6} megapixels”. After saying that, I usually start explaining how can a SPC work. Obviously, the main difference between a conventional camera and a SPC is the sensor. Whereas conventional CCD or CMOS sensors use millions of pixels to measure light coming from a scene, a SPC concentrates all the light in just a pixel¹. Its basic operation principle can be seen in Fig. 2.3.

Conventional cameras create a one-to-one map between the regions of the scene that you want to capture and the regions (pixels) of your sensor by using an optical system (i.e., they create an image of the scene into the sensor plane, where pixels measure the corresponding amount of light for each region). As each zone of the

¹When talking about single-pixel detectors, it is important to have in mind that, even though the sensor only has one pixel, we do not necessarily refer to a point-detector (its size can be arbitrary big). This is why these kind of sensors are also usually labelled as “bucket” detectors.




FIGURE 2.1: **Search trend for the keyword “megapixel” since 2004.** In the graph, we can see a representation of the relevance of the term, given by the number of searches in *Google Search*. The vertical axis goes from 0 to 100, 0 meaning minimum number of searches and 100 meaning maximum number of searches over the full period of time (from 2004 to 2018). We can clearly see several spikes, spaced roughly twelve months apart, and corresponding to the months of November and December, where Christmas shopping season occurs. Data extracted from *Google Trends*.




FIGURE 2.2: **Evolution of pixel count in digital sensors.** In less than 50 years, the number of pixels in digital cameras has grown in more than three orders of magnitude.

scene gets measured by a different set of pixels, the whole image can be acquired in a single exposure (Fig. 2.3.a). Great advances in semiconductor manufacturing technologies, which have entailed miniaturization of both physical size and price of the sensors, have made CCD and CMOS cameras the go-to solution when one wants to capture a digital image.

However, there are other ways of obtaining a picture. Instead of lighting the whole scene at the same time, you can illuminate just one small region, for example, by using a small bright point generated with a laser beam (Fig. 2.3.b). Then, you can measure the light coming from that point with your sensor. By moving the point along the scene, you can measure the light reflected by the object for each position, thus obtaining the same image as before. This is usually known as raster scanning (RS). In a RS setup, you do not need to use an array detector for measuring all the information. However, now the image is not acquired in a single exposure, as you need to scan the scene with your bright spot in a sequential manner. You




FIGURE 2.3: **Single-Pixel Imaging fundamentals.** **a)** Conventional image acquisition using a pixelated detector. The scene is imaged onto the sensor plane with an optical system, and each pixel of the detector measures the amount of photons coming from each region of the scene. **b)** Raster scanning image acquisition. A bright spot moves along the whole scene, and the light reflected is collected with an optical system onto a bucket detector. The image is digitally recovered after the spot has scanned the whole scene. **c)** Single-pixel imaging setup. Instead of scanning the scene with a small point, the whole scene is illuminated with wide-field structured patterns. Light reflected by the scene is collected and measured with a bucket detector. After the full scan is completed, the image is digitally recovered.

might be asking “*why would I like to obtain an image in that way?*”. After all, pixelated sensors are cheap and fast. The answer is quite simple: there are circumstances where a sensor array is neither cheap, nor fast. Moreover, there are cases where even if you have access to a good pixelated detector, their single-pixel counterpart offers better specifications. In my opinion, the paramount example of this is the confocal microscope.

In a conventional wide-field fluorescence microscope, the whole sample is illuminated with a light source (as in a conventional camera setup). All the fluorescent regions of the sample are excited at the same time, and the signal is measured with an array detector. This is problematic because samples are three-dimensional, and light from several depths is mixed in the detector plane. This makes images present a “blurry” aspect, making analysis harder (see Fig. 2.4). In the case of confocal microscopy, the sample is illuminated in a RS fashion (usually by using a set of galvanometer mirrors), where the scanning bright point is optically conjugated with a pinhole that selects light coming from only one depth of the sample. By doing this, the background caused by light coming from different depths is eliminated, and noise-free images can be obtained, as can be seen in Fig. 2.4. So, why is this a good example of the use of SPI? First, the optical setup used to remove background signal makes the technique work in a RS way. In doing so, the detector does not need to be a camera any more. In fact, in fluorescence experiments signals tend

to be very dim (both caused by fluorescence efficiency and by the light levels required to prevent biological damage), and in order to have a good signal-to-noise ratio (SNR), very sensitive detectors need to be used. Even though it is possible to buy CCD and CMOS sensors able to work in this low light level scenarios, the go-to solution for confocal microscopes usually is a photomultiplier tube (PMT) or an avalanche photodiode (APD), both of them in the family of single-pixel detectors. Using single-pixel detectors makes it possible to work with the required high sensitivities, in a wide spectral band (covering the visible and near infra-red), and with a very fast response (in the range of MHz), which is crucial to obtain images of living samples.




FIGURE 2.4: **Wide-field fluorescence versus confocal fluorescence microscopy.** At the top we can see an image of a *Convallaria rhizome* slice acquired by using a conventional microscope (using a pixelated detector). At the bottom, the same sample but imaged using a raster scan (RS) approach with a confocal microscope. It can be clearly seen that the RS system removes the characteristic halo produced by light scattered when travelling through biological tissues. Images kindly provided by E. Irles.

Nevertheless, confocal microscopists do not say that they work with a SPC. This is because the common definition of a SPC is a little bit different. The process of acquiring an image can be easily formulated in the following way. Suppose a pixelated object, \mathbf{x} , each one of its pixels containing an unknown intensity, $x(i)$. Each measurement consists of measuring the total intensity coming from the overlap of the object and a spatial light mask (that represents a mathematical function). The whole measurement process can be expressed in matrix form as

$$\mathbf{y} = \mathbf{S} \cdot \mathbf{x}, \quad (2.1)$$

where \mathbf{y} is the vector containing all the measured intensities; \mathbf{x} is the object, expressed in vector form; and \mathbf{S} is a matrix where the rows contain the masks that are used experimentally.

It is easy to see that the RS approach (Fig. 2.3.b) can be mathematically understood as a measurement of the object into the canonical (or pixel) basis. In this case, the matrix \mathbf{S} is equal to the identity matrix, \mathbf{I} , and each one of our measurements corresponds to the intensity of one of the pixels of the object, \mathbf{x} . After measuring all the pixels, recovering the object is straightforward.

However, once we have introduced this formulation, we can go further and use any orthonormal set of functions as our sensing matrix, not just the identity. Given an object with dimension N pixels, you just need to pick a mathematical basis of the N -dimensional space as your sensing basis. This can be the canonical basis, but also Hadamard, discrete cosine, or wavelet basis, to name just a few. All of these basis share two features. First, each one of the measurements takes information about several regions of the sample at the same time. Instead of asking about the intensity of one pixel, now we ask about several pixels at the same time (Fig. 2.3.c). This increases the SNR of our measurements. Second, the traditional way to recover the object is to solve the equation system given by Eq. 2.1. To do so, you calculate the inverse of the sensing matrix, and then perform the operation: $\mathbf{x} = \mathbf{S}^{-1} \cdot \mathbf{y}$. As any computer science expert can tell, calculating the inverse of a matrix with a big number of entries is quite inefficient, and if you want to work with high resolution images, you will soon start suffering from memory and speed limitations. Nevertheless, using these mathematical functions makes it possible to recover the object by utilizing fast digital transforms instead of calculating inverse matrices, providing a fast and efficient way of obtaining the final image. Last, and most important, the use of structured patterns allows the use of advanced signal processing techniques, such as CS, that greatly boost the performance of SPI systems, as we will see in the following sections.

This structured illumination scheme is what we usually name as a SPC². An example scheme of a SPC can be seen in Fig. 2.5. The main difference from a RS setup is the way of illuminating the scene. In a RS setup, you can use a pair of galvanometer mirrors to rapidly scan your scene (with speeds over tens of kHz). However, when performing structured illumination, more complex systems need to be used. This is where SLMs come into play. Instead of just scanning the scene, now the procedure entails the generation of all the functions of \mathbf{S} with the SLM (Fig. 2.3.c). This process is usually known as basis scan (BS), and can be done with phase-only or amplitude SLMs, depending on the kind of functions (complex or real) that you want to implement. In our experiments, we adopted DMDs for several reasons. First, their high refresh rates make it possible to build SPI systems with frame-rates comparable to the ones achieved by RS systems, which is crucial for practical applications (more information about how DMD technology works can be read on Appendix A). Also, as we will see in the following sections, fast BSs make it possible to take advantage of signal processing techniques that boost the acquisition speed of SPCs while maintaining a very affordable budget.

2.2 Going beyond matrix inversion

As we saw in the previous section, the process of obtaining an image in a SPC is quite straightforward. You choose your basis of functions (the rows of matrix \mathbf{S}). Then, you generate those functions with your SLM (a DMD in our case). After that, you either project those functions to the scene under study, or make an image of your scene into the SLM plane. In this process, both the functions and the scene overlap,

²Although here I am describing the SPC as a structured illumination system, there is an analogous design that does not project structured patterns onto the scene. As can be seen on Fig. 2.3.c, both the patterns and the sample need to physically overlap before collecting the resulting intensity value. However, this overlap can be done in different planes. Depending on the experimental conditions, sometimes it is very practical to image the scene onto the SLM plane, and make the superposition there. These two approaches are usually labelled as “structured illumination” and “structured detection”. During this thesis we will see examples of both.




FIGURE 2.5: **Single-pixel camera setup.** Left: a digital light projector (usually a DMD) generates a set of masks that are projected onto the scene. The light either transmitted or reflected by the object is collected onto a single-pixel detector. After that, the electrical signal is digitized and stored. Using the sequence of patterns projected by the SLM and the measured signal, the final image can be easily recovered. Right: reconstructions obtained with both a CCD camera and a SPC. Images adapted from [53].

and then you measure the corresponding light intensity with your bucket detector. To recover the object, you solve Eq. 2.1, for example, by inverting \mathbf{S} .

Without considering numerical problems (which can appear when the image size gets bigger or measurements are highly affected by noise), this method still presents a big challenge. Take, for example, the recovery of an image with a modest size (at least for today standards) of ~ 1 megapixel (i.e. 1024^2 pixels). To recover such an image, you need to generate a set of 1024^2 functions in your SLM. In our case, using a high-speed DMD module with a refresh rate of 22.7 kHz, this would take $1024^2/22700 \simeq 46$ seconds, which obviously is not fast enough for real-time applications. In fact, primordial SPC systems usually worked with low-resolution images, with sizes between 16^2 and 64^2 pixels.

So, it is evident that there is a trade-off between the number of pixels and the acquisition speed when getting an image using a SPC. However, there are some workarounds that allow us to recover a picture with a lower number of measurements, thus decreasing the measurement time. In this new case, the mathematical problem to solve will be

$$\hat{\mathbf{y}} = \hat{\mathbf{S}} \cdot \mathbf{x}, \quad (2.2)$$

where $\hat{\mathbf{S}}$ is now a rectangular matrix, with less rows than columns. This measurement process entails a dimension reduction, given that the size of our measurement vector, $\hat{\mathbf{y}}$, is lower than the size of our object, \mathbf{x} . This dimension mismatch means that the algebraic problem at hand is ill-posed, and thus it has an infinite number of solutions.

So, how can we retrieve our image in that scenario? Even though it may seem strange, ill-posed problems are quite common in both engineering and physics. For example, any first year physics student is quite used to obtain the trajectory that a

ball would follow when kicked by a football player (Fig. 2.6.a). To do so, they use Newton's laws of motion and a set of initial conditions such as the kick angle or the ball initial speed. This is an example of a direct problem. Using your model (Kinematics) and some parameters (speed, kick angle), you predict the effects (trajectory) of an event (the player kicked the ball). However, there is another interesting problem, maybe not so common among first year students, which is to find those parameters (initial speed, kick angle) from a set of positions of the ball trajectory. This inverse problem can be either ill-posed or well-posed, depending on the number of known positions of the ball. For example, if only the initial and final position are given, there is an infinite number of combinations of kick angles and initial speeds that would produce different trajectories with the same starting and ending positions (see Fig. 2.6.b). However, when adding additional ball positions the problem becomes well-posed (Fig. 2.6.c) and it is quite easy to find the right solution. As a rule-of-thumb, we classify problems as ill-posed when they have more than one possible solution (the more rigorous definition by Hadamard can be found in [56]).




FIGURE 2.6: **Examples of well-posed and ill-posed problems.** **a)** If friction is neglected, knowing the initial speed and the kick angle provides enough information to calculate the full trajectory of a ball kicked by a player (solid line). **b)** If only the initial and final positions are provided, there is an infinite number of combinations of initial speeds and kick angles that make the ball go from the starting to the ending position (dotted lines). **c)** If, in addition to the initial and final locations of the ball, the highest altitude point is given, the full path and the initial parameters can be deduced. **d)** Last, in the case of only knowing starting and ending positions, if we incorporate some assumptions about kick angles (around 45 degrees), we can discard some of the trajectories, obtaining a set of approximate solutions.

This family of ill-posed problems appears in mechanics, but also in optics, acoustics, information theory, astronomy, and medical imaging, among other fields. There are several strategies to overcome the obstacles they provoke. First, you can add additional measurements to make the problem well-posed, but this approach is often discarded if you do not have access to additional experimental data (because it was measured in the past) or it is quite expensive to acquire it (for example, when a patient stays on a computerized tomography scanner). Even if no more data is available, there is still a lot we can do to solve ill-posed problems.

Let's take the two-position-only football problem and think a little bit. Among all the possible speeds and angles, there are some that are quite more probable than others. For example, professional players can kick a ball at around 100 km/h. Then, it becomes evident that solutions with initial speeds much higher than that will be

probably wrong, and thus we can penalize (or directly discard) them. Also, if we know the total distance the ball travelled, we can also favour initial angles that make more sense. A very long shot will probably have initial angles around 45 degrees, which we know is the angle that provides longer travel distances with a given initial speed (Fig. 2.6.d).

This method of favouring some solutions and penalizing others using some extra information is mathematically known as regularization. In the following sections we will see some examples of how to retrieve images with a SPC using a reduced amount of measurements, regularization, similar approaches involving adaptive algorithms, and ideas extracted from recommendation algorithms used in media platforms such as Netflix.

2.2.1 Compressive sensing

Imagine you are playing a game of “Guess who?” (Fig. 2.7). For the ones not familiar with its mechanics, two players play against each other in turns, and the goal of the game is to guess a person the opponent chose from a pool of 24 people with different names. To do so, you ask yes/no questions to your rival, gathering the information you need to win. Of course, there are several strategies you can choose when playing this game. Probably the most naive one is to ask directly the name of a person to your opponent. In this scenario, you have a probability of guessing right of $1/24$ on your first try, then $1/23$, and so on. On average, using this strategy gets you the correct name with $24/2 = 12$ questions ($N/2$ for a pool of N people). However, there is a much better strategy that any kid would use from the start: instead of asking about one person, you can make generic questions that remove more than one person from the full pool of people. For example, if you play with a moderately modern version of the game, asking if the person you are looking for is a male will remove 12 people (either male or female) in just one question. You can follow up this with questions about skin or eyes color, removing more than one possible person each time. If you remove half the remaining people with each one of your questions, you need to make either 4 or 5 questions to end the game ($\log_2 N$ for a pool of N people).



FIGURE 2.7: **Guess who?** Tabletop game where two players try to guess a person by formulating yes/no questions to each other. After each question, the player can eliminate subjects that do not match the answer provided by the opponent.

At first glance, this may seem counter-intuitive. How can we find the solution of a 24-dimensional problem with only 4 or 5 equations (questions)? The answer is simple. Even though we are only measuring (asking) a few times (lower than the dimension of our problem), we have a lot of additional information. First, there is

only one possible solution. We also know that we are guessing people, and people have their own characteristic facial features we can use to produce very good questions (that remove half the pool of possible people each time). All of these concepts are at the heart of compressive sensing (CS) techniques.

For simplicity, imagine a “Guess Who?” game with only three people, each one with their unique name (Xavier, Yuri, and Zeno). In this case, we can say that our problem resides in a three-dimensional space, which I am going to call the *name-space* (see Fig. 2.8). This space can be visually represented using three orthogonal axes (x , y , and z), each one representing one name (or dimension). The solution we are looking for is a point in this three-dimensional space, with a position given by a vector, \mathbf{r} . For example, if we were trying to guess for Yuri, the solution will be $(0, 1, 0)$. In this space, the vector representing Yuri only has one non-zero element (the corresponding to its name axis). However, we can define other different spaces. Let’s define one axis that represents if the person is a male, another one that corresponds to the person hair length, and the last one representing height. So, tall people will lay on the positive axis of the height axis, people with hair on the positive part of the hair axis, etc. In this *body-space*, Yuri will be represented by the point $(1, 1, 1)$ because he was a tall man with nice hair.




FIGURE 2.8: **Vector sparsity in different bases.** In the left, each character (defined by its name) lies along one of the three axes. In the right, the same people does not have such a sparse representation, even though the information content is the same.

This shows a very useful property of data. Depending on the space we choose to work, the amount of information we need to use to define a person changes. In the *name-space*, we only need one number to define Yuri, whereas in the *body-space* we need three. This is commonly known as sparsity (we say that Yuri is sparse in the *name-space*). This idea is at the heart of compression techniques: if you want to store any data, express it in a space where it is very sparse, and then you will only need to store a reduced number of elements of your vector.

However, during the last decade, a small group of mathematicians (mainly Candès, Donoho, Romberg and Tao) made a good reasoning. The conventional process of obtaining and storing information is as follows. First, you acquire your data. Then, you choose a mathematical basis (space) to compress what you measured. After that, you transform your initial data into this basis, where it is sparse, and you only store the relevant parts (the most significant elements of your vector in this new space). Coming back to the last example, the process will be to measure in the *body-space*, transform your vectors to the *name-space*, and only store the non-zero elements of each person. When you want to reproduce your data, you can easily go back to

the original space, as going from one space to another entails just a linear operation. But, why doing it in this way? In order to compress something, first you need to measure all the information, and then you just throw away almost everything you measured. Why not measure the relevant data in the first place? This efficiency is what CS tries to accomplish [57, 58].

Recovering the formulation presented at the beginning of the chapter, our measurements can be expressed as

$$\mathbf{y} = \mathbf{S} \cdot \mathbf{x} = \mathbf{S} \cdot \mathbf{R} \cdot \boldsymbol{\alpha}. \quad (2.3)$$

The first equality is the same we saw on Eq. 2.1, but now with an additional relation. Here I supposed that our object under study, \mathbf{x} , can be expressed in a different basis, \mathbf{R} (usually called recovery basis), where it has a sparse representation (i.e., a lot of the elements of $\boldsymbol{\alpha}$ will be zero or very close to zero). CS shows that, if you make enough measurements, you can use an optimization algorithm to recover your object [14, 17, 59]. The number of measurements needs to fulfil that $m \geq \mathcal{O}[Q \cdot \log(N/Q)]$, where N is the dimension of your object and Q the number of non-zero elements of $\boldsymbol{\alpha}$. The fundamental idea is to find the sparsest object, $\boldsymbol{\alpha}^*$, compatible with your measurements. To do so, the most basic procedure consists of a regularization problem based on the minimization of the l_1 norm:

$$\boldsymbol{\alpha}^* = \arg \min \|\boldsymbol{\alpha}\|_{l_1} \text{ subject to } \mathbf{S} \cdot \mathbf{R} \cdot \boldsymbol{\alpha} = \mathbf{y}, \quad (2.4)$$

where the l_1 norm is defined as $\|\mathbf{x}\|_{l_1} = \sum_{i=1}^N |x_i|$. It is clear that, for vectors with a lot of entries set to zero, the l_1 norm will be low, so using this as a regularization function will favour solutions that are sparse. If you are wondering why choosing the l_1 instead of the l_0 norm (that would be the one giving the most sparse possible solution), the answer is that l_0 -minimization is a NP-hard problem, whereas l_1 -minimization is a linear problem that can be quickly solved [17, 59]. Once we have the result of the minimization problem, the object can be simply recovered with the expression

$$\mathbf{x}^* = \mathbf{R} \cdot \boldsymbol{\alpha}^*. \quad (2.5)$$

So, by using CS, we just need to make enough measurements (always lower than the total number of dimensions of our object), and then run a minimization algorithm. To obtain the best results, CS establishes that you need to choose the recovery space wisely (i.e., choosing a space where our object is very sparse), and also to measure using a basis which is incoherent with the recovery basis. The first condition is easy to meet, as natural data tends to be sparse. For example, natural images (which is the kind of data we are going to work with) are very sparse in many common mathematical bases, such as discrete cosines, Hadamard, or wavelets. Regarding incoherence, it is also easy to achieve. Incoherence can be understood as a measure of dissimilarity between the elements of each basis (any interested reader can go to [59] for a more rigorous definition). For example, in two dimensions, the canonical or pixel basis (formed by delta functions) is extremely incoherent with Hadamard functions (which seem like checker-board patterns). The same can be said for bases formed by random functions. The underlying idea of this incoherence requisite is the same that we saw in the “Guess Who?” example. In that case, we are looking for only one person, which is a solution extremely sparse in the *name-space* (represented by a vector with only one element different from zero). If we ask for one specific person, we only obtain information about one dimension of the whole *name-space* each

time. By measuring in another different incoherent space (making broad questions that eliminate half the pool of people), we get information about the full *name-space* in every measurement, and this helps obtaining a solution with a reduced number of measurements.




FIGURE 2.9: **Compressive sensing by l_1 -minimization.** Top left: ground truth object used in the simulation and its expansion coefficients in the Hadamard basis. It can be clearly seen that a great portion of the 128^2 coefficients have values very close to zero. Bottom left: Measurement process. The superposition of the object and a set of random functions provides the measurement vector shown in the graph, that can be used for recovery using the CS algorithm. Right: compressive recovery using a 50% compression ratio. It can be seen that we recover a very similar coefficient expansion, but with a higher number of elements set to zero due to the regularization process.

We can see a simple example of how this works when recovering an image in Fig. 2.9. In this case, the object under study is a binary Pacman representation with a size of 128^2 pixels. As the object is sparse on the Hadamard domain, we use random functions to obtain our set of measurements (50% of the total, in this case). In this way, we assure the incoherence between the basis where the object is sparse and our measurement basis. After the l_1 minimization is done, we get an image of our object with very good quality. The code used here was extracted from the l_1 -magic package [60], and it is one of the most simple strategies we can use to obtain images in a compressive sensing way. Even though this procedure gives good results, it presents several drawbacks, the main one being long reconstruction times. In the Pacman example, the recovery process took around one minute with a personal computer. For applications where real-time image acquisition is a must, other strategies can be explored [19, 33, 53, 61, 62]. Next, we will see two of them.

2.2.2 Adaptive strategies

As we saw in the last section, CS allows us to recover an object with a reduced number of measurements. In doing so, it recovers the sparsest object compatible with our measurements using iterative algorithms to solve a mathematically ill-posed problem. Even though this approach works well, those iterative algorithms require big

amounts of raw computing power. Also, it is quite common for these kind of algorithms to take longer times to run than the time it takes to make the measurements, which hinders their application in real-time scenarios. This problem also gets worse as image size grows (either by an increase in the number of pixels or by adding additional dimensions, such as wavelength or polarization).

In these scenarios, there is a trend to substitute those optimization problems by fast algorithms that can provide real-time reconstructions. For example, it is possible to sample the scene using orthogonal functions (Fourier, Hadamard, etc.) and recover the image by just adding each function weighted by the measured intensity (coefficient). This is very intuitive, as you get your image as a weighted sum of a basis of the two-dimensional space: patterns similar to your scene will have very big weights, and patterns different from it will have very low weights. An example of this can be seen in Fig. 2.10, where I simulated the measurement process of the Pacman's image using the first 25% of Hadamard functions (ordered by spatial frequency), and recovered the image by doing their weighted sum, which can be done on-the-fly even in low-end computers.




FIGURE 2.10: **Undersampled recovery by weighted addition.** Left: ground truth object. Right: recovery by only measuring the 25% lowest Hadamard frequencies and adding them weighted by the obtained intensity.

However, the problem when using this approach is choosing a good subset of the entire basis to measure. As we do not know which are the most relevant functions, when we perform the subsampling we always lose information (that is why such a simple image as the Pacman in Fig. 2.10 looks pixelated). Other approaches try to use some a priori knowledge about the object we are measuring. For example, one can take an image of the scene, and in subsequent frames, use the knowledge about the coefficient expansion of the scene to measure using only the most relevant functions [19, 26, 32]. In a work developed in 2015, we went a step further. Instead of capturing an image in a traditional way (i.e. measuring all the coefficients), and then using the spatial information to speed-up following acquisitions, we designed an adaptive algorithm that learns about the spatial content of the scene during its acquisition [53]. In order to learn those features, the algorithm uses wavelet transforms to locate regions of the scene with high density of borders, and assumes those regions are the ones where the relevant features are. This is nice for two main reasons: first, it does not need any a priori knowledge about the spatial features of the scene; and second, it only uses fast mathematical operations to retrieve the image (fast wavelet transforms and weighted sums). Before going through its operation principle, let me introduce a few relevant characteristics of wavelet transforms.

A wavelet transform can be intuitively understood as a sequence of filters. For a given N^2 -pixel image, we apply four filters to it. The first one is a low-pass filter in both dimensions (which we name LL). The next two are low-pass in one dimension and high-pass on the other (LH and HL). Last filter is a high-pass filter in both dimensions (HH). In this process we generate four images which size is one fourth of the original image, as can be seen in Fig. 2.11.




FIGURE 2.11: **Wavelet transform definition as spatial filters.** In the top we see the initial image (Steam logo image, property of Valve Corporation). In the bottom row, we can see the four images that result from applying four band-pass filters. This process can be repeated again to the low-pass filtered image (LL), generating a characteristic tree structure.

After the filtering process is done, it is common to generate the transformed image grouping all the levels in one image with the same size as the original one, as can be seen in Fig. 2.12. In this case we can see both the two-level and three-level wavelet transforms. There it is easy to see one of the most relevant characteristics of this transform. In the quadrants labelled H, V, and D we can see that the high-pass filter allows us to retrieve information about edges in the original scene. Also, regions where the scene is flat tend to be zero. This means that for images with a reduced number of edges, their wavelet expansion will be very sparse. Moreover, the pixel values will also give us information about the relevance of the edges in the picture, which we can use as a metric to discern between the most relevant regions of our object.

So, by using wavelet transforms we can locate the positions of the borders of a scene, but, how can this help reduce the number of measurements we need to take? The key resides in the tree structure of the transform. If you take a look, for example, at the three-level transform in Fig. 2.12, you can see that each level mimics the previous one. If we have a region where the coefficients are zero at the first level (painted blue), they will also be zero in the next levels (red and green). This property is used in compression algorithms based on wavelet transforms. First, a full transform of the object is done. After that, an algorithm goes along all the coefficients, and when their value goes beyond a predefined threshold, that coefficient and all




FIGURE 2.12: **Wavelet transform examples.** **Left:** Image sample (Steam logo image, property of Valve Corporation). **Center:** two-level wavelet transform. High values on the transform correspond to edges on the original object. Additional lines are shown to ease visualization of each one of the quadrants, that can be identified with horizontal, vertical and diagonal edges (H, V and D). **Right:** three-level wavelet transform. In this representation we can see the tree structure of wavelet transforms (red, green, and blue branches). At each level (also coloured), we get information with higher and higher spatial resolution (doubling it each time). The number of possible levels is delimited by the size of the original image.

their children coefficients (i.e., coefficients in the next levels that contain information about the same spatial region of the scene) are discarded. Then, only a handful of coefficients needs to be stored. Our adaptive algorithm is very similar in nature, but instead of starting from the full resolution scene, it starts from the low resolution image located the lowest level of the wavelet transform. Lets see how it works with an example.

We are going to simulate a measurement of the object shown in Fig. 2.13.a, which has a size of 256×256 pixels. The first step of our measurement process consists of obtaining a low resolution version of our object. We can choose any resolution we want, and this will determine the number of steps we need to follow to get our final image. In this example, we use the Hadamard basis of the 64×64 space. In order to measure, we resize all our patterns to fit the extent of the scene (256×256 pixels). In this stage we use the complete low-resolution set (64^2) of patterns. After this measurement, we get the coarse picture shown in Fig. 2.13.b. Once we have this coarse image, we calculate its one-level wavelet transform, as can be seen in Fig. 2.13.c. Then we inspect this transform to search for interesting regions. To do so, we divide each one of the three quadrants with the border information into four regions, and we add the absolute value of all the coefficients contained in each region (Fig. 2.13.d). This measurement defines the metric we will use to evaluate the information content of each spatial region. For regions with high density of borders, the sum of coefficients will be high. In contrast, flat regions will be populated with zero-valued coefficients. To classify regions in relevant or not, we just use a threshold defined as a percentage of the total energy of the scene. By doing this we can easily identify regions of the scene which are relevant (Fig. 2.13.e). In this case, in our first iteration we can discard one of the quadrants, so no more patterns will be projected in that spatial region again.

Next, we repeat this process for the regions that we still consider relevant (the three green quadrants in Fig. 2.13.e). In this second step, we get the full set of 64^2 Hadamard patterns and resize them to 128×128 pixels. After that, we get coarse pictures of the three relevant quadrants we classified on the first step, and repeat the process of looking for relevant regions with the aid of a wavelet transform (Fig. 2.14).




FIGURE 2.13: **Initial step of the adaptive recovery process.** **a)** Scene under study, a sample of *Fasciola hepatica* with a size of 256×256 pixels. **b)** Initial coarse observation of the scene, obtained with a set of 64×64 Hadamard patterns. **c)** One-level wavelet transform of the coarse image. Here, it is possible to see three quadrants that contain the border information of the scene. **d)** Border inspection on the wavelet transform. Each one of the quadrants has been divided into four zones (coloured). Each color is linked to the same spatial region of the scene. By adding the values of the coefficients in each region, the amount of borders can be calculated, and then each zone can be defined as relevant or not. **e)** After the transform inspection, the algorithm locates three quadrants (green) with relevant information, and one without. The three relevant zones will be further sampled in the next steps of the algorithm, while the remaining one will stay with the current sampling level.




FIGURE 2.14: **Second step of the adaptive recovery process.** With the information we gathered on the initial step, now we measure in the three remaining relevant regions (coloured). For each one, we acquire a coarse image and, again, compute its one-level wavelet transform. Then, the process of dividing the border quadrants into four regions is repeated, and the total amount of borders is computed. This refines the location of the relevant parts of the scene.

This provides us the last set of relevant regions. In the last step of our algorithm, we use the 64^2 Hadamard patterns (now without resizing) to sample the last batch of zones of interest (the six green regions highlighted in Fig. 2.14). After that, we can compose a full three-level wavelet transform (Fig. 2.15) with all the information we have acquired. By doing a simple inverse transform we retrieve our 256^2 pixel image. In this case, the total number of measurements was $10 \times 64^2 = 40960$, which is roughly a 62% of the total number of pixels of the scene (256^2). In addition to this reduction in the number of measurements, the only mathematical operations needed to recover the object are matrix multiplication, addition, and wavelet transforms. As all those operations can be computed extremely fast, the recovery process can be performed on the fly, as opposed to conventional CS approaches.




FIGURE 2.15: **Final step of the adaptive recovery algorithm.** After all the relevant regions have been measured, the full set of wavelet transforms can be composed into a three-level wavelet transform. By doing a simple inverse transform, the final image is retrieved. It can be seen that with only a number of measurements equal to 62% of the number of pixels, the quality achieved is perfect on the regions where the cells are located, while it is lower on the background.

Different sampling pattern sizes and number of levels can be chosen, depending on the target spatial resolution and size of the final image. In that case, the second step of our example is repeated over and over until the pixel size of the patterns we use to sample the scene is equal to the pixel size of the scene. As an experimental proof of the technique, we performed an experiment with a SPC as the one presented in Fig. 2.5. To test the compression gains of our proposal, we imaged two different objects: a LEGO® minifig and an USAF resolution target. We acquired images of 256×256 and 512×512 pixels, with variable compression rates. In both cases, the

finest details of the scenes could be achieved with undersampling, and the recovery process was practically instant.




FIGURE 2.16: **Adaptive recovery experimental results.** Top row: 256×256 LEGO® Ned Flanders image (left) and the recovery with the adaptive strategy using a number of measurements equal to the 88% of the number of pixels. Bottom row: 512×512 USAF1951 test (left) and the recovery using the adaptive approach. In this case, the number of measurements was 55% of the total number of pixels. Figure extracted from [53].

This adaptive approach presents several benefits. First, it presents real-time recovery. Second, and more important, given the initial parameters, it adaptively decides the total number of measurements (as opposed to a more *democratic* CS approach, that treats all scenes equally). Those two characteristics make it a very good approach when trying to recover information at video rates. However, the technique presented here was only based on the spatial features of the scene. When one wants to acquire multidimensional information, for example in the case of multispectral imaging, other approaches that play with the sparsity in the additional dimensions can be explored. We will see an example next.

2.2.3 Matrix completion

The rise of streaming media platforms has been fuelled by a high number of technological advances. For example, even though movies and TV shows were a common pastime for lots of people decades ago, building an accessible database for millions of users with tens of thousands of titles required both storage and bandwidth capacities that could not be met until recently. The same can be said in the case of music, where even if each song requires very little storage, the huge amount of tunes produced every day makes it very difficult to create such a platform. Reducing the storage needs of each item (whether it be a movie or a song) by using compression techniques (everyone today is familiar with formats such as *mp3* or *avi*) and the tremendous increase of available bandwidth have made possible to build platforms

	Die Hard	Mad Max	Commando	Blade Runner	Moon	Primer
User #1	5	-	5	1	2	-
User #2	1	2	-	5	5	-
User #3	5	-	5	-	-	3
User #4	3	3	3	5	-	1

TABLE 2.1: **Film rating examples for a small group of users.** For each film, ratings can go from 1 (bad) to 5 (very good). Not watched films are marked with a "-". The first three films could be classified in the *action* genre, while the rest could be considered *science-fiction*.

such as Netflix or Spotify, where you are just a click away of watching or listening whatever you like.

However, even if the theme seems to suit this chapter nicely, we are not going to talk about storage or transmission problems. There is another interesting aspect of streaming platforms that some of you may have experienced. As a big fan of *Queen*, some time ago I went looking for *Fat Bottomed Girls*. After hearing it in Spotify, I listened to several more *Queen* songs. Some days later, the software gave me a list of songs that maybe I would like. Among others, the list contained songs from *Boston*, *Led Zeppelin*, and *Deep Purple*. I like all of these bands, so I listened to the playlist, skipping the songs I did not like. Days passed by and I got new recommendations. This time I got songs from a lot of bands (both known and unknown), and the number of songs I did not like was much lower. This kept on going, and I have used this feature every week for a long time to discover music I enjoy listening to.

How can Spotify make good recommendations to its users? In fact, all the streaming services I know do these kind of recommendations. How do they do it? In the past, people relied on suggestions by other people to discover new things. You got books from your friends or your local book store. You went to the movies when someone told you that a film was good. The same with music. Of course, if you like the same music your friends listen, their recommendations will suit you well. After all, they know you and your tastes. However, how do you get an algorithm to provide good recommendations?

This problem was popularized by Netflix some years ago. They wanted to build the best recommendation system possible for their platform, so they offered a one million dollar prize to anyone who could produce a system with better performance than theirs. *The Netflix problem* became extremely well known, and teams from all over the world competed generating better and better algorithms.

In order to build a system like this, we have to understand how the information is stored and what relations it holds. Let's start with a simple example with some film titles and a short list of people (see Table 2.1).

There are several patterns any cinema aficionado can see there. To start, we can clearly divide the films into two main genres: *action* and *science-fiction*. Also, we can classify users here with regard to their preferences. Some will be *action* fans, and thus will tend to rate those films high (as user #1). Some other may be *science-fiction* fans (user #2). Others may be a mix of both, liking or disliking both genres. Adding more users and film titles will expand the categories shown here (adding genres and combination possibilities for different user types), but the basic idea is clear: there is an underlying structure on the ratings that can be used to reduce the dimensionality of the dataset.

For example, we know that user #1 likes *action* films, and *Mad Max* is shown as *not watched*. If we recommend that film, there is a high chance that the rating will be high. Also, we can see that the ratings of *Blade Runner* and *Moon* are low, so we should not recommend *Primer*. Similar predictions can be done for the rest of users. The bigger the database we can get, the better predictions we should be able to do. Moreover, in this example films are just divided into genres, but nothing prevents us to divide into other categories, such as “*the film has classical music in its soundtrack*” or “*the film has a female lead*”, which could be patterns some users like/dislike.

The thing is that, in this example, the ratings are predicted based on my experience and the theory that users like films of the same genre. However, lots of questions can be made about this election. What are good categories to divide users into? The example is based on genres, but is that better than “*female lead*”? Should we use both at the same time? One can clearly see that as soon as we start refining our method by increasing the number of parameters to take into account, this becomes an impossible task for a human. Matrix completion (MC) techniques try to solve this kind of problem in a computational way.

Of course, the problem of filling matrix entries is ill-posed, as there is no a priori relation between its elements. However, if the rank of the matrix is known, the problem can be solved if we know enough entries of the matrix, as can be seen on Fig. 2.17.

1						1	5	-1	1	0	2
		-1				1	5	-1	1	0	2
				0		1	5	-1	1	0	2
	5					1	5	-1	1	0	2
					2	1	5	-1	1	0	2
			1			1	5	-1	1	0	2

FIGURE 2.17: **Matrix completion of an incomplete rank one matrix.** As we know the matrix rank is one, all the rows of the matrix can be easily related, and even though not all the entries are known, it is trivial to fill in the gaps.

Even in the case of unknown rank, it can be demonstrated that we can retrieve a good approximation of the matrix if it is low rank and some assumptions are done [63–65]. This is known as low-rank matrix completion (LRMC), and its one of the most common techniques to solve the recommendation problem we saw before, because user preferences can usually be described with just a few parameters, or in other words, the complete matrix without missing ratings is either low-rank or can be well approximated by a low-rank matrix.

However, if the topic of this thesis is not recommender systems, why did we start to talk about them in the first place? The answer is simple: we can easily take advantage of the techniques developed in MC problems to obtain multidimensional images in a fast way. The reason behind this is easy to understand. In the case of films, MC works because the true dimensionality of the dataset is lower than the one we see at first glance. This is also true in many practical scenarios of multidimensional imaging. Let’s see show it works with a simple example in multispectral imaging.




FIGURE 2.18: **Redundancy example in multispectral imaging.** **a)** Spontaneous Raman image of a polystyrene bead and a silicone oil droplet suspended in water (pseudocolored), with their respective spectral signatures. **b)** Hyperspectral cube expressed in matrix form. Each spatial position corresponds to a row of the matrix, and the columns contain the spectra for each pixel. It can be clearly seen that a high number of rows are similar in structure, given the fact that several spatial positions contain the same chemical species. This becomes evident when the SVD of the matrix is done. In this case, the singular values after the first three ones are almost negligible.

Consider a multispectral image as the one depicted in Fig. 2.18. In this case, the data was acquired using a spontaneous Raman microscope [66], and the full hyperspectral cube has a dimension of $32 \times 32 \times 128$ voxels (32×32 spatial positions and 128 spectral channels). If we take the full datacube and rearrange it so the spectrum of each spatial position is a row of a matrix, we get the matrix shown in Fig. 2.18.b. What is interesting here is that, if we perform the Singular Value Decomposition (SVD) of this matrix (see Appendix B for more information about SVD), we can clearly see that the singular values rapidly decay. The physical meaning of this distribution is that there is a reduced number of distinguishable spectral species in our sample (three in our case). So, for each spatial position, we either have one of the present species in the sample, or a combination of them (as in the first example with film ratings). Given this big redundancy, we can just sample a small amount of entries of the matrix, and use MC techniques to recover the full spectral cube, increasing our measurement speed. In Fig. 2.19.a we can see the reconstruction we get when measuring a random 50% of the entries of the spectral matrix and using MC techniques to fill in the missing values.

In this example, we are using a simple MC algorithm based on the thresholding of the SVD of a matrix, developed by Candès et al. [64]. One of the benefits of this kind of approach is that, as the tools were developed with big datasets in mind, they are quite efficient from the computational point of view and, what is better, really fast for big datasets. This is a real advantage, as the datasets acquired in multispectral imaging (that typically have thousands of spatial positions and hundreds of




FIGURE 2.19: **Multispectral imaging via matrix completion.** For the same polystyrene and silicon sample, here we can see the image recovered by only measuring half of the entries of the hyperspectral cube using MC.

spectral channels) are quite hard to tackle with CS approaches [67]. Also, the output of the algorithm is directly the representation of the data in a low dimension representation (the matrices \mathbf{U} , \mathbf{V} , and $\mathbf{\Sigma}$ from the SVD of our spectral matrix), which is very useful in imaging spectroscopy, where this representation naturally provides the information about the chemical species present in the sample.

With this, we have seen the three main computational techniques I have used along the thesis to improve the performance of SPI systems. However, it is also possible to play with the experimental configuration of a SPC to measure faster and with a better signal-to-noise ratio in challenging scenarios. We will see how next.

2.3 Coping with the limitations of single-pixel imaging systems

In the previous section, we saw the fundamentals of the SPC and some of the computational techniques that can be used to boost its performance. However, software improvements are not the only way to improve a SPC. Of course, SPI systems are subject to improvements in both single-pixel detectors and SLMs. By creating faster and more sensitive sensors, the amount of time spent measuring decreases. In the same way, using faster SLMs naturally increases the frame rate of SPCs.

One of the choices you make when designing a SPC is the SLM. In our case, we work with DMDs because they are fast and cheap. However, they have some limitations. The main one is that they are binary intensity-only modulators. As described before, in order to get an image with a SPC we need to generate a set of functions with our SLM. So, by choosing a DMD we are inherently subject to measure in binary real bases. The binary limitation can be bypassed using either spatial or temporal multiplexing (see Appendix A). In order to measure with negative-valued functions, temporal multiplexing can also be used. Take, for example, the case of a Hadamard matrix (that is one of the most used sets in SPI). Hadamard matrices are binary square matrices valued either $+1$ or -1 . Their rows form an orthogonal basis in which we can measure the projections of our object, which provides the experimental benefits stated in Sec. 2.2.2. In order to codify negative values, we can take into account that the measurement process:

$$\mathbf{y} = \mathbf{H} \cdot \mathbf{x}, \quad (2.6)$$

is linear, and thus can be expressed as a linear combination of two measurements. Also, given the properties of Hadamard matrices, we can write

$$\mathbf{H} = \mathbf{H}^+ - \mathbf{H}^-; \quad (2.7a)$$

$$\mathbf{H}^+ = \frac{1}{2}(\mathbf{I} + \mathbf{H}); \quad (2.7b)$$

$$\mathbf{H}^- = \frac{1}{2}(\mathbf{I} - \mathbf{H}); \quad (2.7c)$$

where \mathbf{I} is a matrix with all its elements set to 1. It can be clearly seen that both \mathbf{H}^+ and \mathbf{H}^- entries are just 1 or 0. By doing this, we can describe the measurement as a two-step process: $\mathbf{y} = \mathbf{H}^+ \cdot \mathbf{x} - \mathbf{H}^- \cdot \mathbf{x}$. To measure one coefficient ($y(i)$), first you project one Hadamard pattern (one row of \mathbf{H}^+), and then its complementary (the corresponding row of \mathbf{H}^- , switching *ones* by *zeros*). After acquiring the electrical signal and converting it into a digital value, you make the subtraction in your computer and get the Hadamard coefficient. This process has been widely used in SPI systems to increase the quality of recovered images [19, 36, 50, 54, 68, 69].

However, it is possible to perform this kind of measurements without reducing your acquisition speed by a factor of two. In a work developed in 2016, we implemented this measurement procedure taking advantage of the dual nature of DMDs and the performance boost that balanced detection provides [54]. The main idea behind this work is to simultaneously use both reflection arms of the DMD to measure the \mathbf{H}^+ and \mathbf{H}^- projections. Also, as the relevant information we want to acquire is located in the difference between those two measurements, a balanced detector provides that exact value with a higher dynamic range and sensitivity than a conventional SPI system.

The experimental setup can be seen in Fig. 2.20. It is a typical SPC working in a structured detection mode, i.e., imaging the scene onto the DMD, where the patterns are generated, and measuring the reflected light with a photodiode. The twist here is to use both reflection directions of the DMD. For each arm, we get the superposition of the scene with one pattern and its complementary one. Now, instead of using two different photodiodes to get both signals and performing the subtraction after the digitization phase, we introduced a balanced detection scheme to increase the performance of the system.

Balanced detection has been widely used in optics to remove background signal and suppress random intensity variations, enabling the measurement in shot-noise limited conditions [70–73]. The fundamental idea behind this implementation is that any signal from a SPI experiment can be divided into two parts: a DC and an AC term. The DC term will be constant, and in our case is related to the total amount of light coming from the scene. On top of this DC term, we can distinguish an AC term that will be caused by the small oscillations of the signal that will be caused when we change the pattern on the DMD and different parts of the object overlap with the bright regions of the patterns. The key here is that the relevant information to obtain an image is contained in the AC term. When using a single photodiode, you convert the whole (DC+AC) signal to digital, and then inspect the signal using a computer. Here, we are using two photodetectors connected in such a way that their photocurrents cancel each other. This directly removes the DC term, and after this difference is done (in the electrical domain), the differential signal can be amplified and digitized using the whole dynamic range of our analog-to-digital converter. By doing this we get a superior SNR, improve the frame rate of our system in a factor of




FIGURE 2.20: **Single-pixel imaging using balanced detection.** A white light lamp (WLS) illuminates the object under study (USAF test), which is imaged into the DMD plane by using a single lens (L_1). The DMD generates the Hadamard patterns, and light is collected in both reflection directions using two collecting lenses (CL) and two optical fibers, that guide the light into the entrance of a balanced photodetector (BPD). After amplification, the signal is digitized with an analog-to-digital converter (ADC) and stored in the computer (PC), which controls all the process. An additional light source (ALS) is used to introduce ambient light into the system in a controlled way using a beam splitter. Figure adapted from [54].

two, and also make it immune to temporal oscillations of the light source or parasite signals such as ambient light.

We can see an example of this in Fig. 2.21. In this experiment we used the setup shown in 2.20 to acquire a set of images using two different configurations. First, we acquired them using balanced detection (Fig. 2.21.a). To show its benefits, we introduced parasite light using an additional light source into our system. It can be clearly seen that even when the signal level is much dimmer (by a factor of two) than the parasite signal we can clearly recover our object. Also, the difference between measuring with this parasite light source is minimal. When using a traditional SPI system with only one photodetector (Fig. 2.21.b), we see that artefacts are present in all signal levels, reducing the capability of the system to obtain a good image. Also, even when the parasite lamp is turned off, the reduced dynamic range of the system causes the quality of the images to be generally lower than their balanced detection counterparts.

With this example we end with the tools I have worked with to improve the general performance of SPI systems, whether it be on the computational or on the hardware side. In the following chapter we will go through the multidimensional systems that can be developed using these tools to obtain images with information about wavelength, polarization and phase.




FIGURE 2.21: **Single-pixel imaging with parasite light.** Group of images with different illumination conditions for a balanced detection (a) and a single detection (b) setup. While maintaining the power level for the ambient light source shown in Fig. 2.20, multiple images are acquired with different illumination levels over the object. We also show the images acquired without introducing parasite light into the system (marked as “lamp OFF”). Figure adapted from [54].

Chapter 3

Experimental results: putting it together

After introducing both the computational and hardware fundamentals of our systems in the previous chapter, we are ready to see some of the applications I have worked on during the last years. I will start with a SPC where the detector is a fibre spectrometer [50]. This allows to obtain multispectral images. After that, adding several polarimetric elements to that design makes it possible to obtain both spectral and polarimetric information simultaneously [51]. With the development of a different detection scheme and using MC techniques, it is also possible to perform multispectral imaging in a Raman microspectroscopy environment. Last, I will demonstrate how to obtain phase information using a lateral position detector as a single-pixel detector [52].

3.1 Spectropolarimetric imaging using a single-pixel camera

Multispectral images are of paramount interest when studying the chemical structure of objects. For example, as can be seen at the beginning of this thesis, this technique can be used to study cellular biology [74], to identify the components of soils by means of aerial imagery [75], or as a tool to measure temperature [76], among others. There are quite a few different approaches when tackling the problem of obtaining multispectral images (an interested reader can go to [77] for a detailed review), but almost all of them struggle in a couple of points. First, using a single array detector in the visible (VIS) region of the spectra is usually quite easy, but when trying to recover information over a wide spectral region, for example in the infrared (IR), costs can be prohibitive or detectors can just be unavailable. Also, when both spatial and spectral resolution start to grow, the size of the datasets increases in such a fast way that both transmission and storage problems start to limit system performance.

To solve these problems, SPI systems offer two main solutions. First, using bucket detectors greatly simplifies the problem of detecting over a wide spectral range, as photodiodes can easily work in the full VIS, IR range, and even in other exotic regions such as THz or X-ray [25, 30, 78, 79]. Also, their size and simplicity makes it even possible to combine multiple detectors for different spectral ranges without needing very complicated systems [21, 26]. Second, the use of SLMs makes it possible to use signal processing techniques, such as CS or MC, to reduce the size of acquired datasets, alleviating both storage and transmission problems.

In a work developed in 2013 [50], we built a simple multispectral system using a fibre spectrometer as a bucket detector. The device is depicted in Fig. 3.1.a. In this case, a white-light lamp illuminates the scene, which is imaged onto the surface of

the DMD. In this plane, we codify the patterns that are used in the SPI reconstruction process (as explained in Chapter 2). After this, an optical fibre guides the light to a beam spectrometer, where its spectrum can be analysed. In this way, we can measure the spatial projections of our scene into our measurement basis for several wavelengths at the same time.




FIGURE 3.1: **Single-pixel multispectral camera.** A white-light source illuminates the scene, which is imaged onto the plane of the DMD. After the structured patterns are generated onto the DMD, the light reflected is concentrated into the entrance of a fiber spectrometer, thus obtaining spectral information (a). If light goes through an analyser and a fiber spectrometer, both polarization and spectral data is gathered (b). Adapted from [51].

By doing this, instead of measuring a single vector \mathbf{y} as in Equation 2.1, we get as many vectors as spectral bands are measured by our spectrometer. Then, we can solve an equation system for each spectral band, obtaining the images composing the spectral cube. Of course, signal processing techniques can be applied to this problem to improve both speed, storage, and bandwidth needs. In a first experiment, we acquired the spectral cube shown in Fig. 3.2 using CS. For a spectral cube with a total size of $256 \times 256 \times 15$ elements, we only used a 10% of the full set of Hadamard functions to sub sample the spatial features of our scene. The object, a couple of cherry tomatoes in two different maturation stages, clearly shows different spectral information along the field of view in both the VIS and NIR region.




FIGURE 3.2: **Multispectral imaging using a SPC.** Cherry tomatoes in two different maturation stages in fifteen different spectral bands going from the VIS to the NIR. Adapted from [50].

In addition to spectral information, it has been demonstrated that adding polarimetric data improves the capability to perform biological analysis of living samples [80, 81] and the results of remote sensing [82, 83], among others. As a first approach

to the problem, a simple linear polarizer can be added to the system to act as an analyser. The modified setup is shown in Fig. 3.1.b., where the polarizer was added just in front of the optical fiber of the spectrometer. The scene under study consisted of two small capacitors (one yellow and one blue), each one illuminated with a different (horizontal or vertical) linear state of polarization (SoP). Now, for four different orientations of the analyzer, a set of spectral images can be acquired. In this case, it is straightforward to derive the spatial distribution of three Stokes parameters (S_0 , S_1 , and S_2).




FIGURE 3.3: **Preliminary polarization analysis of multispectral imagery obtained with a SPC.** Multispectral data of a scene for four different orientations of a linear polarizer. With this simple procedure, the Stokes parameters S_0 , S_1 , and S_2 can be easily deduced. Extracted from [50].

However, this setup has limited capabilities. To start with, in order to get polarization information you need to physically rotate the analyser. The presence of this moving element is detrimental in ranged detection systems, and basically in any setup which purpose is to be used in a real-life scenario. Also, by rotating a simple analyser you only get partial information about the SoP of light (you miss the information contained in the Stokes parameter S_3). In order to solve these problems, we developed a new system with a more sophisticated detection stage [51]. Now, instead of just using a simple analyser, we also added two electronically-controlled liquid crystal variable retarders (see Fig. 3.4).

The setup is basically the same as the one in Fig. 3.1, but now the light passes through a fixed linear polarizer with its transmission axis at 45° placed after two additional polarizing elements: two variable retarder plates with their slow axes oriented at 45° and 0° , respectively. By using Mueller calculus it is easy to prove that the SoP at the entrance of the fiber (\mathbf{S}') can be related to the SoP before the LCVR plates (\mathbf{S}) by:

$$\mathbf{S}' = \mathbf{A} \cdot \text{LCVR}_1 \cdot \text{LCVR}_2 \cdot \mathbf{S},$$

$$\begin{pmatrix} S'_0 \\ S'_1 \\ S'_2 \\ S'_3 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & \sin 2\delta_1 \sin 2\delta_2 & \cos 2\delta_2 & -\cos 2\delta_1 \sin 2\delta_2 \\ 0 & 0 & 0 & 0 \\ 1 & \sin 2\delta_1 \sin 2\delta_2 & \cos 2\delta_2 & -\cos 2\delta_1 \sin 2\delta_2 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cdot \begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix}, \quad (3.1)$$




FIGURE 3.4: **Single-pixel multispectral Stokes imager.** A white light source illuminates the object (a stressed piece of polystyrene), which is imaged by a relay system onto the surface of the DMD. Light reflected by the DMD goes through the LCVR plates and the analyser, and is collected onto the entrance of a fiber spectrometer.

where $2\delta_1$ and $2\delta_2$ are the phase retardances introduced by the LCVR plates. As both δ_1 and δ_2 are wavelength dependent, a previous voltage-to-retardance calibration along the full spectral needs to be done. After the SP reconstruction, we obtain the distribution of irradiance of the scene (S_0), and we can relate it with the full Stokes distribution of the object using Eq. 3.1:

$$S'_0(2\delta_1, 2\delta_2) = \frac{1}{2}[S_0 + \sin(2\delta_1) \sin(2\delta_2)S_1 + \cos(2\delta_2)S_2 - \cos(2\delta_2) \sin(2\delta_2)S_3]. \quad (3.2)$$

Using four different retardance combinations ($\{2\delta_1^i, 2\delta_2^i\}_{i=1,\dots,4}$), we acquire four different images for a given spectral band. After that, we can use the following equation system to obtain the full Stokes information for each spatial position of the scene:

$$\begin{pmatrix} S_0^1 \\ S_0^2 \\ S_0^3 \\ S_0^4 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & \sin 2\delta_1^1 & \sin 2\delta_2^1 & \cos 2\delta_2^1 & -\cos 2\delta_1^1 \sin 2\delta_2^1 \\ 1 & \sin 2\delta_1^2 & \sin 2\delta_2^2 & \cos 2\delta_2^2 & -\cos 2\delta_1^2 \sin 2\delta_2^2 \\ 1 & \sin 2\delta_1^3 & \sin 2\delta_2^3 & \cos 2\delta_2^3 & -\cos 2\delta_1^3 \sin 2\delta_2^3 \\ 1 & \sin 2\delta_1^4 & \sin 2\delta_2^4 & \cos 2\delta_2^4 & -\cos 2\delta_1^4 \sin 2\delta_2^4 \end{pmatrix} \cdot \begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix}, \quad (3.3)$$

$$\mathbf{S}'_0 = \mathbf{M} \cdot \mathbf{S},$$

and repeating this procedure for all the spectral bands, we acquire the full spectro-polarimetric dataset.

As an example of the procedure, we measured the spectral Stokes distribution of a stressed piece of polystyrene. The stress is generated in the process of fabrication, as the plastic is shaped in a predefined form. This stress causes the plastic to present a spatial distribution of birefringence that can be easily seen when the piece is placed between two crossed linear polarizers and illuminated with white light (see Fig. 3.4). In this case, we used a 20% of the total set of 128×128 Hadamard patterns and

CS recovery to obtain the spatial Stokes distribution for eight spectral bands with a width of 20 nm. The results can be seen in Fig. 3.5.




FIGURE 3.5: **Imaging the spectral Stokes distribution of an object using a SPC.** For a small stressed polystyrene piece, the spatial distribution of the Stokes parameters over eight spectral bands can be deduced using a detector built by combining a fibre spectrometer, a linear polarizer, and two LCVR plates. Adapted from [51].

3.2 Computational Raman microspectroscopy

In the previous section, we saw how to use a SPC to obtain multispectral and polarimetric information. In those experiments, a DMD is used to sample the spatial features of the scene under study, while different detection schemes provide the spectral and polarimetric information. However, this way of obtaining multidimensional information is not unique. In this section, we will see a different approach (but with a lot of similarities) to obtain multispectral images in a Raman microspectrometry setup.




FIGURE 3.6: **Dye microscopy examples.** On the left, a microscopic view of a sample of human lung tissue, stained with hematoxylin and eosin (image extracted from [84]). On the right, a composed view of human cheek cells when viewed in a wide-field microscope or by using the fluorescence signal generated when staining the cells with carbon nanoparticles (image adapted from [85]).

When studying biological organisms, it is crucial to obtain images where the internal structure and components of the sample can be easily distinguished. In common scenarios, the specimens under study tend to be transparent, which complicates the imaging process. One of the most common ways of solving this problem is to artificially label the samples to make them easily visible by a traditional optical system, in a process commonly known as staining. By using different dyes and

adequate procedures, it is possible to attach different dyes to different regions of the sample, enhancing image contrast. This process can also be applied using fluorescent dyes, which enhances the depth penetration of microscopy techniques (see Fig. 3.6). However, this methodology has several limitations. First, it is invasive, as the dyes need to physically attach to different structures. The protocols used for dye fixation can damage or sometimes even kill the specimens, limiting the number of applications of the procedure [86]. Second, the number of available dyes is limited, and there are quite a lot of biological structures that do not have known usable dyes. Even though current research is being done in both the fixation procedures and the dye generation, working with techniques without needing to label the samples would make everything easier.




FIGURE 3.7: **Raman effect fundamentals.** **a)** Energy levels involved in several scattering processes. If the photon interacting with the media maintains its energy, we say that the scattering was elastic (Rayleigh scattering). However, when the photon energy changes, we have inelastic (Raman) scattering. In this case, the photon can increase its energy (anti Stokes scattering) or decrease it (Stokes scattering). **b)** Typical Raman spectrum for an interaction between a photon and a molecule. The central peak that goes off the chart corresponds to the wavelength of the incident photons and the ones that suffer from Rayleigh scattering. We can also clearly see that several peaks at different wavelengths appear, which positions depend on the energy level structure of the molecule.

This is where Raman microspectroscopy (RMS) comes into play. RMS is based on the spontaneous Raman effect. When light arrives to an atom or molecule, it can either be absorbed or scattered. Among the scattered photons, most will retain their original energy (Rayleigh scattering), but a small fraction will suffer from inelastic scattering, thus varying its energy (Stokes and anti-Stokes Raman scattering). The energy exchange between the photon and the atom excites the latter, usually to a vibrational level (see the Stokes peaks in Fig. 3.7.b). This is interesting from an imaging point of view: as the chemical species present in a biological sample have different structures, their vibrational energy levels will be distinct, and thus also

their Raman spectra. Given the structural complexity of chemical components, this spectrum is, indeed, a molecular fingerprint that can be used to naturally label any sample. This is a strong point for RMS, as there is no need to generate dyes or to prepare the specimens for imaging. However, the small probability of Raman scattering events (only around one photon in ten million suffers from Raman scattering) hinders its application in real-life scenarios due to very low signal levels. To tackle this, Raman imaging systems usually require high excitation powers, but even in this case the imaging speeds tend to be low [86, 87]. This limits the application of RMS in the study of live biological samples. Also, given the multispectral nature of the technique, data storage and bandwidth bottlenecks are common problems that need to be faced. Even though much stronger Raman signals can be obtained using non-linear processes (as in Coherent Anti-Stokes Raman Spectroscopy, CARS) [87], some novel computational tools have started to be used as a way to increase RMS performance.

In a recent series of works, it has been demonstrated that using a single-pixel computational spectrometer instead of a traditional one using an array detector provides additional flexibility in the acquisition and post-processing stages without a big sacrifice in complexity [66, 88]. Moreover, using a single-pixel detector in combination with a DMD results in an inexpensive setup that enhances sensitivity at low signal levels.




FIGURE 3.8: **Computational Raman microscopy setup.** Light from a laser is focused onto a sample, generating a characteristic Raman spectrum. The sample can be moved using a 3D stage, thus obtaining spatial information. The spectral signal is measured using a computational spectrometer based on a DMD and a PMT. Captions: LS: 532 nm laser source, LLF: laser line filter, NF: notch filter, Di: dichroic mirror. Image adapted from [66].

In contrast with the single-pixel architectures presented before, in this application the DMD is not used to sample the spatial features of our sample, but its spectral content. The setup of the system can be seen in Fig. 3.8. A green laser beam ($\lambda = 532$ nm) is focused onto the sample using a microscope objective. The sample can be displaced in three dimensions using a motorized stage (see Fig. 3.9.a). After

illumination, the generated Raman signal goes through a grating, separating all its spectral components. The diffracted light is guided to the surface of a DMD, where it can be either reflected onto the surface of a single-pixel detector or in a non-guided direction. By doing this, the system is able to obtain spectral information in different modes. First, the DMD can act as a programmable slit, thus measuring a single different wavelength at a time. After all the wavelengths have been measured, the Raman spectrum can be recovered by simple numerical methods (at the expense of a reduced light efficiency). However, the signal level at the detector can be increased by measuring multiple wavelengths at the same time by coding spatial masks on the DMD (Fig. 3.9.b). In this case, the measurement would be analogous to the one in a Hadamard transform spectrometer [89, 90].




FIGURE 3.9: **Sampling fundamentals of a computational Raman microscope.** **a)** The spatial information of the sample is obtained by raster scanning the scene. For each spatial position, a focal spot generated by a microscope objective excites the sample to generate the Raman spectrum, that is measured with a computational spectrometer. **b)** The Raman signal passes through a grating, so each wavelength is mapped to a different spatial position in the DMD. Then, the DMD is used to sample the spectrum, either by acting as a simple moving slit (like in a monochromator) or by implementing structured masks (as is done in the spatial domain with Hadamard functions, for example) and measuring with a single-pixel detector (for example a PMT or a SPAD).

All the sampling fundamentals introduced in the previous sections for the spatial domain can be translated to the measurement of spectral information. In this way, raster scan, basis scan, and also advanced signal processing techniques can be used to obtain the spectral data, while the spatial raster scan provides the features of the sample in two or three dimensions. Recently, this computational approach to Raman microspectroscopy has sparked some interest as a way of increasing acquisition speed and soften the bandwidth requirements for *in vivo* Raman imaging of biological samples by using CS [66, 83], or combining CS with some a priori knowledge of the sample [88]. Usually, the pure CS approach entails long post-processing times, which limits its application in real-time situations. Also, there are scenarios where there is no a priori information about the sample. As a way to achieve real-time acquisition, with no a priori information, novel algorithmic approaches are starting to be used [91]. In a work we have been developing over the last year [92], we introduce the methodology of MC into a single-pixel Raman microscope. Given the technical design of this kind of microscope, the sampling procedure it performs can be directly related to the matrix formalism shown in Sec. 2.2. For each spatial position (a fixed position of the translation stage), the Raman signal is measured by the DMD and the SPD, obtaining a vector which entries correspond to the energies for each wavelength. This process is repeated for all the spatial positions, forming a matrix where each row corresponds to a spatial position, and each column to a

spectral channel. As was described in Sec. 2.2, when a high number of rows of the matrix are linear combinations of a few independent rows, the problem of recovering its missing entries can be solved using MC techniques. Given that the number of different chemical species in a Raman experiment tends to be quite low (dozens, at the most), this computational approach seems very promising as a way to reduce both acquisition and post-processing times.




FIGURE 3.10: **Computational Raman microscopy.** **a)** Top: group of polystyrene beads (green) suspended in water (red). Bottom: Raman spectra of the two present species. In this case, the images were acquired with a measurement ratio of 44%. **b)** Top: Raman microscopy of brain tissue. In this case, both lipids and proteins can be distinguished from the Raman spectra (bottom) analysis. In the inset we highlight the differences between the lipid-rich and protein-rich regions of the sample, that was sampled using a measurement ratio of 56%. In both cases, the eigenimages (only shown for the brain sample), output of the MC algorithm, can be merged to generate the pseudocolored final result.

Using the setup described in Fig. 3.8, we acquired an image of a group of beads made of polystyrene. In Fig. 3.10.a, we can clearly distinguish between the polymer beads (coloured green) in the water suspension (red). From the hyperspectral reconstruction, it is also possible to distinguish the characteristic peaks of the polystyrene vibrational spectrum, together with the water one. In this case, the image size was 199×135 pixels, with 127 spectral channels. The effective pixel dwell time for the acquisition was 8 ms, which is almost one order of magnitude shorter than the traditional approaches [86]. The images shown here were obtained by using only 44% of the total number of entries of the hyperspectral matrix, and the recovery process was done at a rate of 11 ms/pixel (spatial) using a laptop computer and a MC algorithm based on singular value thresholding [64]. To test the procedure in a more challenging scenario, we also obtained images of brain slices, as can be seen in Fig. 3.10.b. In this case, we can distinguish between tubular structures, which correspond to myelins made of lipids (coloured green), and non-tubular structures, which correspond to the axons, rich in proteins (coloured red). For a 215×170 pixels and 127 spectral channels dataset, we only measured a 58% of the total number of entries of the hyperspectral matrix, thus obtaining an effective pixel dwell time of 19.6 ms. The recovery process, using a MC procedure based on non-negative matrix factorization [93] was done at a rate of 8 ms/pixel (spatial) using a laptop computer [92].

These results clearly show the capability of MC to obtain multispectral information in a Raman microspectroscopy setup. This is relevant for several reasons. First, the experimental design of computational Raman microspectroscopes can be naturally done with the working principles of MC in mind, which makes its implementation straightforward. Second, using this kind of technique allows us to obtain images with reconstruction rates equal or even lower to the acquisition time, which is crucial when working in biological scenarios. Last, this presents an unsupervised (i.e., without needing a priori information) approach to Raman microspectroscopy. However, it can also be used as a fast way of obtaining information about a sample, and then using this information to speed-up acquisition processes of future acquisitions, similar to the adaptive procedures that were introduced in SPI in the spatial domain in Chapter 2.

3.3 Phase imaging using a single-pixel camera

In the previous sections we have seen several examples of how to obtain both polarization and spectral information, but those are just two examples of interesting dimensions to measure. In such diverse fields as biomedical imaging, visual optics, astronomy, and three-dimensional imaging, to name just a few, phase is probably the most relevant dimension to add to a conventional imaging system. With the ability to obtain phase information you can see transparent samples, measure thickness and depth, and correct aberrated images.

Obtaining phase information efficiently with a SPC is not trivial. After all, both the detector and the SLM we use in our systems are not sensitive to phase. In the case of SPC based on phase SLMs, although this process is quite straightforward, the systems developed tend to be very slow [47, 94, 95]. However, there are several workarounds that allow us to obtain phase information using a DMD and conventional detectors, such as photodiodes [49, 94, 96]. In a work developed during 2017, we used a lateral position detector and a DMD to build a phase imaging system based on the working principles of Shack-Hartmann wavefront sensors (SHWS) [52]. Let's see how it works.

In a SHWS, an array of lenslets is placed between the wavefront that one wants to measure and a pixelated detector. As the distance between the detector and the lenslet array equals the focal length of the lenslets, a set of foci is generated onto the detector plane. The pixelated sensor is divided in equally sized regions so that each one tracks the position of a focal spot. If the wavefront is plane, you get an ordered array of foci, as can be seen on Fig. 3.11a.

On the contrary, if the wavefront arriving at the SHWS presents some aberrations, the focal spots generated by the lenslet array will be displaced (Fig. 3.11b). It can be easily demonstrated that the displacement of each focal spot is related to the derivative of the phase at each aperture (defined by each lenslet) by the expression:

$$\vec{\Delta} = (\Delta x, \Delta y) = \frac{\lambda f}{2\pi} \vec{\nabla} \varphi, \quad (3.4)$$

where λ is the wavelength of the wavefront, f is the focal length of each lenslet, and φ is the phase of the wavefront at each aperture (supposing that the phase is constant over the whole aperture) [52]. After measuring the displacements with the aid of the pixelated sensor, it is very easy to recover the phase distribution by direct integration. However, this method presents several limitations. First, manufacturing processes limit the lenslets size, which tend to vary around a hundred microns.




FIGURE 3.11: **Shack Hartmann wavefront sensing principles.** **a)** For a plane wavefront, the lenslet array generates a grid of equally spaced focal spots on the sensor. **b)** When an aberrated wavefront arrives at the sensor, the focal spots move around, and the new positions are related to the local phase of the wavefront at each lenslet. Adequate tracking of this position, along with numerical integration, provides the phase distribution arriving at the sensor.

This limits the spatial resolution you can achieve and, given the finite size of the wavefront cross-section, the number of spatial positions you can measure. Moreover, when reducing lenslet size, the amount of light arriving to each region of the sensor decreases, thus reducing signal-to-noise ratio. Second, if the wavefront aberration is big enough, the focal spot generated by one lenslet can move a distance bigger than the region of the sensor delimited to track its displacement, going inside a neighbouring region. This causes a section of the sensor to contain multiple focal spots, generating a cross-talk scenario where it is not easy at all to assign a given displacement to a lenslet (see Fig. 3.12.a) [97]. A manner of tackling this problem is by changing the focal spot of the lenslet array, but there is a tradeoff between sensitivity in the phase measurement and dynamic range that cannot be physically bypassed (although there exist diverse computational methods that try to solve it [97–100]).

In order to solve these problems, one simple approach could be to just block the light from all lenslets except one (see Fig. 3.12.b). In this case, only one focal spot is generated onto the sensor surface, and the displacement can be as big as the full sensor size allows. Then, one would just need to measure sequentially with only one lenslet at a time to recover the full displacement information. In this scenario, it would be possible to just use a lens and an aperture instead of an array of lenslets (Fig. 3.12.c), which would simplify manufacturing and would improve spatial resolution (as apertures can be easily fabricated with arbitrary small sizes). However, the tradeoff between spatial resolution and sensitivity would still be present, as reducing aperture size also reduces the amount of photons arriving to the sensor.

As a way to tackle this tradeoff, it is possible to use coded apertures generated with a DMD. Now, instead of just using a small aperture and performing a scan over the wavefront cross-section, a coded aperture is used to sample the full wavefront in each measurement (Fig. 3.12.d). This fast way of measuring uses the full sensor size and also greatly increases the signal level. After a set of coded apertures are generated, the displacement map can be easily acquired, and again, direct integration provides the phase of the wavefront (see Fig. 3.13 for a comparison of the three methods).




FIGURE 3.12: **Shack Hartmann tradeoffs and possible workarounds.** **a)** When a strongly aberrated wavefront arrives to the sensor, the focal spot generated by one lenslet can lie into the region of the detector associated to a different lenslet, generating a cross-talk problem. **b)** If all the lenslets are blocked, the cross talk problem disappears, and the maximum displacement measurable is greatly increased. **c)** Instead of using a blocked array of lenslets, the same result can be obtained by using a single lens and an aperture. **d)** In order to increase light efficiency, it is also possible to implement a coded aperture. Now, the sensor is used to measure the resulting centroid of the superposition of the wavefront and the known aperture.

Our experimental proposal can be seen in Fig. 3.14. Light coming from a laser source is collimated and directed to the DMD, where the coded apertures (Hadamard patterns in this case) are generated. Using a 4-f system, an image of the DMD plane is projected onto the object under study. Light diffracted by the object goes through a thin lens, and the energy and centroid position of the generated light distribution is measured with a lateral position detector. It has to be noted that, even though the traditional way of measuring these quantities would be to use a pixelated detector, that would yield extreme low frame rates in this sequential system. Using a lateral position detector, which can detect small displacements at speeds of tens of kHz, makes the system able to work at frame rates comparable with the ones of a conventional SHWS, but with increased dynamic range (a factor of 2 when compared to a commercial SHWS) and spatial resolution [52].

A sample of the results that can be obtained with this system is shown in Fig. 3.15. In this case, we imaged a plate simulating a typical coma aberration of a human eye (Fig. 3.15.a) with a resolution of 64×64 pixels. Using CS techniques, we were able to recover a very good estimation of the aberration with a 90% compression ratio. Using a fast lateral position detector and the DMD at its full frame rate (~ 20 kHz) would allow us to perform real-time aberrometry with acquisition times of 20 ms. We also show the spherical phase introduced to a wavefront by a spherical lens (Fig. 3.15.b). Last, as a way of testing the enhanced spatial resolution of the system, we imaged a small sample consisting of a photoresin layer with some bubbles (Fig. 3.15.c). This time the spatial resolution of the images was 128×128 pixels, with a pixel pitch of $41.04 \mu\text{m}$. Both the number of pixels and the pixel pitch are clearly superior to the ones usually found in conventional SHWS, where the number of lenslets lies around a thousand, and the minimum lenslet size tends to be around one hundred microns.




FIGURE 3.13: **Comparison of different wavefront sensing strategies.** **a)** Shack-Hartmann wavefront sensing. The aberrated wavefront is focused by the lenslet array onto the surface of a pixelated detector. The position of each focal spot is related to the phase of the wavefront at each aperture. By measuring the whole displacement distribution, the phase can be recovered by direct integration. **b)** Raster scanning wavefront sensing. Instead of using a lenslet array, now only a small part of the wavefront is measured each time. A small aperture is moved along the transverse direction of the wavefront, and the resulting centroid position is measured. When all the positions have been measured, an analogous map of centroid distributions can be generated, and the phase is obtained by numerical integration. **c)** Coded aperture wavefront sensing. Instead of using a small aperture, the same information can be obtained by using coded apertures that occupy the whole wavefront cross-section, increasing signal-to-noise ratio in the process. Image extracted from [52].




FIGURE 3.14: **Experimental setup of a phase imaging SPC.** Captions: laser source (LS), thin lenses (L_1 , L_2 , and L_3), object (OBJ), cosensing lens (CL), lateral position detector (PD). Image adapted from [52].




FIGURE 3.15: **Phase imaging using a SPC.** **a)** 3D representation of the phase distribution of a coma aberration. **b)** A spherical lens (left), and the 3D representation of the phase distribution it generates. **c)** Phase image of a photoresist layer placed on a transparent plate. On the left we show a picture of the photoresist layer. The black square marks the region to be imaged. On the right we can see the phase image and its 3D representation. Images adapted from [52].

Chapter 4

So long, and thanks for all the fish¹

Along this thesis, we have seen how SPI systems are perfect candidates to obtain multidimensional information using computational imaging techniques. This is due to two main reasons. First, the simple detectors used in a SPC make it possible to work in almost all the electromagnetic spectrum in an inexpensive and straightforward way. While conventional cameras are quite superior in the VIS spectral range, they fail to provide efficient solutions outside of it. Second, a SLM lies at the heart of every SPC, its purpose being to spatially sample one or more of the dimensions that one wants to acquire. The sequential way of measuring information of SPI systems is usually a drawback, as lengthy acquisitions are the norm. However, the presence of the SLM brings an asset to any optical system that cannot be ignored: once you introduce an SLM to your system, it is possible to mold light in almost any imaginable way. This control over light makes it possible, for example, to implement optical aberration corrections, to deal with scattering in biological scenarios, or as we saw along this text, to introduce computational techniques to improve the quality of the obtained images, reduce post-processing times, and increase the acquisition speed of your system.

The main contributions of this thesis can be divided into three main branches. First, I have worked on the development of several procedures on the computational (or software) field of SPI systems. After the introduction of the main algorithmic approaches of SPCs (matrix inversion and CS), I have shown an adaptive approach to SPI, based on wavelet transforms and structured illumination that only uses high resolution sampling in the regions of the scene where there is relevant information. This sampling approach presents two main benefits: a reduction in the total number of required measurements, and very low recovery times. I have also tackled the implementation of novel signal processing techniques, such as MC, into SPI systems. In this direction, I have introduced a first approach in the field of Raman microspectroscopy. In this application, the use of a photon counting detector makes it possible to work in extreme low light level conditions, with high SNR, and with low measurement times. This allows to obtain images in a fast way (with pixel dwell times almost one order of magnitude lower than the conventional approaches based on pixelated detectors), and with even shorter reconstruction times, paving the way to obtain structural chemical analysis of biological samples *in vivo*.

In a second group of applications, I have progressed on the hardware development of SPI systems. In the second chapter I described a SPC using the sensing principles of balanced detection. This novel configuration benefits from increased

¹ “[...] on the planet Earth, man had always assumed that he was more intelligent than dolphins because he had achieved so much –the wheel, New York, wars and so on– whilst all the dolphins had ever done was muck about in the water having a good time. But conversely, the dolphins had always believed that they were far more intelligent than man—for precisely the same reasons. [...]” – Douglas Adams, *The Hitchhiker’s Guide to the Galaxy*.

acquisition speed (by a factor of two when compared to a conventional SPI system) and higher SNR due to the balanced detection configuration, which also solves a problem in light efficiency of SPC systems based on DMDs (as half the light is typically lost after reflection on the DMD surface). This advantages, in combination with CS techniques, allows to obtain a video feed using a SPC.

Finally, I have described multiple contributions in the field of multidimensional imaging. Using a detection system formed by a fibre spectrometer and several polarizing elements, I showed the capability of SPI systems to obtain combined polarimetric and multispectral information in the VIS and NIR spectral range (450–750 nm). Also, I presented a robust, fast, and flexible procedure to obtain phase information using a lateral position detector and a DMD, merging the fundamentals of Shack-Hartmann wavefront sensing with SPI ideas. Using this system, it is possible to double the phase dynamic range obtained with a traditional SHWS. Also, given the spatial sampling procedure based on a DMD, the spatial resolution of the system is not limited any more by manufacturing processes of small lenslets. In a proof-of-principle system, the spatial resolution was one order of magnitude higher than the one of a commercial SHWS, which made possible to obtain phase images of thick samples with spatial resolutions under 100 microns, that could be further improved by using microscopy systems.

Listing the technological applications of the systems developed here would be impossible due to space limitations. Just to name a few, multispectral systems are widely used in industry, as a control tool in manufacturing processes, to provide extra information in metrology systems, weather forecast, plague control and soil identification. Polarization information is of paramount interest when studying surface characteristics, and thus is very used while searching defects in optical elements, windscreens, or when detecting water surfaces in aerospatial imagery. In the case of phase imaging, probably the most relevant applications are in the field of biomedical imaging, as many of the structures that need to be studied only produce phase changes to the light that goes through them. In this same field, the capability of SPI systems to easily work in the IR spectral range opens up a whole realm of possibilities, given the existence of a NIR penetration window in biological tissues. During the last decades, researchers have started to use this spectral window to see inside living tissues with non-ionizing radiation, learn about them, and develop novel medical treatments. However, current optical systems still suffer from limited spatial and temporal resolution when trying to see objects at big depths inside living specimens. The techniques explained here, both in the hardware and software side, make a promising starting point to tackle this challenge, given their benefits in price, their flexibility in optical design, and the additional information they can provide (phase, wavelength, polarization or even non-invasive chemical labelling).

Last, another interesting aspect is the strong computational charge present in every SPI system. Data inversion, compressive sensing, matrix completion, and multiple flavours of artificial intelligence (AI) have been or are starting to be explored in SPI systems as a way to increase their performance. In the following decades, it is foreseen that the use of computational systems based on AI will be extended in almost every aspect we can imagine of our lives: automated driving cars, face recognition systems, optical encryption, entertainment, medical diagnosis, etc. To reach that level of technological development, scientists will need to learn these ideas, and SPI systems form a simple, confined, and controllable experimental playground to explore them from an academical point of view.

Appendix A

Taming light: Digital Micromirror Devices

When you go to the movies, or watch a presentation at your university, there is a very high chance that you are seeing the work of a DMD. The first DMD was developed by Larry J. Hornbeck (and others) in 1987, while working at Texas Instruments. Since then, DMDs have been extensively used in metrology, head-mounted displays, and digital cinema projection, to name just a few applications. They have a very high spatial resolution, having millions of very small pixels. This, coupled with their physical specifications (very fast refresh rates and high light power damage thresholds), make them a very good choice in audiovisual applications, where generating bright, high quality images is key to a good user experience. Also, they have also benefited from the advances in semiconductor manufacturing processes. This has cheapen their production costs, and made them an established solution in the fields I mentioned earlier.

The basic idea behind a DMD is to build a reflective SLM that uses small mirrors (with current sizes in the order of ten microns) to either reflect or block light in one direction. To do so, each mirror is placed on top of a yoke, that allows the mirrors to move (see Fig. A.1.a). Two electrodes are placed under the mirror-yoke structure. These electrodes are connected to a CMOS memory cell, where two possible states can be stored (1 or 0). Both states correspond to the two possible micromirror positions ($\pm\sim 10^\circ$, the total angle depending on the DMD model), and are generated by applying either zero or a determined bias voltage to the structure. Releasing the mirror is done by removing the bias voltage.

With these characteristics in mind, it is easy to see some of the benefits of DMDs over other types of SLMs. Their working principle is based on semiconductor technologies, providing very high refresh rates (~ 22 kHz) when compared, for example, to SLMs based on liquid crystals, which were the standard technology at the time for video projection systems. Also, given that each pixel is made from metal (typically aluminium), it is quite easy to work with very low power losses over the whole visible spectrum (and also the IR range).

However, DMDs still present drawbacks for some applications. First, given that each micromirror only has two possible positions, light can be either reflected or not, providing binary light modulation only (Fig. A.1.b). In multimedia applications this is not very problematic, as human eye limitations can be used to generate gray levels easily. For each frame to be displayed, mirrors can be a different amount of time in the bright/dark states, and the brain will integrate all the information as a single image with different light levels for each pixel. If you want to generate a 30 fps video, each frame will be displayed $1/30$ of a second. At 22 kHz, you can achieve $22000/30 \simeq 733$ gray levels (more than a 9-bit depth). Modern DLP projectors in cinemas use a combination of three DMDs (each one with a different




FIGURE A.1: **Fundamentals of a DMD and DMD-based video projector.** **a)** Exploded view of a DMD pixel. Each micromirror is placed on top of a CMOS memory cell that controls its movement by applying different voltages. **b)** In a video projector, a light source illuminates the DMD screen. Bright pixels reflect light into the screen direction, where the image is formed by using an optical system. Dark pixels direct light to a stop. Figures adapted from US Patent No. 7518781 B2 (Apr. 14, 2009).

light source color) to generate gigantic color gamuts. Moreover, there is a different solution that, instead of trading temporal resolution for gray levels, does the same by giving away spatial resolution. By grouping several micromirrors it is possible to generate a “superpixel” and, by putting a different number of mirrors in the bright state, project pixels with different brightness levels [101–103].

Also, given their nature, DMDs are amplitude-only SLMs. Nevertheless, there are workarounds to control both the amplitude and phase of a light beam using a DMD, which is extremely useful in biological imaging scenarios. To do so, you can generate binary holograms in the DMD screen, which after propagation and correct spatial filtering generate arbitrary complex light distributions [104–106]. This is interesting because it is a really fast method of obtaining phase control over a light beam, but comes at the cost of losing both spatial resolution and power efficiency.

Appendix B

Singular Value Decomposition

The Singular Value Decomposition (SVD) of an $m \times n$ matrix \mathbf{M} provides a triplet of matrices, \mathbf{U} , $\mathbf{\Sigma}$, and \mathbf{V} that hold:

$$\mathbf{M} = \mathbf{U} \cdot \mathbf{\Sigma} \cdot \mathbf{V}^\dagger, \quad (\text{B.1})$$

where \mathbf{U} is an $m \times m$ unitary matrix; $\mathbf{\Sigma}$ is a $m \times n$ diagonal matrix with only non-negative real numbers in its diagonal (σ_i), sorted in descending order; and \mathbf{V} is an $n \times n$ unitary matrix with conjugate transpose \mathbf{V}^\dagger . The elements on the diagonal of $\mathbf{\Sigma}$ are known as the singular values of \mathbf{M} . Moreover, the number of non-zero singular values is equal to the rank of \mathbf{M} .

There are several intuitive interpretations of the meaning of the SVD of a matrix (as rotating and scaling matrices, geometrical, etc.). In our case, the most practical one for us is to understand the SVD as a decomposition of \mathbf{M} into a weighted sum of separable matrices. It can be demonstrated that we can write

$$\mathbf{M} = \sum_i \sigma_i \cdot \mathbf{U}_i \otimes \mathbf{V}_i^\dagger = \sum_i \mathbf{A}_i, \quad (\text{B.2})$$

where \mathbf{U}_i and \mathbf{V}_i^\dagger are the i -th columns of \mathbf{U} and \mathbf{V}^\dagger , and \otimes is the outer product of two column vectors ($\mathbf{u} \otimes \mathbf{v} = \mathbf{u} \cdot \mathbf{v}^\top$). This is interesting because we can divide any operator (\mathbf{M}) into two separate operators (\mathbf{U} and \mathbf{V}). For example, an image filter can be divided into horizontal and vertical filters. Also, if a multispectral cube is expressed in matrix form, where each row contains the spectrum of one spatial position, the columns of \mathbf{U} and \mathbf{V} contain the spatial features of the scene and the spectral signatures of each element present, respectively.

One interesting use of the SVD decomposition is that it can be easily used to compress images. Take, for example, the sample picture from the Eiffel Tower in Fig. B.1.a. If we calculate its SVD, we can clearly see that almost all its singular values are negligible. To compress the image, we can eliminate the lowest singular values of the decomposition, and then recover the image using

$$\mathbf{M}' = \mathbf{U} \cdot \mathbf{\Sigma}' \cdot \mathbf{V}^\dagger, \quad (\text{B.3})$$

where $\mathbf{\Sigma}'$ is the diagonal matrix with only a subset of the biggest singular values of \mathbf{M} . In the example, I just used the first 200 singular values of the original image, and set the rest of the elements to zero (Fig. B.1.b). It can be clearly seen that with this compression level ($\sim 80\%$), we still get a very good quality image.




FIGURE B.1: **Image compression by SVD.** **a)** Original image and its singular values. From a rough inspection, we easily identify that most of the information of the picture is stored in a reduced number of singular values. **b)** Compressed image recovered by using only the first 200 singular values.

Appendix C

List of publications

The following publications are a part of this thesis. They have been cited throughout the text, and contain specific information about the experimental design of the presented systems, along with additional results and studies. An interested reader might find useful information in them, so they are attached in the following pages.

- Soldevila, F. et al. Single-pixel polarimetric imaging spectrometer by compressive sensing. *Appl. Phys. B* **113**, 551–558 (2013).
- Soldevila, F. et al. Spectropolarimetric Imaging Techniques with Compressive Sensing, in *Multi-Dimensional Imaging* 371–390 (John Wiley & Sons, Ltd, 2014). doi:10.1002/9781118705766.ch16
- Soldevila, F., Salvador-Balaguer, E., Clemente, P., Tajahuerce, E. & Lancis, J. High-resolution adaptive imaging with a single photodiode. *Sci. Rep.* **5**, 14300 (2015).
- Soldevila, F. et al. Computational imaging with a balanced detector. *Sci. Rep.* **6**, 29181 (2016).
- Soldevila, F., Durán, V., Clemente, P., Lancis, J. & Tajahuerce, E. Phase imaging by spatial wavefront sampling. *Optica* **5**, 164 (2018).
- Soldevila, F., Dong, J., Tajahuerce, E., Gigan, S. & de Aguiar, H. B. Fast compressive chemical bio-imaging via matrix completion. to be published. (2018).

Also, during the time I have been working on this thesis, I developed some contributions in parallel which are very related to the SPI systems presented here, but with applications in imaging through scattering media and phase imaging. They can provide a wider view of the range of applications of SPI systems, so I also list them here.

- Tajahuerce, E. et al. Image transmission through dynamic scattering media by single-pixel photodetection. *Opt. Express* **22**, 16945–16955 (2014).
- Durán, V. et al. Compressive imaging in scattering media. *Opt. Express* **23**, 14424–14433 (2015).
- Soldevila, F., Clemente, P., Tajahuerce, E. & Lancis, J. Single-Pixel Imaging Using the Hadamard Transform, in *Encyclopedia of Modern Optics* **4**, 193–198 (Elsevier, 2018).
- González, H. et al. High sampling rate single-pixel digital holography system employing a DMD and phase-encoded patterns. *Opt. Express* **26**, 20342 (2018).

Single-pixel polarimetric imaging spectrometer by compressive sensing

F. Soldevila · E. Irlles · V. Durán · P. Clemente · Mercedes Fernández-Alonso · Enrique Tajahuerce · Jesús Lancis

Received: 1 February 2013 / Accepted: 6 May 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract We present an optical system that performs polarimetric spectral imaging with a detector with no spatial resolution. This fact is possible by applying the theory of compressive sampling to the data acquired by a sensor composed of an analyzer followed by a commercial fiber spectrometer. The key element in the measurement process is a digital micromirror device, which sequentially generates a set of intensity light patterns to sample the object image. For different configurations of the analyzer, we obtain polarimetric images that provide information about the spatial distribution of light polarization at several spectral channels. Experimental results for colorful objects are presented in a spectral range that covers the visible spectrum and a part of the NIR range. The performance of the proposed technique is discussed in detail, and further improvements are suggested.

1 Introduction

Multispectral imaging (MI) is a useful optical technique that provides two-dimensional images of an object for a set of specific wavelengths within a selected spectral range.

F. Soldevila · E. Irlles · V. Durán (✉) · P. Clemente · M. Fernández-Alonso · E. Tajahuerce · J. Lancis
GROC-UJI, Departament de Física, Universitat Jaume I,
12071 Castelló, Spain
e-mail: vduran@sg.uji.es

V. Durán · M. Fernández-Alonso · E. Tajahuerce · J. Lancis
Institut de Noves Tecnologies de la Imatge (INIT),
Universitat Jaume I, 12071 Castelló, Spain

P. Clemente
Servei Central d'Instrumentació Científica, Universitat Jaume I,
12071 Castelló, Spain

Dispersive elements (as prisms or gratings), filter wheels or tunable band-pass filters are typical components used in MI systems to acquire image spectral content [1]. In certain applications, MI can be improved by adding spatially resolved information about the light polarization. Multispectral polarimetric imaging facilitates the analysis and identification of soils [2], plants [3] and surfaces contaminated with chemical agents [4]. In the field of biomedical optics, multispectral polarimetric imaging has been applied to the characterization of human colon cancer [5] or the pathological analysis of skin [6]. In many cases, polarimetric analysis can be performed by just including a linear polarizer in an imaging system to record images for various selected orientations of its transmission axis [6, 7]. A simple configuration that includes two orthogonal polarizers integrated in a spectral system has been used for noninvasively imaging of the microcirculation through mucus membranes and on the surface of solid organs [7]. An illustrative example of a spectral camera with polarimetric capability is a system that combines an acousto-optic tunable filter with a liquid-crystal-based polarization analyzer [8].

In an apparently different context, compressive sampling (CS) has emerged in recent years as a novel sensing theory that goes beyond the Shannon–Nyquist limit [9]. In the field of imaging, CS states that an N -pixel image of an object can be reconstructed from $M < N$ linear measurements. This sub-Nyquist condition is achieved by exploiting the “sparsity” of natural images. According to this property, when images are expressed in a proper function basis, most terms are negligible or zero-valued. CS theory ensures that it is possible to reconstruct the object images from a relatively small collection of well-chosen measurements, typically by an iterative acquisition process. The object reconstruction is obtained from experimental data by solving a convex optimization program.

One of the most outstanding applications of CS is the design of a single-pixel camera [10, 11], which offers promising benefits at spectral regions where image sensors are impractical or inexistent [12]. In contrast to conventional image sensors, which typically perform intensity measurements, single-pixel detectors can provide information about other properties of an incoming light field, as its spectrum or its polarization. The substitution of the photodiode of a CS single-pixel camera by a spectrometer without spatial resolution permits to perform hyperspectral imaging [13, 14]. In the same way, single-pixel polarimetric imaging has been demonstrated with a CS camera that includes a commercial beam polarimeter [15]. CS has also been applied to biological fluorescence microscopy [16]. In this case, the CS fluorescence microscope includes a photomultiplier tube as a point detector, since biological samples often have low fluorescence.

In this work, we present a CS imaging system able to provide spatially resolved information about the spectrum and the polarization of the light reflected by an object. As a detector, we use a polarization analyzer followed by a fiber spectrometer with no spatial resolution. The key element of our system is a digital micromirror device (DMD), which makes possible the CS acquisition process. To this end, a set of binary intensity patterns is sequentially generated by the DMD to sample the image of an object of interest. The experimental data are subsequently processed to obtain a set of multispectral data cubes, one for each selected configuration of the analyzer. For a given spectral channel, the corresponding polarimetric images can be linearly combined to derive the spatial distribution of the Stokes parameters of light [8, 15]. In this sense, the single pixel described in this paper is a first step toward the realization of an imaging Stokes polarimeter like in Ref. [15], but with the ability of performing polarization analysis for a large variety of wavelengths.

2 Outline of compressive sampling

The basis of single-pixel imaging by CS can be briefly presented as follows. Let us consider a sample object, whose N -pixel image is arranged in an $N \times 1$ column x . This image is supposed to be compressible when it is expressed in terms of a basis of functions, $\Psi = \{\Psi_l\}$ ($l = 1, \dots, N$). From a mathematical point of view, x can be written as $\mathbf{x} = \Psi\mathbf{s}$, where Ψ is a $N \times N$ matrix that has the vectors $\{\Psi_l\}$ as columns and \mathbf{s} is the $N \times 1$ vector composed of the expansion coefficients. The assumed sparsity of the image implies that only a small group of these coefficients is nonzero. In order to determine \mathbf{x} , we implement an experimental system able to measure the projections of the object image on a basis of M intensity

patterns φ_m ($m = 1, \dots, M$) of N -pixel resolution. This acquisition process can be written as

$$\mathbf{y} = \Phi \mathbf{x} = \Phi (\Psi \mathbf{s}) = \Theta \mathbf{s}, \quad (1)$$

where \mathbf{y} is the $M \times 1$ column formed by the measured projections, and Φ is the $M \times N$ sensing matrix. Each row of Φ is an intensity pattern φ_m , and the product of Φ and Ψ gives the $M \times N$ matrix Θ acting on \mathbf{s} . The underlying mathematical formalism of CS states that there is a high probability of reconstructing \mathbf{x} from a random subset of coefficients ($M < N$) in the Ψ domain. Equation (1) constitutes an underdetermined matrix relation, so it must be resolved by means of a proper reconstruction algorithm. The best strategy to perform this step is based on the minimization of the l_1 -norm of \mathbf{s} subjected to the restriction given by Eq. (1). As the measurements $\{y_m\}$ are affected by noise, the CS recovery process is usually reformulated with inequality constrains [9, 10]. In this case, the proposed reconstruction \mathbf{x}^* is given by $\mathbf{x}^* = \Psi \mathbf{s}^*$, where \mathbf{s}^* is the solution of the optimization program

$$\min \|\mathbf{s}'\|_{l_1} \quad \text{such that} \quad \|\mathbf{y} - \Theta \mathbf{s}'\|_{l_2} < \varepsilon, \quad (2)$$


where ε is an upper bound of the noise magnitude and the l_2 -norm is used to express the measurement restriction.

3 Description of the polarimetric imaging spectrometer

3.1 Optical system

The scheme of our polarimetric spectral camera is shown in Fig. 1a. A white light source illuminates a sample object, and a CCD camera lens images the object on a DMD, which is a reflective spatial light modulator that selectively redirects parts of an input light beam [17]. A DMD consists of an array of electronically controlled micromirrors that can rotate about a hinge, as is schematically depicted in Fig. 1b. Every micromirror is positioned over a CMOS memory cell. The angular position of a specific micromirror admits two possible states ($+12^\circ$ and -12° respect to a common direction), depending on the binary state (logic 0 or 1) of the corresponding CMOS memory cell contents. In this way, the light can be reflected at two angles depending on the signal applied to the mirror. The DMD of our system is a Texas Instrument device (DLP Discovery 4100) with a resolution of $1,920 \times 1,080$ micromirrors and a panel size of $0.95''$. The mirror pitch is $10.8 \mu\text{m}$, and the fill factor is greater than 0.91. The optical system 1 has its optical axis forming an angle respect to the orthogonal direction to the DMD panel that approximately corresponds to twice the tilt angle of the device mirrors (24°). As is shown in Fig. 1c, in such a configuration, a micromirror oriented at $+12^\circ$ orthogonally reflects the light

Fig. 1 **a** Setup for single-pixel polarimetric multispectral imaging, **b** transverse view of an individual DMD micromirror showing its two possible orientations and **c** scheme of the DMD operation mode



into the next part of the system, appearing as a bright pixel (ON state). In their turn, micromirrors oriented at -12° result to be dark pixels (OFF state). The light emerging from the bright pixels of the DMD is collected by a second lens system similar to the first one (optical system 2 in Fig. 1a). This lens system couples the light into a silica multimode fiber with a diameter of $1,000\ \mu\text{m}$ and a spectral range of $220\text{--}1,100\ \text{nm}$, which is connected to a commercial concave grating spectrometer (Black Comet CXR-SR from StellarNet). The wavelength resolution of this spectrometer is $8\ \text{nm}$ (with a slit of $200\ \mu\text{m}$), and the maximum signal-to-noise ratio (SNR) is $1,000:1$. Just before the fiber, the light passes through a configurable polarization analyzer. In our camera, this analyzer consists of a film polarizer mounted in a rotating holder.

3.2 Operation principle

The CS single-pixel camera shown in Fig. 1a performs the iterative acquisition process synthesized in Eq. (1). The DMD sequentially produces the set of M irradiance patterns of N -pixel resolution that composes the sensing matrix Φ . The collected data consist of a succession of spectra, one for each pattern sent to the DMD. The M irradiances measured by the spectrometer for each

spectrum channel form the vector \mathbf{y} of Eq. (1). As spectrum channels have a prefixed bandwidth, the quantities $\{y_i\}$ that feed the CS algorithm are actually an average of the measured irradiances within the considered range.

For the practical implementation of the CS acquisition process, it is essential to determine which incoherent basis should be selected (when no prior information on the object is accessible). A suitable choice for image basis results to be the Hadamard–Walsh functions, which constitute a basis known to be incoherent with the Dirac basis [10]. Hadamard matrices are square arrays of plus and minus ones, whose rows (and columns) are orthogonal relative to one another. Each row of a Hadamard matrix can be interpreted as a rectangular wave ranging from ± 1 (Walsh function). In this context, the Hadamard matrix performs the decomposition of a function by a set of rectangular waveforms, instead of the usual harmonic waveforms associated with the Fourier transform [18]. From an experimental point of view, a CS acquisition process using the Hadamard–Walsh basis can be carried out by generating a collection of binary intensity patterns, easily imprinted on a DMD. To represent a Hadamard function H on the DMD, we use two binary patterns $H^+ = (J + H)/2$ and $H^- = (J - H)/2$ that are related by $H = H^+ - H^-$. Here, J is the matrix consisting of all 1s. Generating H^+

and H^- sequentially and subtracting the measured intensities, we obtain the one that corresponds to H . In principle, when the total light intensity is known, the acquisition process can be performed by using only one of the above sequences of binary patterns, taking a unique measurement for every Hadamard function. This would lead to a reduction in the acquisition time by a factor two. However, this procedure is sensitive to the existence of source intensity fluctuations. In such a common situation, the method based on two successive measurements increases the signal-to-noise ratio (SNR) of the acquired data.

The sequential measurement process requires the synchronization between the DMD and the fiber spectrometer with the aid of a computer (not shown in Fig. 1). For each intensity pattern generated by the DMD, the spectrum of the light coming from the object is measured by the spectrometer, configured with an integration time that ensures an acceptable SNR for the acquired data. The minimum integration time provided by our spectrometer (1 ms) represents the main limiting factor on the measurement rate of our device, since DMDs are modulators that can work at relatively high frequencies (upper than 1 kHz). For a given acquisition frequency, the total time required to take image data increases with the number of measurements, which is, in accordance with the CS theory, a pre-established fraction of the image resolution. Therefore, a single-pixel camera shows a trade-off between image resolution and image acquisition time.

In our experimental setup, the measurement process is executed and controlled by means of custom software written with LabVIEW. The programming code used in the off-line CS reconstruction is called *l1eq-pd*, which solves the standard basis pursuit problem using a primal-dual algorithm [19]. This code includes a collection of MATLAB routines and is a well-tested algorithm for CS problems. However, other selections are possible and, in fact, the search of improved CS algorithms (more robust to data noise, with lower computation time, etc.) is currently an active area in the field of convex optimization (e.g., see Ref. [20]).

4 Experimental results

4.1 Numerical analysis

The aim of CS is to provide an accurate reconstruction of an object from an undersampled signal, but the exact number of measurements that allows attaining it is not a priori known. In addition, this number strongly depends on the features of the object under consideration. For this reason, when CS single-pixel imaging is attempted, it is useful to begin with relatively low-resolution reconstructions to

estimate the parameters of the acquisition process. In accordance with this approach, and also to evaluate the image quality achievable with our camera, we performed multispectral imaging sending to the modulator Hadamard–Walsh patterns of 64×64 unit cells ($N = 4,096$), each one composed of 8×8 DMD pixels. The covered square window on the modulator panel had a width of ~ 5.5 mm. As a sample object, we used two square capacitors with a width of 7 mm. The illumination source was a Xenon white light lamp, and the polarization analyzer was removed. The number of measurements was $M = 4,096$, which allowed us to fulfill the Nyquist criterion. Eight central wavelengths λ_0 were selected in the visible spectrum. The bandwidth of the corresponding spectral channels was 20 nm ($\lambda_0 \pm 10$ nm). In order to determine the object spectral reflectance, a spectrum was taken from a white reference (Spectralon diffuse 99 % reflectance target from Labsphere, Inc.) to normalize the measured spectra during the CS acquisition process. The integration time of the spectrometer was set at 300 ms.

For each spectral channel, we resolved the off-line CS algorithm with the complete set of measurements ($M = N$). After a suitable filtering, the recovered matrix served as a reference (lossless) image $I_{\text{ref}}(i, j)$, where (i, j) indicates the location of an arbitrary image pixel. The reconstruction process was then repeated using decreasing fractions of the total number of pixels. Concretely, the value of M was varied from 5 to 90 % of N . The fidelity of the reconstructed images was estimated by calculating the mean square error (MSE), given by

$$\text{MSE} = \frac{1}{N} \sum_i \sum_j [I(i, j) - I_{\text{ref}}(i, j)]^2, \quad (3)$$

where $I(i, j)$ is the noisy image obtained for a given value of M . We used another metric to evaluate the quality of the reconstruction, the so-called peak signal-to-noise ratio (PSNR), which is defined as the ratio between the maximum possible power of a signal and the power of the noise that affects the fidelity of its representation. In mathematical terms, [21]

$$\text{PSNR} = 10 \log \left(\frac{I_{\text{max}}^2}{\text{MSE}} \right) = 20 \log (I_{\text{max}}) - 10 \log (\text{MSE}). \quad (4)$$

Here, I_{max} is the maximum possible pixel value of the reference image. For each spectral channel, the reference images were represented by 2^8 gray-levels, so $I_{\text{max}} = 255$. Figure 2a, b shows the curves for the MSE and the PSNR versus M for the different values of λ_0 . As is expected, both figures point out that the image quality improves as the number of measurements grows and approximates to the Nyquist limit. However, when $M \geq 0.4 N$, the slope of




Fig. 2 Image quality of the CS reconstructions. **a** MSE and **b** PSNR of recovered images versus the number of measurements. Each curve corresponds to a spectral channel. A reference gray-scale image ($\lambda_0 = 520$ nm) is also included in the MSE graph




Fig. 3 PSNR versus λ_0 for object reconstructions with $M = 0.2 N$. Each line corresponds to a different spectrometer integration time

both curves becomes visibly smoother for all spectral channels. In the case, for instance, of $\lambda_0 = 610$ nm, $MSE = 0.13 (I_{\max})^2$ and $PSNR = 28.72$ dB for $M = 0.4 N$ (the PSNR for $0.9 N$ is only somewhat greater, 29.10 dB).

If images of higher resolution are considered, the priority is to minimize the value of M , since even small fractions of N can imply a huge volume of data feeding the CS algorithm, which, in addition, operates with higher dimensional matrices. The consequence is a dramatic increase in the total computation time. In our system, it is

possible to change the SNR of the reconstructed image (for a given value of M) by controlling the spectrometer integration time, since the noise of the data used in the CS algorithm considerably depends on this parameter. To illustrate this point, the whole CS acquisition process was repeated for different values of the integration time. The number of measurements in all series was $M = 0.2 N$. The resulting curves for the PSNR versus λ_0 are shown in Fig. 3. As can be observed, the PSNR is incremented in approximately 8 dB when the integration time is varied from 50 to 300 ms.

4.2 Multispectral imaging

As a first application of our camera, we performed multispectral imaging of a sample object composed of an unripe cherry tomato together with a red one. The Walsh–Hadamard patterns addressed to the DMD had a resolution of 256×256 unit cells ($N = 65,536$). Each unit cell was composed of 2×2 DMD pixels. With this resolution, in accordance with the discussion of the previous section, the number of measurements was chosen to be $M = 81^2$, which corresponds to $\sim 10\%$ of N (i.e., a compression rate of 10:1). The integration time of each spectrometer measurement was 300 ms.

The object spectral reflectance was determined by means of the white reference used in Sect. 4.1. In the case of plants, this magnitude has been used, for example, to investigate the chlorophyll content in leaves [22, 23]. The noisy data collected by the spectrometer for wavelengths lower than 500 nm imposed an inferior boundary to the useable spectral range. The results of the CS reconstruction for 15 spectral channels are shown in Fig. 4. The selected central wavelengths λ_0 in the visible spectrum (VIS) range from 510 to 680 nm. The bandwidth of each spectrum channel was 10 nm ($\lambda_0 \pm 5$ nm). The recovered images were pseudo-colored, and the color assignment (the wavelength to RGB transform) was carried out with the aid of standard XYZ color-matching functions [24].

In the near-infrared spectrum, the CS algorithm provided an acceptable reconstruction around 860 nm, which is presented by means of a gray-level image. Figure 4 also includes a colorful image of the object obtained from the combination of the conventional three RGB channels.

4.3 Polarimetric multispectral imaging

In this case, the sample object was the same as that used in Sect. 4.1, but the light emerging from each element of the scene had different linear polarizations. This effect was achieved by locating a linear polarizer after the object with its area split in two parts, each of which had its transmission axis oriented at orthogonal directions (0° and 90° ,





Fig. 4 Multispectral image cube reconstructed by means of the CS algorithm. In the VIS spectrum, the reflectance for each spectral channel is a 256×256 pseudo-colored image. A *gray-scale* representation is used for the CS reconstruction in the NIR spectrum range. A colorful image of the scene made up from the conventional RGB channels is also included

respectively). The resolution of the patterns addressed to the DMD was 128×128 unit cells ($N = 16,384$) composed of 4×4 DMD pixels. The number of measurements was $M = 57^2$, which corresponds to $\sim 20\%$ of N (i.e., a compression rate of 5:1). As the source light was in principle unpolarized, half of light was lost after the object polarizer, so the integration time of the spectrometer was increased until 500 ms. The white reference of Sect. 4.1 was employed again to normalize the measured spectra.

Fig. 5 Multispectral image cube for different configurations of the polarization analyzer. The RGB image of the object is also included. Except for the wavelength closer to the NIR spectrum, all channels in the VIS range are represented by pseudo-colored images



Eight central wavelengths λ_0 were selected in the VIS spectrum. The bandwidth of the channels was 20 nm ($\lambda_0 \pm 10$ nm). Aside from the channels at the boundaries of the spectral range, the values of λ_0 correspond to the peak emissions of commercial light-emitting diodes. For each channel, four orientations of the polarization analyzer were sequentially considered in separated measurement series. To simplify data display, image reconstructions are arranged in a table, as can be observed in Fig. 5. Each column corresponds to a spectral channel, and each row shows the results for a given orientation of the analyzer. As shown in Fig. 4, a colorful image of the object is also presented. This RGB image was made up from the data taken for the second configuration of the analyzer (45°). The result for 680 nm is presented by means of a gray-level image due to its proximity to the near-infrared range.

5 Discussion and conclusions

We have performed polarimetric multispectral imaging by using a detector with no spatial resolution, which is composed of a configurable polarization analyzer and a commercial spectrometer. This single-pixel camera employs a DMD that generates a collection of binary intensity patterns that samples the image of the object under study. For a given analyzer configuration, a succession of spectra is sequentially acquired (one for each DMD realization). From this data, the object spectral image cube is recovered off-line by means of a CS algorithm, which makes possible to achieve a sub-Nyquist limit, that is, the total number of measurements is a fraction of the number of image pixels.

In contrast to cameras based on tunable band-pass filters, which carry out a wavelength sweep to measure the spectral content, our system collects the spectral information of all

channels at once (albeit at the expense of sequentially acquiring the spatial information). As a result, the number of channels, their spectral resolution and the total wavelength range are those provided by the spectrometer integrated in the system. This fact facilitates to exploit the high performance of commercial devices. Our spectral system can in principle cover the whole VIS spectrum and part of the NIR range (up to 1.1 microns). In the infrared region, conventional multispectral systems require pixelated sensors specifically designed for that wavelength range (like InGaAs cameras).

As in other multispectral systems, the illumination is a main question to ensure a minimum signal along the selected spectral range. We have used a high-power Xe arc lamp, which produces a continuous and roughly uniform spectrum across the VIS region and a complex line spectrum in the 750–1,000 nm region of the NIR range. However, the decreasing source irradiance at the “blue” side of the VIS spectrum, as well as the low reflectance of samples at that region, limited our spectral range to wavelengths higher than 470 nm.

As is discussed in the previous section, our single-pixel camera presents a trade-off between image resolution and total acquisition time. Increasing the illumination level or reducing the spectral resolution (which permits lower integration times) could make the acquisition time to drop in at least one order of magnitude. A comparable drawback can be found in cameras that employ acousto-optic or liquid crystal tunable filters. In such systems, the higher spectral resolution (number of channels), the longer acquisition time, with a strong dependence on the exposure time of the pixelated sensor used as a detector. For image resolutions similar to those presented in this work, a hyperspectral camera (i.e., with more than 100 spectral channels) can take a few minutes in acquiring a data cube [25].

The single-pixel spectral system presented here also provides spatially resolved information about light polarization. To this end, the camera detector includes a polarizing film mounted in a rotating holder. This element limits the total spectral range, since the optical behavior of polarizing films is wavelength dependent. As a consequence, when the polarimetric multispectral imaging is carried out, the upper boundary of the spectral range is ~ 700 nm. The use of high-grade crystalline polarizers can resolve this limitation. For the successive configurations of the analyzer, a separated series of measurements must be taken. From the polarimetric images recovered for each spectral channel, it is possible to obtain information about the spatial distribution of the Stokes parameters of light, S_i ($i = 0, \dots, 3$). If a linear polarizer is used as analyzer, the spatial distribution of S_1 and S_2 can be straightforward derived. A full Stokes polarization analysis should be performed by means of a rotating circular (or elliptic) polarizer. It is also possible to avoid mobile polarization

elements by using an analyzer that combines voltage-controlled linear retarders (as those based on liquid crystal technology) with linear polarizers [8]. In that case, all polarimetric information could be acquired in a unique series of measurement, changing sequentially the configuration of the variable retarders for each DMD realization.

Acknowledgments This work has been partly funded by the Spanish Ministry of Education (project FIS2010-15746) and the Excellence Net from the Generalitat Valenciana about Medical Imaging (project ISIC/2012/013). Also funding from Generalitat Valenciana through Prometeo Excellence Programme (project PROMETEO/2012/021) is acknowledged.

References

1. R.G. Sellar, G.D. Boreman, Classification of imaging spectrometers for remote sensing applications. *Opt. Eng.* **44**(1), 013602 (2005)
2. K.L. Coulson, Effects of reflection properties of natural surfaces in aerial reconnaissance. *Appl. Opt.* **5**, 905–917 (1966)
3. V.C. Vanderbilt, L. Grant, L.L. Biehl, B.F. Robinson, Specular, diffuse, and polarized light scattered by two wheat canopies. *Appl. Opt.* **24**, 2408–2418 (1985)
4. S.M. Haugland, E. Bahar, A.H. Carrieri, Identification of contaminant coatings over rough surfaces using polarized infrared scattering. *Appl. Opt.* **31**, 3847–3852 (1992)
5. A. Pierangelo, A. Benali, M-R. Antonelli, T. Novikova, P. Validire, B. Gayet, A. De Martino, Ex-vivo characterization of human colon cancer by Mueller polarimetric imaging. *Opt. Express* **19**, 1582–1593 (2011)
6. Y. Zhao, L. Zhang, Q. Pan, Spectropolarimetric imaging for pathological analysis of skin. *Appl. Opt.* **48**, D236–D246 (2009)
7. W. Groner, J.W. Winkelman, A.G. Harris, C. Ince, G.J. Bouma, K. Messmer, R.G. Nadeau, Orthogonal polarization spectral imaging: a new method for study of the microcirculation. *Nat. Med.* **5**, 1209–1213 (1999)
8. N. Gupta, D.R. Suhre, Acousto-optic tunable filter imaging spectrometer with full Stokes polarimetric capability. *Appl. Opt.* **46**, 2632–2637 (2007)
9. E.J. Candès, M.B. Wakin, An introduction to compressive sampling. *IEEE Signal Process. Mag.* **25**, 21–30 (2008)
10. M.F. Duarte, M.A. Davenport, D. Takhar, J.N. Laska, T. Sun, K.F. Kelly, R.G. Baraniuk, Single-pixel imaging via compressive sampling. *IEEE Signal Process. Mag.* **25**, 83–91 (2008)
11. F. Magalhaes, F.M. Araújo, M.V. Correia, M. Abolbashari, F. Farahi, Active illumination single-pixel camera based on compressive sensing. *Appl. Opt.* **50**, 405–414 (2011)
12. W.L. Chan, K. Charan, D. Takhar, K.F. Kelly, R.G. Baraniuk, D.M. Mittleman, A single-pixel terahertz imaging system based on compressed sensing. *Appl. Phys. Lett.* **93**, 121105 (2008)
13. T. Sun, K. Kelly, *Compressive Sensing Hyperspectral Imager, Computational Optical Sensing and Imaging, OSA Technical Digest (CD)* (Optical Society of America, 2009)
14. Y. Wu, I.O. Mirza, G.R. Arce, D.W. Prather, Development of a digital-micromirror-device-based multishot snapshot spectral imaging system. *Opt. Lett.* **36**, 2692–2694 (2011)
15. V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, J. Lancis, Single-pixel polarimetric imaging. *Opt. Lett.* **37**, 824–826 (2012)
16. V. Studer, J. Bobin, M. Chahid, S.H. Shams Mousavi, E. Candès, M. Dahan, Compressive fluorescence microscopy for biological and hyperspectral imaging. *PNAS* **109**, E1679–E1687 (2012)

17. J. Sampsell, An overview of the digital micromirror device (DMD) and its application to projection displays, *SID Int. Symp. Digest of Technical Papers* **24**, 1012 (1993)
18. W.K. Pratt, J. Kane, H.C. Andrews, Hadamard transform image coding. *Proc. IEEE* **57**, 58–67 (1969)
19. <http://users.ece.gatech.edu/~justin/1Imagic>
20. M.A.T. Figueiredo, R.D. Nowak, S.J. Wright, Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems. *IEEE J. Sel. Top. Signal Process.* **1**, 586–597 (2007)
21. W. K. Pratt, *Digital Image Processing*, 4th edn. (Wiley, 2007)
22. J. Vila-Francés, J. Calpe-Maravilla, J. Muñoz-Mari, L. Gómez-Chova, J. Amorós-López, E. Ribes-Gómez, V. Durán-Bosch, Configurable-bandwidth imaging spectrometer based on an acousto-optic tunable filter. *Rev. Sci. Instr.* **77**, 073108 (2006)
23. Z. Xiaobo, S. Jiyong, H. Limin, Z. Jiewen, M. Hanpin, C. Zhenwei, L. Yanxiao, M. Holmes, In vivo non-invasive detection of chlorophyll distribution in cucumber (*Cucumis sativus*) leaves by indices based on hyperspectral imaging. *Anal. Chim. Acta* **706**, 105–112 (2011)
24. <http://www.mathworks.com/matlabcentral/fileexchange/7021-spectral-and-xyz-color-functions>
25. K.J. Zuzak, M.D. Schaeberle, E.N. Lewis, I.W. Levin, Visible reflectance hyperspectral imaging: characterization of a noninvasive, in vivo system for determining tissue perfusion. *Anal. Chem.* **74**, 2021–2028 (2002)

16

Spectropolarimetric Imaging Techniques with Compressive Sensing

Fernando Soldevila¹, Esther Irlés¹, Vicente Durán^{1,2}, Pere Clemente^{1,3},
Mercedes Fernández-Alonso^{1,2}, Enrique Tajahuerce^{1,2} and Jesús Lancis^{1,2}

¹*GROC·UJI, Departament de Física, Universitat Jaume I, Spain*

²*Institut de Noves Tecnologies de la Imatge (INIT), Universitat Jaume I, Spain*

³*Servei Central d'Instrumentació Científica, Universitat Jaume I, Spain*

16.1 Chapter Overview

The information that the human eye can provide is limited. Although we are able to see in a wide range of distances, under different light conditions, and in a relatively broad spectral range, in many applications it is necessary to acquire information far beyond the limits imposed by the human eye. To this end, a great variety of image techniques have been developed [1]. As an archetypical example, microscopy, which is essential in fields like biology or medicine, provides a tool for obtaining high-resolution images of very close objects [2]. Many of these imaging techniques share a common feature: they measure the intensity of the light coming from the scene under consideration. However, it is sometimes required to measure other physical quantities, like the phase of the optical field, its spectral content, or its polarization state. The spectral content of a sample is normally used to obtain information about its material components. Polarization, that is, the knowledge of the vector nature of light, gives information about surface features such as shape, shading, and roughness of an object [3]. Advanced imaging techniques make it possible to acquire multi-dimensional images, which provide information not only about the spatial distribution of intensity but also about the previously mentioned primary physical quantities associated with an optical field.

In general, the measurement of multi-dimensional images involves the acquisition of a huge amount of information, which causes both storage and transmission difficulties [4]. In addition, techniques such as multispectral or hyperspectral imaging, require a sequential acquisition of images in the spectral domain, leading to a dramatic increase in measurement time. A recent approach to hyperspectral and polarimetric imaging is based on the use, respectively, of miniaturized spectral and polarimetric filters [5,6] that are incorporated to each pixel of the sensor, which allows acquiring multi-dimensional images in one shot. However, the development of such systems implies the use of high-end micro-optical components.

In this chapter, we describe several single-pixel multi-dimensional imaging systems based on *compressive sensing* (CS), a new sampling paradigm that has revolutionized data acquisition protocols, enabling us to start the signal compression at the measurement stage. In Section 16.2 we show how single-pixel imaging techniques work and how CS can boost their performance. In Sections 16.3 to 16.5 we describe single-pixel architectures that use off-the-shelf components in the fields of polarimetry, multispectral imaging, and spectropolarimetry.

16.2 Single-Pixel Imaging and Compressive Sensing

The operation principle of single-pixel imaging can be briefly described as follows. Let us consider a sample object, whose N -pixel image is arranged in an $N \times 1$ column vector, \mathbf{x} . That image can be expressed in terms of a basis of functions, $\Psi = \{\Psi_\ell\}$ ($\ell=1, \dots, N$). In mathematical terms, $\mathbf{x} = \Psi \cdot \mathbf{s}$, where Ψ is a $N \times N$ matrix that has the vectors $\{\Psi_\ell\}$ as columns and \mathbf{s} is the $N \times 1$ vector which contains the expansion coefficients of \mathbf{x} in the chosen basis. Single-pixel cameras exploit the fact that those coefficients can be measured by using detectors with no spatial resolution. The acquisition process is governed by a spatial light modulator (SLM), which generates a set of patterns directly related to the selected basis. The irradiance corresponding to the inner product between the patterns and the object provides the coefficients of the image expansion.

In recent years, the introduction of CS has dramatically improved the performance of these single-pixel architectures. CS exploits the fact that natural images tend to be sparse, that is, only a small set of the expansion coefficients is nonzero when a suitable basis is chosen [7]. In this way, images can be retrieved without measuring all the projections of the object on the chosen basis. The mathematical formulation behind CS ensures that the object under study, \mathbf{x} , can be reconstructed from just a random subset of the expansion coefficients that make up \mathbf{s} . To this end, we randomly choose M different functions of the basis ($M < N$) and measure the projections of the object. This process can be expressed in matrix form as

$$\mathbf{y} = \Phi \cdot \mathbf{x} = \Phi(\Psi \cdot \mathbf{s}) = \Theta \cdot \mathbf{s}, \quad (16.1)$$

where \mathbf{y} is a $M \times 1$ vector which contains the measured projections and Φ is a $M \times N$ matrix called sensing matrix. Each row of Φ is a function of Ψ chosen randomly, and the product of Φ and Ψ gives the matrix Θ acting on \mathbf{s} . If the chosen basis is orthonormal, every row of Θ randomly selects a unique element of \mathbf{s} . As $M < N$, the underdetermined matrix relation obtained after the measurement process is resolved through an off-line algorithm. The best approach to recover the object is based on the minimization of the ℓ_1 -norm of \mathbf{s} subjected to the constrain given by Eq. (16.1), that is, that the solution given by the algorithm has to be compatible with the performed measurements. In this case, the proposed reconstruction \mathbf{x}^* is

given by the optimization program

$$\min_{\mathbf{x}^*} \|\Psi^{-1} \mathbf{x}^*\|_{\ell_1} \text{ subject to } \Phi \mathbf{x}^* = \mathbf{y}. \quad (16.2)$$

In the experiments described in this paper, the chosen basis is a family of binary intensity patterns derived from the Walsh–Hadamard basis. This basis has proved to be suitable for single-pixel architectures due to being easily implemented on a SLM. A Walsh–Hadamard matrix of order N (\mathbf{H}_N) is a $N \times N$ matrix with ± 1 entries that satisfies $\mathbf{H}_N^T \mathbf{H}_N = N \cdot \mathbf{I}_N$, where \mathbf{I}_N is the identity matrix and \mathbf{H}_N^T denotes transposed matrix. Walsh–Hadamard matrices form an orthonormal basis that was first proposed in image coding and transmission techniques [8]. By shifting and rescaling the different \mathbf{H}_N , it is possible to generate binary waveforms taking values 0 or 1 that can be simply encoded onto the SLM as an intensity modulation.

A suitable SLM for a single-pixel camera is a display formed by voltage-controlled liquid-crystal (LC) cells, such as those found in video projection systems [9]. Another option is a digital micromirror device (DMD), composed of an array of micromirrors that can rotate between two positions. In this way, only selected portions of the incoming light beam are reflected in a given direction [10]. Both devices are used in the different optical systems described in the following sections. Regarding detection, in general, a photodiode is used as single-pixel camera, which measures the irradiance of the light coming from an object for each pattern generated by the SLM. In the optical systems described in this chapter other single-pixel detectors, such as a beam polarimeter or a fiber spectrometer, are used.

16.3 Single-Pixel Polarimetric Imaging

Polarimetric imaging (PI) has the aim of measuring spatially resolved polarization properties of a light field, an object, or an optical system [11]. These properties are usually the Stokes parameters of light (passive imaging polarimeters) or the Mueller matrix that characterizes a sample or a system (active imaging polarimeters). The use of PI includes a great variety of optical applications, like scene analysis, target detection [3], polarization-sensitive microscopy [12], or segmentation of rough surfaces [13], among others. Polarimetric techniques have been used in the field of biomedical imaging for enhanced visualization of biological samples at different depths [14], as well as *in vivo* detection and diagnosis of cancerous tumors in tissues [15,16]. PI can be also combined with optical coherence tomography [17] and ophthalmic adaptive optics [18].

In this chapter, we describe how the concept of single-pixel imaging by CS has been extended to the design of a passive polarimetric camera. In particular, we describe a PI system able to measure spatially resolved Stokes parameters by means of a commercial beam polarimeter [19]. This commercial beam polarimeter is designed for free-space and fiber-based measurements, and provides the state of polarization (SOP) of an optical beam as a whole; that is, without spatial resolution. The PI system exhibits high dynamic range (up to 70 dB), broad wavelength range, and high accuracy on the Poincaré sphere, thanks to the use of the beam polarimeter. This fact simplifies the design and optimization procedures of current polarimetric cameras based on pixelated image sensors [16,20]. A programmable SLM is at the heart of this imaging polarimeter. This modulator controls the time-multiplexed acquisition process required by a single-pixel imaging scheme. The amount of acquired data is minimized by the

application of a CS algorithm, which implies a proper selection of light patterns generated by the SLM, in accordance with the theory briefly described in Section 16.2.

A *Stokes polarimeter* (SP) is a device that measures the irradiance of a light beam whose SOP is modulated by a *polarization state analyzer* (PSA). In the commercial SP used here, which is sketched in Fig. 16.1, the PSA is formed by two voltage-controlled liquid-crystal variable retarders (LCVR₁ and LCVR₂) and a polarizing beam splitter (PBS). Two photodiodes (PD₁ and PD₂) are respectively located at the output ports of the PBS. The sum of the signals of PD₁ and PD₂ gives the total irradiance I_0 impinging onto the SP, despite of slight (and measurable) losses. The application of the Stokes–Mueller formalism allows obtaining the SOP of the input light, which is given by the Stokes vector $(I_0, S_1, S_2, S_3)^T$. If the retardances of LCVR₁ and LCVR₂ are δ_1 and δ_2 , respectively, the irradiance I_{PD} measured by one photodiode is given by

$$I_{PD}(\delta_1, \delta_2) = m_{00}(\delta_1, \delta_2)I_0 + \sum_{i=1}^3 m_{0i}(\delta_1, \delta_2)S_i. \quad (16.3)$$

In this expression m_{0k} ($k = 0, \dots, 3$) are the voltage-dependent elements of the first row of the PSA Mueller matrix. A proper calibration process, usually performed by the manufacturer, can be used to determine these elements. The description of such a process is out of the scope of the present study [21]. By a sequential reconfiguration of the PSA, the SOP of the incoming light is derived through the measurement of at least three values of I_{PD} , together with the irradiance I_0 . In commercial devices, the LCVRs perform a wide retardance sweep. In this way, the input SOP is obtained through a least-squares fitting routine to minimize measurement errors [22]. The quantities registered by the SP are usually the normalized version of the Stokes parameters, $\sigma_i = S_i/I_0$ ($i = 1, \dots, 3$). It should be noted that the retardance of the LCVRs strongly depends on the light frequency, so the device calibration is valid for a given wavelength and must be repeated if the light source spectrum is changed.




Figure 16.1 Scheme of the Stokes polarimeter acting as a single-pixel detector

A polarimetric detector with no spatial structure, such as the SP detector in Fig. 16.1, can be adapted to perform PI with the aid of the single-pixel architecture discussed on Section 16.2. The idea is simple: the problem of measuring a spatial-dependent Stokes vector is equivalent to resolving three times the CS algorithm of single-pixel imaging. This is possible because the linearity of Eq. (16.3) implies that each Stokes parameter S_i^{SP} provided by the SP is the sum of the values taken by S_i at each point of the input light beam. As a consequence, the measurement process expressed by Eq. (16.1) can be separately applied to each Stokes parameter whose spatial distribution (described by an N -pixel matrix) is recovered using $M < N$ polarimetric measurements. A layout of the PI system is depicted in Fig. 16.2. A collimated (unpolarized) laser beam passes through an LC-SLM, programmed to generate a set of intensity patterns. Just after the modulator there is a polarization object (PO), which produces a space-variant Stokes vector. As an LC-SLM is a polarization-dependent device; it is sandwiched between properly oriented linear polarizers (P_1 and P_2), so the object is illuminated with linearly polarized light. The light emerging from the object is guided to the SP by means of an afocal optical system, like an inverted beam expander, which fits the beam width to the typically small entrance window of the SP. This coupling optic ensures that all the light emerging from the object is collected by the SP and it preserves the normal incidence, which contributes to the optimal performance of the polarimeter.

The light source used in this experiment was an He-Ne laser emitting at 632.8 nm. The LC-SLM was a transmissive twisted nematic LCD (TNLCD) with SVGA resolution (800×600 pixels) and a pixel pitch of $32 \mu\text{m}$. To reach an intensity modulation regime, the LC-SLM was sandwiched between two linear polarizers, respectively oriented parallel and normal to the input molecular director of the TNLCD, which was previously determined by a polarimetric technique [23]. In this configuration, the LC-SLM worked as a spatial intensity modulator. Pixels were individually addressed by sending gray-level images to the TNLCD. Each gray level corresponded to a transmitted intensity level, ranging from the dark state (extinction) to the bright state (maximum transmission).




Figure 16.2 Experimental setup for the polarimetric single-pixel camera. One of the binary intensity patterns displayed by the SLM is also shown

Table 16.1 Technical specifications of the polarimetric camera

Wavelength	632.8 nm
Image resolution	64×64 pixels
Compression ratio	3:1
Pixel pitch	$64 \mu\text{m}$

In order to perform CS, the Walsh–Hadamard functions were chosen as the reconstruction basis Ψ . This election was particularly useful because the intensity patterns $\{\Phi_m\}$ generated by the TNLCD were binary masks (see Fig. 16.2). The corresponding images addressed to the display had a resolution of 64×64 cells, and the cell pitch was $64 \mu\text{m}$. The number of binary patterns displayed onto the TNLCD was 1225, which represents $\sim 30\%$ of the Nyquist criterion. Custom software written with LabVIEW was used to synchronize the SP with the modulator. These technical specifications are summarized in Table 16.1. For each realization, the values of the Stokes parameters $\{S_i^{\text{SP}}\} (i = 1, \dots, 3)$, as well as the signals of PD_1 and PD_2 , are measured. The maximum measurement rate of the SP (10 Stokes vectors per second) was the speed limiting factor, since the refreshing frequency of the TNLCD was 60 Hz.

The selected object, shown in Fig. 16.3(a), was a cellophane film, acting as linear retarder, attached to an amplitude mask, which reproduces the logotype of the university UJI. Linearly




Figure 16.3 (a) High-resolution representation of the object under study, which is an amplitude mask representing the logo of the university UJI, with a cellophane film over the letter in gray shading (was yellow). Pseudocolored pictures showing the distribution of the Stokes parameters are shown in (b), (c), and (d). *Source:* V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis 2012, Figure 3. Reproduced with permission from The Optical Society

polarized light emerging from the polarizer P_2 illuminated the object. An inhomogeneous polarization distribution was generated by covering just the capital letter J with cellophane. With the introduction of this element, the polarization of the light coming from this letter was approximately rotated by the cellophane film. The parameters of the polarization ellipse (the azimuth α and the ellipticity e) of the light passing through the J were measured by the SP (blocking the light emerging from the other part of the object). For this measurement, the TNLCD was configured in its bright state. The results were $\alpha_J = 8.62^\circ$ and $e_J = -0.07$. Repeating the process for letters U and I, the measured parameters were $\alpha_{U,I} = 42.22^\circ$ and $e_{U,I} = 0.003$.

Figures 16.3(b)–N.3(d) show pseudocolor plots for the normalized Stokes parameters. These images exhibit a clear uniformity within the different parts of the object. The spatial distributions for α and e were calculated from the Stokes parameters through conventional expressions (see, for example, [24]). The mean values of the ellipse parameters for each part of the object were ($\langle\alpha_J\rangle = 2.5^\circ \pm 1.4^\circ$, $\langle e_J\rangle = -0.08 \pm 0.02$) and ($\langle\alpha_{U,I}\rangle = 43.6^\circ \pm 1.1^\circ$, $\langle e_{U,I}\rangle = -0.01 \pm 0.04$). The assigned uncertainties were the standard deviations of each distribution. These results were in good agreement with the values previously measured by the SP. The major discrepancy was found for α_J ($\sim 6^\circ$), which only represents $\sim 3\%$ of the total range of values (from -90° to 90°) that can be taken by the azimuth.

These results demonstrate the possibility of performing spatially resolved Stokes polarimetry with the aid of CS. In particular, the system described here converts a commercial beam SP into a polarimetric imager. Although this system is based on liquid crystal elements, the method is valid for other types of polarimeters, provided that the selected device is itself spatially homogeneous, and the relationship between the measured signals and the Stokes parameters is linear, as in Eq. (16.3). Concerning the acquisition process, a TNLCD is used to project the intensity patterns over the object. Another possibility is to employ SLMs insensitive to polarization, like a DMD, as is done in the optical systems described in the following sections. The combination of DMDs with fast SPs may lead to the design of PI systems working at very high frequencies (~ 1 KHz), opening the door to near-real-time applications.

16.4 Single-Pixel Multispectral Imaging

Multispectral imaging (MI) is a useful optical technique that provides two-dimensional images of an object for a set of specific wavelengths within a selected spectral range [1]. Dispersive elements (prisms or gratings), filter wheels, or tunable band-pass filters, are typical components used in MI systems to acquire image spectral content [25]. Multispectral imaging provides both spatial and spectral information of an object and represents a powerful analysis tool in different scientific fields as medicine [26], pharmaceuticals [27], astronomy [28], and agriculture [29]. In industry, new techniques have emerged that use VIS and NIR imaging to make quality and safety control, for example, in the detection of surface properties on fruits [30].


The second optical system described in this chapter is a CS imaging system able to provide spatially resolved information about the spectrum of the light reflected by an object [31]. A fiber spectrometer with no spatial resolution is used as a single-pixel detector. Now, the key element of the system that makes possible the CS acquisition process, is a digital micromirror device (DMD). The modulator sequentially generates a set of binary intensity patterns that sample the image of the object under consideration. The acquired data is subsequently processed to obtain a multispectral data cube.

A layout of the spectral camera is shown in Fig. 16.4(a). A white-light source illuminates a sample and a CCD camera lens images the object on a DMD, which is a reflective spatial light modulator that selectively redirects parts of an input light beam [32]. The DMD consists of an array of electronically controlled micromirrors, positioned over a CMOS memory cell, which can rotate about a hinge, as is schematically depicted in Fig. 16.4(b). The angular position of each micromirror admits two possible states ($+12^\circ$ and -12° respect to a common direction), depending on the binary state (logic 0 or 1) of the corresponding CMOS memory cell contents. As a consequence, light can be reflected at two angles depending on the signal applied to the mirror. The DMD used in this system is a Texas Instrument device (DLP Discovery 4100) with a resolution of 1920×1080 micromirrors, the panel size of the display is $0.95''$, the mirror pitch is $10.8 \mu\text{m}$ and the fill factor is greater than 0.91. The optical axis of the “optical system 1” forms an angle with respect to the orthogonal direction to the DMD panel that approximately corresponds to twice the tilt angle of the device mirrors (24°). As is shown in Fig. 16.4(c), a micromirror oriented at $+12^\circ$ orthogonally reflects the light, appearing as a bright pixel (ON state). In their turn, micromirrors oriented at -12° work as dark pixels (OFF state). The light emerging from the bright pixels of the DMD is collected by “the optical system 2” (see Fig. 16.4a). This lens system couples the light into a silica multimode fiber with a diameter of $1000 \mu\text{m}$, which is connected to a commercial concave grating spectrometer (Black Comet CXR-SR from Stellarnet). The spectral range of the fiber ranges from 220–1100 nm. The wavelength resolution of this spectrometer is 8 nm (with a slit of $200 \mu\text{m}$) and the maximum *signal-to-noise ratio* (SNR) is 1000:1. Technical specifications for this setup are summarized in Table 16.2.


An example of spectral image with resolution 256×256 pixels was performed with a sample object composed of an unripe cherry tomato together with a red one. The illumination source was a xenon white light lamp. The Walsh–Hadamard patterns addressed to the DMD had $N = 65536$ unit cells. Each unit cell was composed of 2×2 DMD pixels. With this resolution, the number of measurements was chosen to be $M = 6561$, which corresponds to a compression rate of 10:1 ($M \approx 0.1N$). The integration time of each spectrometer measurement was 300 ms.

In order to determine the object spectral reflectance, a spectrum was taken from a white reference (Spectralon diffuse 99% reflectance target from Labsphere, Inc.) to normalize the measured spectra during the CS acquisition process. In the case of plants, the chlorophyll content in leaves can be recovered from the spectral reflectance [33,34]. The data collected by the spectrometer for wavelengths lower than 500 nm, clearly affected by noise, imposed an inferior boundary to the usable spectral range. The results of the CS reconstruction for 15 spectral channels are shown in Fig. 16.5 (Plate 26). The selected central wavelengths λ_0 in the visible spectrum (VIS) range from 510–680 nm. The bandwidth of each spectrum channel was 10 nm ($\lambda_0 \pm 5$ nm). The recovered images were pseudo-colored and the color assignment (the wavelength to RGB transform) was carried out with the aid of standard XYZ Color-Matching Functions [35]. The CS algorithm provided an acceptable reconstruction in the near-infrared spectrum, around 860 nm. This is presented by means of a gray-level image. Figure 16.5 also includes a RGB image of the object.


The quality of the images obtained with the single-pixel camera in Fig. 16.4(a) is evaluated by performing a multispectral imaging experiment, by sending Walsh–Hadamard patterns of



(a)



(b)



(c)

Figure 16.4 (a) Optical system for multispectral imaging using a single-pixel detector, (b) Individual DMD micromirror showed in the transverse view indicating its two possible orientations. (c) Working mode scheme of the DMD

Table 16.2 Technical specifications of the multispectral camera

Wavelength range	505–865 nm
Number of channels	15
Image resolution	256 × 256 pixels
Compression ratio	10:1
Integration time	300 ms
Pixel pitch	21.6 μm




Figure 16.5 (Plate 26) Multispectral data cube reconstructed using CS. In the VIS band, the reflectance for each spectral channel is represented by means of a 256×256 pseudo-color image. In the NIR band we show a gray-scale representation. A colorful image of the scene made up from the conventional RGB channels is also included. *Source:* F. Soldevila, E. Irlés, V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis 2013, Figure 4. Reproduced with permission from Springer. *See plate section for the color version*

64×64 unit cells ($N = 4096$) to the modulator, each one composed of 8×8 DMD pixels. In this case, the sample scene is constituted by two small square color objects. The number of measurements was $M = 4096$ ($M = N$), which allowed us to fulfill the Nyquist criterion. Eight central wavelengths, λ_0 , were selected in the visible spectrum. The bandwidth of the corresponding spectral channels was 20 nm ($\lambda_0 \pm 10$ nm). Aside from the channels at the boundaries of the spectral range under consideration, the values of λ_0 correspond to peak emissions of commercial light-emitting diodes. The object spectral reflectance was determined again by means of the previously used white reference. The integration time of the spectrometer was set at 300 ms.

For each spectral channel, the off-line CS algorithm was resolved with the complete set of measurements. After a suitable filtering, the recovered matrix served as a reference (lossless) image $I_{ref}(i, j)$, where (i, j) indicates the location of an arbitrary image pixel. The reconstruction process was then repeated using decreasing fractions of the total number of pixels. In particular, the value of M was varied from 5 to 90% of N , and the fidelity of the reconstructed images was estimated by calculating the *mean square error* (MSE), given by

$$MSE = \frac{1}{N} \sum_i \sum_j [I(i, j) - I_{ref}(i, j)]^2, \quad (16.4)$$

where $I(i, j)$ is the noisy image obtained for a given value of M . The *peak signal-to-noise ratio* (PSNR) was used to evaluate the quality of the reconstruction. It is defined as the ratio between the maximum possible power of a signal and the power of the noise that affects the fidelity of its representation. In mathematical terms, [36]

$$PSNR = 10 \log \left(\frac{I_{max}^2}{MSE} \right) = 20 \log(I_{max}) - 10 \log(MSE) \quad (16.5)$$

Here, I_{max} is the maximum pixel value of the reference image. For each spectral channel, the reference images were represented by 2^8 gray-levels, so $I_{max} = 255$. Figures 16.6(a) and (b) show the curves for the MSE and the PSNR versus M for the different values of λ_0 . Both figures point out that the image quality improves as the number of measurements grows and approximates to the Nyquist limit. However, it should be noted that the slope of both curves becomes visibly smoother for all spectral channels when $M \geq 0.4N$. In the case, for instance, of $\lambda_0 = 610$ nm, $MSE = 0.13I_{max}^2$ and $PSNR = 28.72$ dB for $M = 0.4N$ (the PSNR for $0.9N$ is only somewhat greater, 29.10 dB).

Although this single-pixel camera acquires sequentially the spatial information of the input object, it allows collection of all the spectral content at once, in contrast to those cameras based on tunable band-pass filters, which perform a wavelength sweep to measure the spectral information. In addition, the number of channels, their spectral resolution, and the total wavelength range of the single-pixel system are those provided by the spectrometer working as detector. This fact makes possible to exploit the high performance of commercially available devices. Thus, the spectral system can, in principle, cover the whole VIS spectrum and part of the NIR range (up to 1.1 microns), while conventional multispectral systems require pixelated sensors specifically designed for the infrared range (like InGaAs cameras).

Apart from the detector, the illumination is another key element to ensure a minimum signal along the selected spectral range. The use of a high power Xe arc lamp provides a continuous




Figure 16.6 (a) MSE and (b) PSNR of the recovered images versus the number of measurements. Each curve corresponds to a spectral channel. *Source:* F. Soldevila, E. Irlés, V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis 2013, Figure 2. Reproduced with permission from Springer

and roughly uniform spectrum across the VIS region. However, the decreasing source irradiance at the “blue” side of the VIS spectrum, as well as the low reflectance of samples at that region, limits the spectral range to wavelengths higher than 500 nm.

This single-pixel multispectral camera presents a trade-off between image resolution and acquisition time. Increasing the illumination level or considering lower integration times (by a reduction of the spectral resolution) can make the acquisition time to drop by at least one order of magnitude. A comparable trade-off can be found in cameras based on acousto-optic or liquid crystal tunable filters, where the higher spectral resolution (number of channels), the longer acquisition time, with a strong dependence on the exposure time of the pixelated sensor used as a detector. A hyperspectral camera (i.e., a camera with more than 100 spectral channels) can take a few minutes in acquiring a data cube for image resolutions similar to those presented in this work [37].

16.5 Single-Pixel Spectropolarimetric Imaging

In certain applications, MI can be improved by adding spatially resolved information about the light polarization. Multispectral polarimetric imaging facilitates the analysis and identification

of soils [38], plants [39], and surfaces contaminated with chemical agents [40]. In the field of biomedical optics, multispectral polarimetric imaging has been applied to the characterization of human colon cancer [41] or the pathological analysis of skin [42]. In many cases, polarimetric analysis can be performed by just including a linear polarizer in an imaging system to record images for various selected orientations of its transmission axis [42,43]. A simple configuration that includes two orthogonal polarizers integrated in a spectral system has been used for noninvasively imaging of the microcirculation through mucus membranes and on the surface of solid organs [43]. An illustrative example of a spectral camera with polarimetric capability is a system that combines an acousto-optic tunable filter with a liquid-crystal based polarization analyzer [44].

In this section we describe two different optical architectures for spectropolarimetric imaging. In the first one, polarimetry is performed by using a rotating linear analyzer in front of the detector, which leads to a linear polarization spectral imager. In the second one, the optical system is constituted by a fixed polarizer and two voltage-controlled variable retarders to spatially resolve the circular polarization component of light. In this way, this single-pixel multispectral system works as an imaging full-Stokes meter for each spectral channel.

16.5.1 Multispectral Linear Polarimetric Camera

A scheme of the multispectral linear polarimetric camera is depicted in Fig. 16.7. The optical system is similar to that described in the previous section, see Fig. 16.4(a), but now includes a linear polarizer. The sample scene is constituted by two square capacitors with a width of 7 mm. A xenon white light lamp is used again as illumination source. The light emerging from each




Figure 16.7 Optical system to obtain polarimetric multispectral imaging by using a single-pixel detector

Table 16.3 Technical specifications of the multispectral linear polarimetric camera

Wavelength range	470–700 nm
Number of channels	8
Image resolution	128 × 128 pixels
Compression ratio	5:1
Integration time	500 ms
Pixel pitch	43.2 μm

element of the scene had a different linear polarization. This spatial distribution of polarization was achieved by means of a linear polarizer located after the object, which had its area split in two parts, each of which with its transmission axis oriented at orthogonal directions (0 and 90°, respectively). The resolution of the patterns generated by the DMD was 128 × 128 unit cells ($N = 16384$) composed of 4 × 4 DMD pixels. The number of measurements was $M = 3249$, which corresponds to ~20% of N (i.e., a compression rate of 5:1). The integration time of the spectrometer was fixed to 500 ms. Eight central wavelengths λ_0 were selected in the VIS spectrum. The bandwidth of the channels was 20 nm ($\lambda_0 \pm 10$ nm). For each channel, four orientations of the polarization analyzer were sequentially considered in separated measurement series. The technical specifications of this camera are outlined in Table 16.3.

Figure 16.8 (Plate 27) shows the image reconstructions with the optical system in Fig. 16.7. Each column of the figure corresponds to a spectral channel and each row shows the results for a given orientation of the analyzer. A colorful image of the object is also shown (see Plate 27 in the plate section). This RGB image was made up from the data taken for the second configuration of the analyzer (45°). The result for 680 nm is presented by means of a gray-level image due to its proximity to the near infrared range.

The polarizer included in the single-pixel optical system in Fig. 16.7 limits the total spectral range, since the optical behavior of polarizing films is wavelength dependent. As a consequence, the upper boundary of the spectral range is ~700 nm. However, the use of high grade crystalline polarizers can solve this limitation.

16.5.2 Multispectral Full-Stokes Imaging Polarimeter

In principle, it is possible to obtain information about the spatial distribution of the Stokes parameters of light, S_i ($i = 0, \dots, 3$) from polarimetric images recovered for each spectral channel. In the previous optical system, as a linear polarizer is used as analyzer, the spatial distribution of S_0 , S_1 , and S_2 can be straightforwardly derived. However, a full Stokes polarimeter requires adding at least a linear retarder.

The scheme of a full-Stokes polarimeter is shown in Fig. 16.9. A white light beam generated by a xenon lamp is collimated by a lens and illuminates a sample object, whose image is formed on a digital micromirror device (DMD) by a pair of lenses. The light emerging from the DMD is guided to a single-pixel detector with the aid of a fourth lens. In order to achieve both polarimetric and spectral information, the single-pixel detector consists of two *liquid crystal variable retarders* (LCVR) (liquid crystal variable retarder from Meadowlark) with




Figure 16.8 (Plate 27) Multispectral image cube reconstructed by CS algorithm for four different configurations of the polarization analyzer. The RGB image of the object is also included. In the VIS spectrum all channels are represented by pseudo-color images and a gray-scale representation is used for the wavelength closer to the NIR spectrum. *Source:* F. Soldevila, E. Irlés, V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis 2013, Figure 5. Reproduced with permission from Springer. *See plate section for the color version*




Figure 16.9 Scheme of the multispectral full-Stokes imaging polarimeter

their slow axis oriented at 45° and 0° , followed by a linear analyzer with its transmission axis oriented at 45° and a commercial fiber spectrometer (Black Comet CXRSR from StellarNet). Each LCVR is precalibrated to introduce controlled retardances in each chromatic channel of interest. The commercial fiber spectrometer is the same used in the preceding section.

By acquiring four images for different retardances of the LCVRs, it is possible to compute the Stokes parameters of each pixel of the scene. The CS algorithm provides an intensity map for the scene, which corresponds to the spatial distribution of the Stokes parameter S'_0 . By using Stokes–Mueller calculus, it is possible to relate, for each pixel, the value of the Stokes vector with the measured irradiance S'_0 . From the Mueller matrix expressions of a retarder wave plate and a linear polarizer, it is possible to show that the relationship between the recovered irradiance and the original Stokes parameters is

$$S'_0(2\delta_1, 2\delta_2) = \frac{1}{2}S_0 + \frac{1}{2} \sin(2\delta_1) \sin(2\delta_2)S_1 + \frac{1}{2} \cos(2\delta_2)S_2 - \frac{1}{2} \cos(2\delta_1) \sin(2\delta_2)S_3, \quad (16.6)$$

where $2\delta_1$ and $2\delta_2$ are the phase retardances introduced, respectively, by the two LCVRs. Equation 16.6 establishes an undetermined system with four unknown quantities (the Stokes parameters of the incident light). In order to solve that system, a minimum of four pairs of phase retardances must be applied to the two LCVRs. After the off-line reconstructions, the Stokes vector in each point of the scene is given by $\mathbf{S}'_0 = \mathbf{M} \cdot \mathbf{S}$, where

$$\mathbf{M} = \frac{1}{2} \begin{pmatrix} 1 & \sin(2\delta_1^{(1)}) \sin(2\delta_2^{(1)}) & \cos(2\delta_2^{(1)}) & -\cos(2\delta_1^{(1)}) \sin(2\delta_2^{(1)}) \\ 1 & \sin(2\delta_1^{(2)}) \sin(2\delta_2^{(2)}) & \cos(2\delta_2^{(2)}) & -\cos(2\delta_1^{(2)}) \sin(2\delta_2^{(2)}) \\ 1 & \sin(2\delta_1^{(3)}) \sin(2\delta_2^{(3)}) & \cos(2\delta_2^{(3)}) & -\cos(2\delta_1^{(3)}) \sin(2\delta_2^{(3)}) \\ 1 & \sin(2\delta_1^{(4)}) \sin(2\delta_2^{(4)}) & \cos(2\delta_2^{(4)}) & -\cos(2\delta_1^{(4)}) \sin(2\delta_2^{(4)}) \end{pmatrix}. \quad (16.7)$$

The subscripts in the elements of \mathbf{M} relate to each one of the LCVRs and the superscripts denote each one of the four acquisitions. The solution of this linear system provides the spatial distribution of the Stokes parameters.

As a direct application of the single-pixel spectral Stokes polarimeter, a photoelasticity measurement on a piece of polystyrene was carried out. In the process of fabrication, the piece of polystyrene is shaped in a certain form. Due to this, the material presents stresses that cause a spatial distribution of birefringence. This distribution can be seen when the piece is placed between crossed linear polarizers and illuminated with white light, as can be seen in Fig. 16.10.

The Walsh–Hadamard patterns addressed to the DMD had a resolution of 128×128 unit cells ($N = 16384$). Each unit cell was composed of 4×4 DMD pixels. The number of measurements was chosen to be $M = 3249$, which corresponds to $\sim 20\%$ of N (a measurement rate of 5:1). The integration time of each spectrometer measurement was set to 20 ms. These technical specifications of the camera shown in Fig. 16.9 are summarized in Table 16.4. In Fig. 16.11 (Plate 28) we show the experimental results of the distribution of the normalized Stokes parameters for eight chromatic channels, each one with 20 nm width ($\lambda_0 \pm 10$ nm). To simplify data display, image reconstructions are arranged in a table. Each column corresponds




Figure 16.10 Color picture of the polystyrene sample. It is placed between two crossed linear polarizers and illuminated with white light. Color fringes are a consequence of the different states of polarization produced by the stress in the piece. The square indicates the region of interest imaged by the spectral camera

Table 16.4 Technical specifications of the multispectral full-Stokes imaging polarimeter

Wavelength range	450–730 nm
Number of channels	8
Image resolution	128 × 128 pixels
Compression ratio	5:1
Integration time	20 ms
Pixel pitch	43.2 μm




Figure 16.11 (Plate 28) Spatial distribution of the Stokes parameters of the polystyrene piece. Each distribution is represented by a pseudo-colored 128 × 128 pixels picture. The values range from -1 (blue) to 1 (red). See plate section for the color version

to a spectral channel and each row shows the spatial distribution of a normalized Stokes parameter.

To simplify data display, image reconstructions are arranged in a table. Each column corresponds to a spectral channel and each row shows the spatial distribution of a normalized Stokes parameter.

As can be seen in Fig. 16.11, the expected fringe distribution of the Stokes parameters is recovered, if we compare this result with that in Fig. 16.10. For wavelengths near the IR, reconstructions are a bit noisy. This is caused by the low amount of light the source emits in this zone of the spectrum, which causes the reconstructions on these channels to have lower SNR. This problem can be solved by increasing the integration time of the spectrometer, but this makes measurement times much greater and the reconstructions on the channels in the visible region of the spectrum do not improve their SNRs. Using a light source with a flatter spectrum will solve the quality drop near the IR region.

16.6 Conclusion

We have described several multi-dimensional single-pixel imaging techniques providing the spatial distribution of multiple optical properties of an input scene. In all cases, the key element of the optical system is a SLM that sequentially generates a set of intensity light patterns to sample the input scene. In this way, it is possible to apply the theory of compressive sampling to data acquired with a single-pixel sensor. In particular, we have described a single-pixel hyperspectral imaging polarimeter. This system is able to provide spatially resolved measurements of Stokes parameters for different spectral channels. In this case the spatial light modulator is a digital micromirror device, and the sensor is composed by polarizing elements followed by a commercial fiber spectrometer. Experimental results for color objects with an inhomogeneous polarization distribution show the ability of the method to measure the spatial distribution of the Stokes parameters for multiple spectral components.

Acknowledgments

This work has been partly funded by the Spanish Ministry of Education (project FIS2010-15746) and the Excellence Net from the Generalitat Valenciana about Medical Imaging (project ISIC/2012/013). Also funding from Generalitat Valenciana through Prometeo Excellence Programme (project PROMETEO/2012/021) is acknowledged.

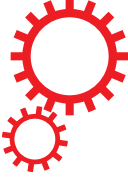
References

- [1] Brady D., *Optical Imaging and Spectroscopy*, 1st edn. John Wiley & Sons, Ltd, 2009.
- [2] Weissleder R. and M. J. Pittet, "Imaging in the era of molecular oncology," *Nature*, vol. **452**, no. 7187, pp. 580–589, Apr. 2008.
- [3] Tyo J. S., D. L. Goldstein, D. B. Chenault, and J. A. Shaw, "Review of passive imaging polarimetry for remote sensing applications.," *Appl. Opt.*, vol. **45**, no. 22, pp. 5453–5469, Aug. 2006.
- [4] Brady D. J., M. E. Gehm, R. A. Stack, D. L. Marks, D. S. Kittle, D. R. Golish, *et al.*, "Multiscale gigapixel photography," *Nature*, vol. **486**, no. 7403, pp. 386–389, Jun. 2012.
- [5] Geelen B., N. Tack, and A. Lambrechts, "A snapshot multispectral imager with integrated tiled filters and optical duplication," pp. 861314–861313, 2013.

- [6] Zhao X. and F. Boussaid, "Thin photo-patterned micropolarizer array for CMOS image sensors," *Photonics Technol.*, vol. **21**, no. 12, pp. 805–807, 2009.
- [7] Candès E. and M. Wakin, "An introduction to compressive sampling," *Signal Process. Mag. IEEE*, no. March 2008, pp. 21–30, 2008.
- [8] Pratt W., J. Kane, and H. Andrews, "Hadamard transform image coding," *Proc. IEEE*, vol. **57**, no. 1, 1969.
- [9] Magalhães F., F. M. Araújo, M. V. Correia, M. Abolbashari, and F. Farahi, "Active illumination single-pixel camera based on compressive sensing," *Appl. Opt.*, vol. **50**, no. 4, pp. 405–414, Feb. 2011.
- [10] Duarte M. and M. Davenport, "Single-pixel imaging via compressive sampling," *Signal Process.*, March 2008, pp. 83–91, 2008.
- [11] Solomon J. E., "Polarization imaging," *Appl. Opt.*, vol. **20**, no. 9, pp. 1537–1544, May 1981.
- [12] Oldenbourg R., "A new view on polarization microscopy," *Nature*, vol. **381**, no. 6585, pp. 811–82, Jun. 1996.
- [13] Terrier P., V. Devlaminck, and J. M. Charbois, "Segmentation of rough surfaces using a polarization imaging system," *J. Opt. Soc. Am. A. Opt. Image Sci. Vis.*, vol. **25**, no. 2, pp. 423–430, Mar. 2008.
- [14] Demos S. G. and R. R. Alfano, "Optical polarization imaging," *Appl. Opt.*, vol. **36**, no. 1, pp. 150–155, Jan. 1997.
- [15] Baba J. S., J.-R. Chung, A. H. DeLaughter, B. D. Cameron, and G. L. Coté, "Development and calibration of an automated Mueller matrix polarization imaging system," *J. Biomed. Opt.*, vol. **7**, no. 3, pp. 341–349, Jul. 2002.
- [16] Laude-Boulesteix B., A. De Martino, B. Drévilion, and L. Schwartz, "Mueller polarimetric imaging system with liquid crystals," *Appl. Opt.*, vol. **43**, no. 14, pp. 2824–2832, May 2004.
- [17] de Boer J. F. and T. E. Milner, "Review of polarization sensitive optical coherence tomography and Stokes vector determination," *J. Biomed. Opt.*, vol. **7**, no. 3, pp. 359–371, Jul. 2002.
- [18] Song H., Y. Zhao, X. Qi, Y. T. Chui, and S. A. Burns, "Stokes vector analysis of adaptive optics images of the retina," *Opt. Lett.*, vol. **33**, no. 2, pp. 137–139, Jan. 2008.
- [19] Durán V., P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis, "Single-pixel polarimetric imaging," *Opt. Lett.*, vol. **37**, no. 5, pp. 824–826, Mar. 2012.
- [20] Sabatke D. S., M. R. Descour, E. L. Dereniak, W. C. Sweatt, S. A. Kemme, and G. S. Phipps, "Optimization of retardance for a complete Stokes polarimeter," *Opt. Lett.*, vol. **25**, no. 11, pp. 802–4, Jun. 2000.
- [21] Meadowlark Optics. *Liquid Crystal Polarimeter User Manual*. Available at www.meadowlark.com/store/PMI_Users_Manual_2.10.pdf, 2012.
- [22] Davis S., R. Uberna, and R. Herke, "Retardance sweep polarimeter and method," *US Pat. 6,744,509*, vol. **2**, no. 12, 2004.
- [23] Durán V., J. Lancis, E. Tajahuerce, and Z. Jaroszewicz, "Cell parameter determination of a twisted-nematic liquid crystal display by single-wavelength polarimetry," *J. Appl. Phys.*, vol. **97**, no. 4, p. 043101, 2005.
- [24] Brosseau C., *Fundamentals of Polarized Light: A Statistical Optics Approach*, 1st edn. John Wiley & Sons, Inc., 1998.
- [25] Boreman G. D., "Classification of imaging spectrometers for remote sensing applications," *Opt. Eng.*, vol. **44**, no. 1, p. 013602, Jan. 2005.
- [26] Stamatias G. N., M. Southall, and N. Kollias, "In vivo monitoring of cutaneous edema using spectral imaging in the visible and near infrared," *J. Invest. Dermatol.*, vol. **126**, no. 8, pp. 1753–60, Aug. 2006.
- [27] Hamilton S. J. and R. A. Lodder, "Hyperspectral imaging technology for pharmaceutical analysis," in *Proc. SPIE 4626, Biomedical Nanotechnology Architectures and Applications*, pp. 136–147, 2002.
- [28] Scholl J. F., E. K. Hege, M. Hart, D. O'Connell, and E. L. Dereniak, "Flash hyperspectral imaging of non-stellar astronomical objects," in *Proc. SPIE 7075, Mathematics of*

- Data/Image Pattern Recognition, Compression, and Encryption with Applications XI*, vol. **7075**, p. 70750H–70750H–12, 2008.
- [29] Dale L. M., A. Thewis, C. Boudry, I. Rotar, P. Dardenne, V. Baeten, and J. A. F. Pierna, “Hyperspectral imaging applications in agriculture and agro-food product quality and safety control: a review,” *Appl. Spectrosc. Rev.*, vol. **48**, no. 2, pp. 142–159, Mar. 2013.
- [30] Mehl P. M., Y.-R. Chen, M. S. Kim, and D. E. Chan, “Development of hyperspectral imaging technique for the detection of apple surface defects and contaminations,” *J. Food Eng.*, vol. **61**, no. 1, pp. 67–81, Jan. 2004.
- [31] Soldevila F., E. Irlles, V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis, “Single-pixel polarimetric imaging spectrometer by compressive sensing,” *Appl. Phys. B*, vol. **113**, no. 4, pp. 551–558, 2013.
- [32] Sampell J. B., “Digital micromirror device and its application to projection displays,” *J. Vac. Sci. Technol. B Microelectron. Nanom. Struct.*, vol. **12**, no. 6, p. 3242, Nov. 1994.
- [33] Vila-Francés J., J. Calpe-Maravilla, J. Muñoz-Mari, L. Gómez-Chova, J. Amorós-López, E. Ribes-Gómez, and V. Durán-Bosch, “Configurable-bandwidth imaging spectrometer based on an acousto-optic tunable filter,” *Rev. Sci. Instrum.*, vol. **77**, no. 7, p. 073108, 2006.
- [34] Zou X., J. Shi, L. Hao, J. Zhao, H. Mao, Z. Chen, *et al.*, “*In vivo* noninvasive detection of chlorophyll distribution in cucumber (*Cucumis sativus*) leaves by indices based on hyperspectral imaging,” *Anal. Chim. Acta*, vol. **706**, no. 1, pp. 105–112, Nov. 2011.
- [35] Mather J., “Spectral and XYZ color functions,” 2005. [Online]. Available at: www.mathworks.com/matlabcentral/fileexchange/7021-spectral-and-xyz-color-functions (accessed December 6, 2013).
- [36] Pratt W. K., *Digital Image Processing*, 4th edn. John Wiley & Sons, Inc., 2007.
- [37] Zuzak K. J., M. D. Schaeberle, E. N. Lewis, and I. W. Levin, “Visible reflectance hyperspectral imaging: characterization of a noninvasive, *in vivo* system for determining tissue perfusion,” *Anal. Chem.*, vol. **74**, no. 9, pp. 2021–2028, May 2002.
- [38] Coulson K. L., “Effects of reflection properties of natural surfaces in aerial reconnaissance,” *Appl. Opt.*, vol. **5**, no. 6, pp. 905–917, Jun. 1966.
- [39] Vanderbilt V. C., L. Grant, L. L. Biehl, and B. F. Robinson, “Specular, diffuse, and polarized light scattered by two wheat canopies,” *Appl. Opt.*, vol. **24**, no. 15, pp. 2408–2418, Aug. 1985.
- [40] Haugland S. M., E. Bahar, and A. H. Carrieri, “Identification of contaminant coatings over rough surfaces using polarized infrared scattering,” *Appl. Opt.*, vol. **31**, no. 19, pp. 3847–3852, Jul. 1992.
- [41] Pierangelo A., A. Benali, M.-R. Antonelli, T. Novikova, P. Validire, B. Gayet, and A. De Martino, “*Ex-vivo* characterization of human colon cancer by Mueller polarimetric imaging,” *Opt. Express*, vol. **19**, no. 2, pp. 1582–1593, Jan. 2011.
- [42] Zhao Y., L. Zhang, and Q. Pan, “Spectropolarimetric imaging for pathological analysis of skin,” *Appl. Opt.*, vol. **48**, no. 10, pp. D236–246, Apr. 2009.
- [43] Groner W., J. W. Winkelman, A. G. Harris, C. Ince, G. J. Bouma, K. Messmer, and R. G. Nadeau, “Orthogonal polarization spectral imaging: a new method for study of the microcirculation,” *Nat. Med.*, vol. **10**, no. 10, pp. 1209–1212, 1999.
- [44] Gupta N. and D. R. Suhre, “Acousto-optic tunable filter imaging spectrometer with full Stokes polarimetric capability,” *Appl. Opt.*, vol. **46**, no. 14, pp. 2632–2637, May 2007.

SCIENTIFIC REPORTS



OPEN

High-resolution adaptive imaging with a single photodiode

F. Soldevila^{1,2,*}, E. Salvador-Balaguer^{1,2,*}, P. Clemente^{1,2,3}, E. Tajahuerce^{1,2} & J. Lancis^{1,2}

Received: 16 June 2015

Accepted: 18 August 2015

Published: 18 September 2015

During the past few years, the emergence of spatial light modulators operating at the tens of kHz has enabled new imaging modalities based on single-pixel photodetectors. The nature of single-pixel imaging enforces a reciprocal relationship between frame rate and image size. Compressive imaging methods allow images to be reconstructed from a number of projections that is only a fraction of the number of pixels. In microscopy, single-pixel imaging is capable of producing images with a moderate size of 128×128 pixels at frame rates under one Hz. Recently, there has been considerable interest in the development of advanced techniques for high-resolution real-time operation in applications such as biological microscopy. Here, we introduce an adaptive compressive technique based on wavelet trees within this framework. In our adaptive approach, the resolution of the projecting patterns remains deliberately small, which is crucial to avoid the demanding memory requirements of compressive sensing algorithms. At pattern projection rates of 22.7 kHz, our technique would enable to obtain 128×128 pixel images at frame rates around 3 Hz. In our experiments, we have demonstrated a cost-effective solution employing a commercial projection display.

In the latter years, single-pixel imaging (SPI) has been established as a suitable tool in life sciences. One of the main characteristics of the technique is that it uses very simple sensors (bucket detectors such as photodiodes or photomultiplier tubes) and mathematical algorithms to recover an image¹. This reduction in complexity on the sensing device gives these systems the capability to work efficiently in conditions where light is scarce². Furthermore, single-pixel cameras have been demonstrated to obtain images at shallow depth overcoming the scattering problem^{3–5}. There are also several approaches that exploit the simplicity of the detectors in order to acquire multidimensional information, such as 3D, polarization and spectral images^{6–10}. However, this complexity reduction in the sensor entails an increase of computational time to recover an image when compared to conventional techniques. In the visible region of the spectrum, where pixelated sensors have acquired very high performances with low costs, this may not seem like a good trade-off. Nevertheless, in other regions of the spectra, such as infrared and terahertz, where pixelated sensors do not have such good specifications, this technique can provide huge benefits^{11–14}.

To recover an image, SPI needs to overlap a set of masks onto the scene under study and recover the total intensity of light transmitted or reflected by the scene. The size of this set depends on the desired resolution of the image. Even for low resolution images of 64×64 pixels, this requires a huge amount of projections (64^2). In spite of the fact that fast spatial light modulators (digital micromirror devices, or DMD) are usually used in these single-pixel camera architectures, this limits the speed of the acquisition process.

In order to solve this problem, compressive sensing (CS) techniques provide a method to recover the images with a number of measurements lower than the total number of pixels. This is possible because natural images tend to be sparse (i.e. only a small fraction of these projections have relevant information) in some basis of functions¹⁵. Despite lowering the total number of projections, the reconstruction algorithms require high computational power, which also limits the technique to low resolution images if

¹GROC.UJI, Departament de Física, Universitat Jaume I, E12071 Castelló, Spain. ²Institute of New Imaging Technologies (INIT), Universitat Jaume I, E12071 Castelló, Spain. ³Servei Central d'Instrumentació Científica (SCIC), Universitat Jaume I, E12071 Castelló, Spain. *These authors contributed equally to this work. Correspondence and requests for materials should be addressed to F.S. (email: fsoldevi@uji.es)




Figure 1. Spirogyra algae image (512×512 pixels), its first level and third level wavelet transforms, respectively. The bright pixels on the wavelet transform represent the edges of the scene. In the wavelet representation, a region of the scene is represented by a set of wavelet coefficients arranged in a tree structure, as shown in the right panel.

the user wants high speed acquisition and high-speed display in real-time. Some other techniques tackle this problem by using temporal sparsity^{16,17}.

Recently, a technique has been reported operating at sub-Hz rates with image sizes up to 128×128 pixels¹⁸. The key element of the technique is to project only the functions of the basis that have most of the information about the scene and recover the picture without using CS algorithms, thus speeding-up the display stage. Even though this approach solves the problem of the reconstruction time of the CS algorithms, to project only the important masks one must know beforehand which ones have the relevant information about the scene. Moreover, if the scene changes quickly, the a priori information of the relevant functions is rapidly lost. In order to achieve real time frame rates, it would be preferable to have a technique that does not use a priori information about the scene maintaining the benefits discussed above.

Here we propose an adaptive method for SPI that recovers images with a number of measurements lower than the number of pixels of the scene, with finer details than CS techniques (given the same number of measurements) and lower reconstruction times. This approach does not need to have a priori knowledge of the object and only collects information about the relevant parts of the scene in an adaptive way. It is based on smart sampling of the scene with a small set of masks. These masks are adaptively resized when the part of the scene needs to be recorded with higher resolution. The image is finally recovered by using fast wavelet transforms. Similar ideas have been proposed to improve the performance of ghost and dual photography imaging systems^{19–21}. This method is very suitable when the user wants to capture big size images for two main reasons. One, even though the size of the final image can be big (in our experiments, 2048×2048 pixels), the number of projected masks remains small due to the nature of the adaptive algorithm. And two, even if this number cannot be reduced due to object characteristics (i.e. objects with very complex spatial features or texture-like images), with this approach only a low resolution set of masks needs to be stored. This characteristic is very suitable from the computational point of view, as small deterministic matrices require low amounts of memory to be stored and can be easily used in fast mathematical operations, providing massive computational gains^{22,23}. For example, in the recovery stage of the aforementioned image classical SPI techniques would have to store 2048^2 different masks with 2048^2 elements each. Nevertheless, with the adaptive approach we are able to recover the image operating with a set of 128^2 masks with 128^2 elements each. With these characteristics in mind, we have been able to design a high-resolution fast-operation SPI system with an off-the-shelf DMD and a mid-range laptop. This kind of system can be easily coupled with available commercial microscopes to take advantage from the SPI benefits discussed above.

Results

Adaptive Compressive Imaging (ACI). To better understand the ACI algorithm one has to figure out how the 2D wavelet transform technique works. The process is depicted in Fig. 1. Given an $N \times N$ image, the wavelet transform consists of applying four bandpass filters to the image. As a result, four $N/2 \times N/2$ quadrants are obtained; a low resolution version of the image and three more quadrants with the information of the horizontal, vertical and diagonal edges. This process can be applied again to the low resolution version of the image, giving the tree-structured image shown in the right image of Fig. 1 where the upper left quadrant (the low resolution image) has been replaced by its wavelet transform. This procedure can be repeated up to \sqrt{N} times, when the pixel in the upper left corner contains the total energy of the scene and the rest of the image has the information about the edges of the scene. As it can be seen, the number of pixels containing information of sharp edges is scarce thus few coefficients are enough to get an image similar to the original one. Wavelet compression algorithms choose a number of iterations (levels) and only store the coefficients with values higher than a predetermined threshold,




Figure 2. Adaptive Compressive Imaging operation scheme. Object: 256×256 *Fasciola hepatica* cells image captured with a commercial microscope. Stage 1: Coarse picture and its level one wavelet transform inspection. As the fourth quadrant has no relevant information, it is discarded. Stage 2: Higher resolution images of the non-discarded zones with their level one wavelet transforms inspection. In this stage, six regions are discarded. Stage 3: Highest resolution images of the non-discarded zones and their level one wavelet transforms. As this is the last stage, no more zones will be discarded so there is no wavelet inspection process. Using all the level one wavelet transforms, the algorithm builds the level three wavelet transform. By doing its inverse wavelet transform, the reconstruction of the scene is acquired. In this example, the total number of measurements to recover the scene was 62% of the 256^2 measurements established by the Nyquist criterion.

reducing in this way the size of the digital file without a significant loss in quality. As the operations needed to calculate this transform are linear, they are an excellent option for fast algorithms because they require low computational power and memory usage.

The idea behind the ACI technique is to reduce the number of masks needed to reconstruct an image, using high resolution masks only on the high resolution regions of the image. To that end, we first sample the scene with a low resolution set of masks and only continue sampling with higher resolution masks the parts of the image with regions of interest (i.e. the regions with high density of sharp edges). If the algorithm detects a zone in the scene with no borders, it does not project patterns again on this region, as more resolution is not needed. This process is repeated until we arrive at the final desired resolution. As a demonstration of the technique, in Fig. 2 we show a simulation done with a real microscopy image.

In this case, the goal is to recover an image of 256×256 pixels of a group of cells. The set of masks chosen is the 2D Hadamard basis of 64×64 pixels. In the first stage, the algorithm acquires a coarse picture of the full scene with our set of masks (resizing the 64×64 Hadamard patterns to 256×256 pixels). Once the image is recovered, it calculates its level one wavelet transform and searches for the quadrants with higher density of borders. If one of the quadrants has a number of borders below a predefined threshold, it will be discarded on the next stages, thus reducing the total number of masks projected. Once this step is complete, the second stage of the algorithm starts. Now, the 64×64 patterns are resized to 128×128 pixels, therefore occupying a quarter of the original scene. If none of the quadrants has been discarded in the previous stage, here the algorithm takes four more pictures, recovering then the scene with finer details. If one or more of the quadrants were discarded on the previous stage, the algorithm does not project the set of patterns in the discarded quadrants. In this stage, the algorithm repeats the level one wavelet transformations to each image in order to search again for new zones with no borders. As the algorithm goes on, the search zones get smaller and smaller, and the following sets of smaller masks are only projected on the high spatial resolution zones of the scene. In the last stage, masks are




Figure 3. Fasciola hepatica 4 MP image (left panel) and its ACI reconstruction (right panel). The ACI reconstruction is acquired with roughly a 25% of the 2048^2 measurements established by the Nyquist ratio.

projected onto the regions with finer details of the scene. Once all the regions have been measured with the required resolutions, all the wavelet transforms are used to build the third level wavelet transform shown in Fig. 2. The final 256×256 image is recovered via inverse wavelet transform. In contrast with traditional CS techniques, which usually require off-line reconstruction, this step is not computationally consuming. As it will be further discussed in the text, when using ACI this process can be done on the fly with low end computers.

The ACI approach has several remarkable benefits when compared with the traditional SPI-CS approach. These can be grouped into three categories: image size, resolution and temporal benefits. In order to prove those benefits, throughout this paper, we will compare our technique with the GPSR-Basic algorithm by Figueredo *et al.*²⁴. First, we will discuss the image size benefits. As stated before, to recover a $N \times N$ picture, SPI needs to project $M = N^2$ masks. When using CS techniques, those M projections are reduced (typical compression ratios tend to be between 10% and 40% without a significant quality loss). However, this reduction entails the use of convex optimization algorithms to recover the image. The memory and time requirements of the algorithms increase with the size of the image and the number of measurements made. Even if speed is not a crucial requirement for some applications, the memory requirements limit the maximum size of the images recovered. In our experiments, carried on a computer with 24GB of RAM and a Intel Xeon Processor X5690 at 3.47 GHz, the maximum image size that can be reconstructed due to memory limitations with CS algorithms is 256×256 pixels, with compression ratios around 50%. When using the ACI algorithm, only a small set of masks of low resolution and the measurements vector need to be stored. In practice, we did simulations of images with sizes up to 4 Megapixel (MP), even though this is not the limit of our equipment. As can be seen in Fig. 3, the regions of interest of the scene are recovered with perfect resemblance, and the zones with no information have lower detail.

When dealing with higher resolutions like the one shown above, the reconstruction time starts to be a crucial factor to be reckoned. A general approach made by researchers is to sacrifice some quality in their reconstructions in order to achieve higher frame rates. This can be made by applying high compression rates or by novel approaches like the one proposed by Radwell *et al.*¹⁸. This procedures either aren't fast enough to achieve high resolution real time imaging or need a priori information about the scene to speed-up the reconstruction process. In SPI systems, the image acquisition time, t_a , depends on the number of projected masks, M , which is determined by the size of the image. Defining the projection rate of the SLM as R_{SLM} , the image acquisition time is given by $t_a = M \cdot (1/R_{SLM} + t_{int}) + t_p$, where t_p is the post-processing time to recover the picture from the measurements made, M , and t_{int} is the integration time of the bucket detector. Ideally, both M and t_p should be as low as possible. As single-pixel detectors work at higher frequencies than SLM's, t_{int} is negligible in all the scenarios considered here where lighting conditions are not extreme. Whereas traditional single pixel imaging requires M to be equal to the number of pixels of the image, N^2 , and has negligible t_p , CS techniques reduce M but increase t_p . Furthermore, CS techniques need to solve a convex optimization problem to recover an image, which requires high amounts of memory. Adaptive imaging techniques are known to reduce M while keeping negligible t_p ^{19,20}. Nevertheless, due to the nature of ACI, increasing the size of the scene does not necessarily imply increasing memory requirements. If there is a memory limitation, the number of stages to reconstruct a scene will be increased, and consequently the size of the masks reduced.

In Fig. 4 we show two comparisons between CS and ACI algorithms. We have verified the PSNR and reconstruction time to behave similar for several biological images. To carry out the simulation we have used three different biological test images (shown in Fig. 5), with a size of 128×128 pixels. The ACI algorithm number of stages was set to three, thus using the 32×32 Hadamard masks. The first graph




Figure 4. Full frame quality and time comparison between CS and ACI techniques. The full scene resemblance is measured with the PSNR comparison in the top graph. Given the same amount of measurements, the quality is similar in both methods. However, the time comparison shows that the ACI technique outspeeds the traditional CS technique.




Figure 5. ROI quality comparison between CS and ACI techniques. Three biological samples are selected for the analysis. In each sample, one ROI is studied. Due to the adaptive nature of the ACI approach, its PSNR curves present a steplike behaviour, acquiring higher resemblance than the traditional CS technique. In order to ease the visualization, the CS curve is the average of the three ROI's, as the results were almost equal in all the images.




Figure 6. Experimental setup used for ACI reconstructions. The digital projector sends a predefined set of masks to different parts of the object. The light reflected from the object is measured with a bucket detector. The signal is digitalized and used to recover an image of the scene.

shows the average of the PSNR versus number of measurements for ACI and CS algorithms when the reconstructions are compared with the original images. In the second graph we show a time comparison between both methods. The time includes not only the CS and ACI algorithm computational time but also the masks projecting time. Even though the quality is similar, it is clear by watching at the lower graph, that the ACI technique has a great benefit in reconstruction times. As can be deduced from the results, high quality pictures can be achieved with sub-Nyquist measurement rates around 50%. With those measurement ratios and state-of-the-art DMD's, images of 128×128 pixels can be acquired at frame rates around 3 Hz (~ 8000 measurements at 22.7 kHz). Similar tests were made with higher resolutions with results even more favorable to the ACI technique. Images with a resolution of 256×256 pixels are reconstructed in less than a minute with the ACI algorithm in comparison with several days with the GPCR routine. Higher resolution images could not be compared as our equipment is not able to reconstruct images with resolutions above 256×256 pixels using CS.

Another huge benefit of the technique is that it recovers the regions of interest at high resolution using sub-Nyquist measurement rates. In Fig. 5 we show an example with three biological samples. For each sample, we chose a region of interest and we compare the PSNR of that region when using both techniques. As CS projects masks covering the entire scene, the quality of the whole image gradually improves with the number of projected masks, independently of the region of interest chosen. However, this does not happen with the ACI algorithm because masks are sent to different regions of the scene. In this case, the ACI curves have a steplike behaviour, where each step corresponds to a stage of the algorithm. If there are few regions of interest in the scene, steps are concentrated in the initial measurements (see green curve), while if the scene is plenty of sharp edges, the highest quality of the region is achieved later (see orange curve). In microscopy setups, where the samples usually lay onto specific regions of a slide, this characteristic can be used to recover specimens with very low number of measurements or to locate regions of interest in a fast way. Once those regions are located, the amplification of the system can be changed so the sample fills the full field of view of the system.

Experimental results. In order to test those ideas, we conducted a proof of concept experiment with a projector and a mid-range laptop. The experimental setup is shown in Fig. 6. It consists of a digital light projector, a photodiode, an analog-to-digital converter and a computer. The DLP sends the set of masks onto different regions of the object, resizing them when needed. Each set of resized masks is pre-computed, and custom software written in Labview chooses the suitable one for each stage of the algorithm. As the number of pixels of each mask remains the same in all stages (the only change is the pixel size), the reconstruction algorithm computational charge is alleviated. By means of an optical collecting system, light reflected by the object is measured with the photodiode. Being the quantum efficiency of




Figure 7. ACI experimental results. In the upper row, we show a 256×256 pixels LEGO® Ned Flanders picture (left) and its reconstruction via the ACI technique (right). The ACI reconstruction was acquired by using 88% of measurements of the total number of pixels. In the lower row, we show a 512×512 pixels USAF1951 test (left) and its ACI reconstruction (right). In this case, the reconstruction was acquired with a 55% of measurements of the total number of pixels. The only post-processing made to the ACI images was white balance correction.

photodetectors higher than CCD/CMOS sensors and given that more photons reach the detector at each measurement, signals acquired suffer less distortion from dark and read-out noise¹. The analog-to-digital converter digitalizes the signal, and the computer reconstructs the images with custom code written in Matlab. The experimental process of projecting the patterns and measuring the electrical signal is controlled by custom software written in Labview.

In Fig. 7, two experimental reconstructions are shown. The first scene reconstructed is a small LEGO® object. We start the ACI algorithm with 64×64 Hadamard patterns and the number of stages is set to 3. Then, the final resolution achieved is 256×256 pixels. Unlike standard CS techniques the time needed to reconstruct the scene only depends on the SLM refresh rate and not on the compressive strategy. For this particular reconstruction we have sent 88% of the total 256^2 measurements established by the Nyquist criterion. In our experiments, we use a DLP LightCrafter 4500 from Texas Instruments. Even though the maximum refresh rate of this device is 4225 Hz, when the number of patterns to be projected gets higher, memory limitations arise. Then, the patterns have to be sent using the video input of the device instead of being preloaded on the internal memory. The speed of this video input is limited to 120 Hz. By encoding 24 different binary patterns in each video frame as a 24-bit image, the maximum speed acquired is 2880 Hz^{25} . Bearing that in mind, the acquisition time with our equipment was 20.02 seconds. If state of the art SLM's are used, with refresh rates around 22.7 kHz and high capacity internal memory, these limitations can be avoided, and reconstruction times of 2.54 seconds can be attained with this method.

As stated before, ACI stands out when capturing big resolution images. Due to the DLP specifications, the biggest square masks that can be projected have a size of 512×512 pixels. In the second row of Fig. 7, we show a reconstruction of an USAF test with that resolution. In order to achieve this resolution, four stages of the ACI algorithm were used. As the picture gets bigger, the number of discarded regions tends to get higher, so greater compression ratios are achieved while maintaining good resemblance with the scene. In particular, for this second example we only used 55% of measurements of the 512^2 established

by the Nyquist criterion. With our setup, we acquired the image in 50.06 seconds. By using high-end DMD's, acquisition times of 6.35 seconds could be achieved.

Discussion

We have designed an ACI algorithm that allows recovering high resolution images at low time costs by using SPI. Compared to traditional CS approaches, we acquire images with equivalent quality in much lower times^{1,3}. In fact, with the same number of measurements, CS needs post-processing to recover the image while ACI can do it live. Novel approaches like multi-diode design cameras or adaptive ghost imaging have also tackled the speed-resolution limitation with success^{20,26}. However, our approach maintains the use of a single photodetector, which improves the SNR of the measurements when working on low-light level scenarios, such as biological environments. Furthermore, by using deterministic matrices as the basis of our masks, the method is better suited for fast mathematical operations or even using CS in each stage of the acquisition process²². We have also computationally demonstrated the effectiveness of the technique to perfectly recover the regions of interest of an image at sub-Nyquist measurement rates. For the scenes investigated here, containing relatively few regions of interest, ACI has been shown to provide higher quality reconstructions in less time than the GPSR-Basic CS approach. Compared to other single-pixel techniques, we do not need a priori knowledge of the scene to achieve this speed¹⁸.

Huge technological efforts are focused on increasing resolution of optical devices. Up to now, SPI has failed to provide high resolution images due to time restrictions. The only limiting factor of ACI technique lies in the SLM refresh rate, enabling us to present high resolution experimental images with dimensions comparable to the SLM number of pixels, without needing stitching techniques and using a single photodiode as a detector.

It has to be also noted that ACI is a very flexible technique. For example, we can use traditional compressive techniques in each ACI stage to reduce even more the number of measurements needed. ACI can also be used to improve the performance of single-pixel cameras working in different regions of the visible, infrared and terahertz spectrum^{3,4,11}. Future work in the ACI technique should be directed towards improving the adaptive scheme. This will involve searching edges in a more intelligent way. Instead of inspecting quadrants regularly placed in the scene, we could freely situate them on high density border zones to gather the information more efficiently. This will improve the quality of the recovered images at even lower measurement rates.

Methods

The DLP used in the experiments is a DLP LightCrafter 4500 from Texas Instruments. It contains a DMD and three coloured light sources (red, green and blue). A built-in optical system is used to project the patterns onto the scene. The photodetector used is PDA36A-EC from Thorlabs, and the electrical signal is digitalized with NI USB-6001 DAQ. All the experimental results were acquired with a Lenovo ThinkPad E531 laptop with 12GB of RAM and an Intel Core i7 2.20 GHz processor. The biological images used in the simulations correspond to different samples from two slide sets from Carolina (#292148A and #293708).

ACI algorithm. Here we attach the pseudo-code of the ACI algorithm used in the experimental setup.

Algorithm 1 ACI
Input: $L \leftarrow$ number of stages
$R \leftarrow$ final image resolution
$Q \leftarrow$ Hadamard patterns of resolution: $\frac{R}{2^{L-1}}$
Output: $\text{finalImg} \leftarrow R \times R$ image of the scene
for $i = L$ to 1 do
Divide the scene into 4^{L-i} quadrants.
for $j = 1$ to 4^{L-i} do
if $i \neq L$ and $\text{quadrantRelevance}(j) < \text{threshold}$
$\text{childrenQuadrantRelevance}(j) \leftarrow$ Set to 0 the quadrantRelevance of all its children
$\text{quadrantImg}(j) \leftarrow$ Use the information of its father to recover the image of quadrant j .
else
$\text{quadrantImg}(j) \leftarrow$ Use patterns Q to sample the quadrant j of the scene and recover its image.
end if
$\text{quadrantWavelet}(j) \leftarrow$ perform one-level wavelet transform of $\text{quadrantImg}(j)$
Continued

childrenQuadrantRelevance(j) ← borders information extracted from quadrantWavelet(j).
end for
stageWavelet(i) ← compose the wavelet transform of stage i.
end for
finalWavelet ← use the information of all stageWavelet to recover the L-wavelet transform information of the scene.
finalImg ← Perform a L-level inverse wavelet transform to recover the $R \times R$ image of the scene.

References

- Duarte, M. F. *et al.* Single-Pixel Imaging via Compressive Sampling. *IEEE Signal Process. Mag.* **25**, 83–91 (2008).
- Howland, G.A., Lum, D.J., Ware, M. R. & Howell, J. C. Photon counting compressive depth mapping. *Opt. Express* **21**, 23822–23837 (2013).
- Tajahuerce, E. *et al.* Image transmission through dynamic scattering media by single-pixel photodetection. *Opt. Express* **22**, 16945 (2014).
- Durán, V. *et al.* Compressive imaging in scattering media. *Opt. Express* **23**, 14424–14433 (2015).
- Ntziachristos, V. Going deeper than microscopy: the optical imaging frontier in biology. *Nat. Methods* **7**, 603–614 (2010).
- Sun, B. *et al.* 3D computational imaging with single-pixel detectors. *Science* **340**, 844–847 (2013).
- Durán, V., Clemente, P., Fernández-Alonso, M., Tajahuerce, E. & Lancis, J. Single-pixel polarimetric imaging. *Opt. Lett.* **37**, 824–826 (2012).
- Soldevila, F. *et al.* Single-pixel polarimetric imaging spectrometer by compressive sensing. *Appl. Phys. B* **113**, 551–558 (2013).
- Welsh, S. S., Edgar, M. P., Bowman, R., Sun, B. & Padgett, M. J. Near video-rate linear Stokes imaging with single-pixel detectors. *J. Opt.* **17**, 025705 (2015).
- Edgar, M. P. *et al.* Simultaneous real-time visible and infrared video with single-pixel detectors. *Sci. Rep.* **5**, 10669 (2015).
- Watts, C. M. *et al.* Terahertz compressive imaging with metamaterial spatial light modulators. *Nat. Photonics* **8**, 605–609 (2014).
- Chen, H., Xi, N., Song, B. & Lai, K. Single pixel infrared camera using a carbon nanotube photodetector. *Sensors, 2011 IEEE* (2011).
- Shrekenhamer, D., Watts, C. M. & Padilla, W. J. Terahertz single pixel imaging with an optically controlled dynamic spatial light modulator. *Opt. Express* **21**, 12507 (2013).
- Chan, W. L. *et al.* A single-pixel terahertz imaging system based on compressed sensing. *Appl. Phys. Lett.* **93**, 121105 (2008).
- Romberg, J. Imaging via Compressive Sampling. *IEEE Signal Process. Mag.* **25**, 14–20 (2008).
- Noor, I. & Jacobs, E. L. Adaptive compressive sensing algorithm for video acquisition using a single-pixel camera. *J. Electron. Imaging* **22**, 021013 (2013).
- Sankaranarayanan, A. C., Studer, C. & Baraniuk, R. G. CS-MUVI: Video compressive sensing for spatial-multiplexing cameras. in *2012 IEEE International Conference on Computational Photography (ICCP)* 1–10 (IEEE, 2012).
- Radwell, N. *et al.* Single-pixel infrared and visible microscope. *Optica* **1**, 285–289 (2014).
- Aßmann, M. & Bayer, M. Compressive adaptive computational ghost imaging. *Sci. Rep.* **3**, 1545 (2013).
- Yu, W. K. *et al.* Adaptive compressive ghost imaging based on wavelet trees and sparse representation. *Opt. Express* **22**, 7133–7144 (2014).
- Sen, P. *et al.* Dual photography. *ACM Trans. Graph.* **24**, 745 (2005).
- Blanchard, J. D. Toward deterministic compressed sensing. *Proc. Natl. Acad. Sci. USA* **110**, 1146–1147 (2013).
- Monajemi, H., Jafarpour, S., Gavish, M. & Donoho, D. L. Deterministic matrices matching the compressed sensing phase transitions of Gaussian random matrices. *Proc. Natl. Acad. Sci. USA* **110**, 1181–1186 (2013).
- Figueiredo, M. A. T., Nowak, R. D. & Wright, S. J. Gradient Projection for Sparse Reconstruction: Application to Compressed Sensing and Other Inverse Problems. *IEEE J. Sel. Top. Signal Process* **1**, 586–597 (2007).
- Thorlabs, Inc. DLP[®] LightCrafter 4500TM Evaluation Module User's Guide. at <http://www.ti.com/lit/ug/dlpu011d/dlpu011d.pdf> (2014) Date of access: 03/08/2015.
- Herman, M. A., Tidman, J., Hewitt, D., Weston, T. & McMackin, L. A higher-speed compressive sensing camera through multi-diode design. in *Proceedings of the SPIE* (ed. Ahmad, E.) **8717**, 871706 (2013).

Acknowledgements

This work was supported in part from MINECO (grant FIS2013-40666-P), Generalitat Valenciana (grants PROMETEO2012-021 and ISIC 2012/013), and Universitat Jaume I (P1-1B2012-55).

Author Contributions

F.S. and E.S. conceived, conducted the experiments and wrote the manuscript; F.S., E.S., P.C., E.T. and J.L. discussed the results and revised the manuscript.

Additional Information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Soldevila, F. *et al.* High-resolution adaptive imaging with a single photodiode. *Sci. Rep.* **5**, 14300; doi: 10.1038/srep14300 (2015).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

SCIENTIFIC REPORTS

OPEN

Computational imaging with a balanced detector

F. Soldevila¹, P. Clemente², E. Tajahuerce¹, N. Uribe-Patarroyo³, P. Andrés⁴ & J. Lancis¹

Received: 02 March 2016

Accepted: 15 June 2016

Published: 29 June 2016

Single-pixel cameras allow to obtain images in a wide range of challenging scenarios, including broad regions of the electromagnetic spectrum and through scattering media. However, there still exist several drawbacks that single-pixel architectures must address, such as acquisition speed and imaging in the presence of ambient light. In this work we introduce balanced detection in combination with simultaneous complementary illumination in a single-pixel camera. This approach enables to acquire information even when the power of the parasite signal is higher than the signal itself. Furthermore, this novel detection scheme increases both the frame rate and the signal-to-noise ratio of the system. By means of a fast digital micromirror device together with a low numerical aperture collecting system, we are able to produce a live-feed video with a resolution of 64×64 pixels at 5 Hz. With advanced undersampling techniques, such as compressive sensing, we can acquire information at rates of 25 Hz. By using this strategy, we foresee real-time biological imaging with large area detectors in conditions where array sensors are unable to operate properly, such as infrared imaging and dealing with objects embedded in turbid media.

Over the past decade, single-pixel imaging (SPI) has demonstrated its viability in scenarios where traditional imaging techniques struggle to provide images with acceptable quality in practicable times and reasonable costs. Examples of this can be found in exotic regions of the light spectrum, such as infrared¹, terahertz² or X-rays³. Moreover, single-pixel setups have also been proposed in challenging frameworks such as photon-counting regimes⁴, imaging in scattering media^{5,6}, multidimensional imaging^{7–11}, two-photon microscopy¹², photoacoustic imaging^{13,14}, and spatial entanglement imaging¹⁵.

In SPI, the resolving power of the system is shifted from the sensor to a set of microstructured spatial masks that are codified onto a programmable spatial light modulator (SLM). The masks are optically projected onto the sample and the whole intensity is collected onto a bucket (single-pixel) sensor. The photodetector provides an electrical signal, proportional to the total amount of light that leaves the object. Measurements are sequentially made by changing the spatial mask. If many different masks are used, their shapes and the electrical signal can be combined to retrieve the sample¹⁶.

However, SPI still has several limitations inherent to the technique. If unmodulated light arrives to the object and it is collected, the signal provided by the detector is corrupted and the recovered image gets affected. Depending on the quantity of ambient light, recovery can be unattainable. Furthermore, the use of SLMs to measure the projections of the scene onto a set of spatial masks places an upper bound to the acquisition speed of the devices. The nature of SPI enforces a reciprocal relationship between the frame rate and the image size as the time required to capture an image scales with the number of pixels in the image. As a matter of fact, SPI usually relies on the use of fast SLMs such digital micromirror devices (DMDs) to codify the projecting masks. DMDs permit highly flexible codification of binary masks at frame rates above 20 kHz. For images about 128×128 pixels, typical acquisition times tend to be around one frame per second¹.

Two different approaches can be employed to overcome the above issue. On the one hand, given some reasonable assumptions about the sparsity of the signal, compressive sensing (CS) dramatically reduces the number of measurements needed well below the number of pixels of the sample^{17–19}. What is remarkable here is that with the only assumption that the signal is sparse, it is possible to avoid the measurement of the full-length signal, saving measurement time. More recently, adaptive sensing has been introduced as a way to circumvent the computational complexity in convex optimization or greedy algorithms used in CS^{1,20,21}. The idea is to make some

¹GROC-UJI, Institute of New Imaging Technologies (INIT), Universitat Jaume I, E12071 Castelló, Spain. ²Servei Central d'Instrumentació Científica (SCIC), Universitat Jaume I, E12071 Castelló, Spain. ³Wellman Center for Photomedicine, Harvard Medical School and Massachusetts General Hospital, 40 Blossom Street, Boston, Massachusetts 02114, USA. ⁴Facultat de Física, Universitat de València, E46100 Burjassot, Spain. Correspondence and requests for materials should be addressed to F.S. (email: fsoldevi@uji.es)




Figure 1. Single-pixel balanced detection camera. The object illuminated with incoherent light is imaged onto the surface of the DMD via a lens. Light reflected in both directions by the DMD is gathered by two collecting lenses and focused onto the entrance of two optical fibers, which are connected to the balanced photodetector. The signal is digitized and stored in the computer, which controls all the process. An external light source is used to introduce ambient light to the system in a controlled way with the aid of a beam splitter. ALS: ambient light source (halogen lamp), WLS: white-light source, OBJ: object, BS: beam splitter, L_1 : lens, CL: collecting lens, DMD: digital micromirror device, BPD: balanced photodetector, ADC: analog-to-digital converter, PC: computer.

preliminary low-resolution measurements, and take advantage of the obtained information to avoid irrelevant acquisitions in further measurements. In this way, high resolution images can be obtained in a fast and computationally efficient way.

When dealing with light coming from the surroundings, measuring directly in the Fourier space and computationally filtering the frequencies of the ambient light has been proven to increase the quality of the images even when the ambient light has more power than the signal itself²². Moreover, the use of complementary (or differential) measurements, which has been reported in several works^{1,9,23–25}, also increases the signal-to-noise ratio of the system, helping the devices in this type of unfavorable environmental conditions.

Here we present a novel approach, using complementary measurements and a single balanced detector. By using balanced detection, we improve the frame rate of the complementary measurement architectures by a factor of two, and we make use of all the photons arriving from the object. Furthermore, the use of a balanced detector provides environmental light immunity to the method. The key to success is the use of both reflecting arms provided by the DMD simultaneously, via the two sensors in the balanced detector. In addition to this, balanced detection also improves the signal digitization process. When we consider a single detector the detected light will be, on average, half the energy of the object with oscillations caused by the overlapping of each pattern with the object intensity distribution. The electrical signal as a function of time can be visualized as a DC term, which depends on the total energy of the object, and the AC term caused by the varying overlap of the patterns with the object. There are situations where using this approach causes loss of information because the sensitivity and dynamic range of both detectors and digitization systems are limited⁶. When using balanced detection, the DC term can be eliminated, and the full dynamic range of the detection system can be used to measure the small oscillations that are of interest to recover the image. This is another example of SPI taking advantage of the use of dedicated sensors that are difficult to implement in an array^{8,9,26}.

Results

Experimental setup. Our single-pixel camera can be seen in Fig. 1. The system resembles a conventional camera but the object is imaged over an array of micromirrors as opposed to the traditional array of photosensors. The object is illuminated by a white-light lamp (Thorlabs HPLS200). Then, the image is created onto the DMD surface (DLP Discovery 4100 V-7000 from ViALUX). A DMD consists of an array of electronically controlled micromirrors that can rotate about a hinge. The angular position of a specific micromirror admits two possible states ($+12^\circ$ and -12° respect to a common direction which will be referred ON and OFF respectively). In this way, the light can be reflected in two different directions depending on the signal applied to the mirror. If one looks at the DMD from one reflection direction, only the mirrors in the ON state will reflect light, and the mirrors in the OFF state will appear dark. In the other reflection direction there will be a complementary illumination pattern, as previous OFF mirrors will reflect light, whereas previous ON mirrors will not. A sequence of sampling patterns is codified onto the DMD and the irradiance striking the large-area single-pixel photodetector is stored. In one reflection direction, we have the superposition of one pattern with the object. In the other direction, we have the superposition of the complementary pattern and the object. Light is collected by two identical collimating lenses (Thorlabs F810SMA-635), and enters two optical fibers (Thorlabs FT200EMT). The output of the balanced photodetector (Thorlabs PDB210A/M) corresponds to the subtraction of the two signals in the analog domain, amplified before the analog-to-digital conversion made by our digital acquisition system (National Instruments NI6251). In order to introduce ambient light in a controlled way, we used a commercial halogen

lamp. Custom software written in LabVIEW controls both the generation of the patterns in the DMD and the digitization process of the analog-to-digital converter (ADC).

Concerning the measurement patterns, various matrices can be employed. For instance, raster-scan style masks stem from the well-known raster-scan technique in which spatial pixels are measured sequentially. Random matrices can also be used in which each mask has a random distribution of binary values. On the other hand, Hadamard matrices provide a convenient codification framework because binary non-negative elements can easily be displayed onto the DMD (see Supplementary Information). Structured matrices avoid the need to store entire matrices as the entries can be computed on the fly and permit the use of recovery algorithms with a lower computational penalty than random matrices. Note that operation at video rates is of paramount significance in potential applications of SPI like optical microscopy.

One important feature of the setup is the light collecting system. In our setup, light reflected by the DMD is collected using two identical optical fiber collimating systems. Those systems are coupled to two identical optical fibers. As both the optical fibers and the collimating systems have low numerical aperture (0.39 and 0.25 NA, respectively), light coming from directions other than the object is partially removed from the system before arriving to the detector. Even though this is a good feature and improves the quality of the presented images, it is not possible to reduce the NA of the system arbitrarily. A lower NA will decrease the field of view of the system, and also the total quantity of light at the detector, decreasing the overall quality of the reconstructions.

The other key element of our setup is the balanced detector. The use of balanced detection in optical experiments dates from the late 60's²⁷. The simplest way of doing balanced detection consists of connecting two photoreceptors in a way that their photocurrents cancel. If one is able to equalize the optical power arriving at each sensor, the effective photocurrent generated by the pair will be zero. Changes in the optical signal between both photoreceptors will unbalance the detector, generating a signal at the detector output. In this way, even if the optical power arriving at the detector is high, the output signal will only depend on the intensity of the fluctuations. This is of paramount usefulness when the information that one wants to measure relies in the light fluctuations, and not in the absolute value of the signals. Several optical techniques that require high signal-to-noise ratio have been using the principles of balanced detection for decades now. Examples range from microstructure metrology²⁸, light scattering microscopy²⁹, high sensitive measurements using noisy sources³⁰, and what may be the most well-known optical technique using balanced detection, optical coherence tomography³¹.

In our technique, the information about the object is encoded in the difference between two signals (see Methods). In one reflection direction of the DMD, we have the result of overlapping the object and one set of patterns. On the other reflection arm, we have the result of overlapping the object and the complementary set of patterns (switching white parts by black parts and vice versa). The difference between these two signals provides the mathematical projection of the object in our basis of functions: measuring this difference using balanced detection is straightforward. Using this idea, not only we capture all the photons reflected by the object and measure both signals at the same time (reducing acquisition time), but we also use the full dynamic range of our ADC. In a traditional SPI scheme (whether using complementary illumination or not), electrical signals provided by detectors will always be positive. ADC systems use a determined number of bits in order to digitize the signal. Once the voltage range is set, the number of quantization bits determines the possible digitized voltage levels and their precision. As voltage signals have both positive and negative values, using any ADC to measure an always positive electrical signal implies losing half the number of quantization values. However, in our setup we work directly with the subtraction of two signals, which will have both positive and negative values depending on the patterns generated in the DMD. The use of both reflection arms of the DMD improves the signal strength because we utilize the full flux of photons coming from the scene, in contrast with the classical single-pixel approach, which loses, on average, half of the photons.

Lastly, the adoption of balanced detection also increases the ambient light immunity of the optical system. Since both sensors of the balanced detector work at the same time, temporal fluctuations of the signal are subtracted automatically. This is very relevant in experimental conditions where light coming from unstable sources can corrupt the signal, making unfeasible the recovery of the object. This is illustrated in Fig. 2. We show several snapshots of the electrical signal when using balanced detection or single detection and switching ON or OFF an ambient halogen light source with intensity fluctuations at 100 Hz. In all four images, we see the temporal square signal generated by the DMD when switching between two different Walsh-Hadamard patterns at a low rate (around 10 Hz). When we turn on the halogen lamp, an undesirable ripple appears in the signal, with a frequency of 100 Hz. When using balanced detection, as the oscillating signal introduced by the halogen lamp is the same in both sensing arms, it does not contribute to the balanced signal. However, in the single detection case, this ripple persists. Using the corrupted signal to recover an object via SPI techniques would provide low quality images, as we will see in the following section.

Experimental results. As a proof-of-concept, we used the experimental setup shown in Fig. 1 to record a set of images with different illumination conditions. The results are shown in Fig. 3. For a given fixed power of the halogen lamp, we changed the intensity of the light illuminating the sample object (a small part of an USAF resolution test chart with 3.6 lines/mm) and acquired an image both using balanced detection (Fig. 3(a)) and only with one sensor (Fig. 3(b)). In both cases, we also took images with the halogen lamp turned off. The images have a size of 64×64 pixels, and each pixel was arranged by grouping 4 DMD mirrors, with $13.68 \mu\text{m}$ width each. In order to acquire the images, the full set of 4096 (64^2) Walsh-Hadamard patterns was generated in the DMD. To ease visualization, all the images were normalized between 0 and 255 $\left(I_{norm} = \frac{I - \min(I)}{\max(I) - \min(I)} \cdot 255\right)$.

It is clear, even when the halogen lamp is much dimmer than the object source, that the non-balanced image has much lower quality. One can see a checkerboard-like artifact at the top of the images. This is caused by the temporal oscillation of the signal at a frequency of 100 Hz. As the sampling rate of the ADC is known, one could




Figure 2. Electrical signals in presence of ambient light. Four snapshots from the oscilloscope with different illumination and sensing conditions. In order to acquire the signals, two different Walsh-Hadamard patterns were generated periodically onto the DMD surface, thus giving a low frequency square wave. The first row of pictures corresponds to balanced detection with the halogen lamp ON (a) and OFF (b). It is clear that no fast variations appear in any case. The second row of pictures corresponds to sensing with only one of the detectors when the lamp is ON (c) or OFF (d). It can be seen that a ripple at a frequency of 100 Hz appears in (c). If one needs to measure the difference between the two intensity levels, this oscillation can corrupt the measurement, and thus decrease the quality of the recovered image.

digitally filter the 100 Hz component of the signal in the Fourier domain, go back to the time domain and recover the object suppressing this checkerboard pattern. However, this requires *a priori* knowledge of the power spectrum of the ambient and parasitic light sources. Balanced detection makes this post-processing of the electrical signal unnecessary. This is advantageous as post-processing procedures are time consuming, decreasing the total acquisition speed of the system.

Readers may observe a reduction in the quality of the images as the WLS power decreases while that of the halogen lamp remains constant. Here we point out two different effects. On the one hand, we have the effect of taking an image using a small quantity of light. On the other hand, we identify the effect of varying the power ratio between the two sources. For the first effect, we can look at both the second and fourth rows of Fig. 3. In this case, the halogen lamp is turned off and we take images decreasing the power of the lamp illuminating the object. We can see that as illumination decreases, noise starts to appear covering the image. This is due to the signal power approaching the level of the technical noise of the setup, which ultimately limits the working range of illumination powers in which the camera can acquire images. In order to work at different illumination levels or other regions of the electromagnetic spectrum, sensors can be exchanged with high sensitive photodiodes (or avalanche photodiodes) or even detectors with different spectral sensitivities (outside the visible region). The second effect depends on the ratio between the power of the two light sources. As we stated before, the halogen lamp produces a checkerboard-like pattern on top of the image when measuring using only one detector. If one compares this pattern in the third row of pictures, it is clear that decreasing the power of the lamp illuminating the object also increases the weight of the checkerboard pattern in the recovered image. So, with ambient lights much dimmer than the light coming from the object, one can recover an image which resembles the object with a small artifact coming from the ambient light. However, if the signal coming from the object decreases, artifacts will hide the object under study. This does not happen in the balanced detection case, as the sensor eliminates the fluctuating signal coming from the halogen lamp. Due to this, all the recovered images present similar quality, no matter whether the halogen lamp was turned on or off.

At this point, it is interesting to compare the results of both complementary single-pixel imaging as has been reported^{1,9,20,23,32} and our approach using balanced photodetection. Some previous work on complementary measurements worked with only one reflection arm of the DMD^{1,9,20,32}, which entails the generation of each pattern and its complementary sequentially. In addition of doubling the acquisition time of the system, if the scene changes between a pair of complementary patterns, the coefficient associated with that pair will be corrupted. Even in the case of using both reflection arms of the DMD and two detectors²³, measuring in a non-balanced




Figure 3. Single-pixel imaging in presence of ambient light. Series of measurements using balanced detection (a), and a single sensor (b). For each experiment, both measurements with the halogen lamp turned on and off are shown. Each column represents a different power level of the source illuminating the object, as the power of the halogen lamp is either fixed (lamp on) or zero (lamp off). In order to ease visualization, all images have been normalized between 0 and 255.

scheme introduces both electrical and quantization errors in the process. As the complementary measurements are made at the exact same time using the balanced detector, the technique presents a natural double frame rate and temporal immunity when compared to the sequential complementary approach. Furthermore, the SNR increase stated in the previous section also improves the quality of the recovered images. We show a comparison between non-complementary, sequential, and simultaneous complementary imaging in Fig. 4.

Images shown in Fig. 4 have a size of 64×64 pixels. The total number of projections onto the Walsh-Hadamard basis established by the Nyquist-Shannon criterion is 4096. In order to take the single-pixel traditional image, we generate the 4096 patterns by placing white and black pixels onto the $+1$ s and -1 s entries of the Walsh-Hadamard matrix (shifting and rescaling approach, see Methods). In the sequential case, we generate 8192 patterns and measure the projections only by using one of the reflection arms of our setup. In the balanced setup, the first set of 4096 patterns already generates the full set of 8192 projections by using both reflection arms of the DMD. It is clear that, as stated on the published reports, the complementary approach improves the SNR of the images²³. Both sequential and balanced approaches provide very similar results, but the balanced approach works at double frame rate and it is insensitive to temporal oscillations caused by ambient and parasitic light.

Up to this point, we have remarked the improvement in quality produced by the use of balanced detection. However, there is another key feature of the technique, and it is the increased acquisition speed of the system. By using both reflection arms of the DMD in conjunction with the balanced detector, one can carry out complementary measurements at the same speed as single measurements. By doing this, we gather more information about the object at every pattern generated on the DMD, without decreasing the overall frame rate of the system. To show this benefit, we acquired a video of a moving scene with our setup (see Supplementary Information). We show some of the frames in Fig. 5. In this case, the scene consists of a moving Pac-Man eating a pellet. Every frame of the video is a 64×64 pixel image. We show both SPI and CS acquisitions. The total number of patterns needed to recover a frame is 4096 (64^2), as stated by the Nyquist-Shannon criterion. As our DMD is working at a frequency of 20 kHz, roughly 200 ms are needed to acquire a frame. By doing this, we can achieve loss-less frame rates of 5 Hz.

If more speed is needed, there are several approaches that can boost the frame rate at expense of some quality loss. Here we show two more acquisitions by using undersampling and a CS approach. In order to capture a frame, a subset of the total Walsh-Hadamard basis with low spatial frequencies was generated on the DMD. In




Figure 4. Single-pixel imaging and complementary single-pixel imaging comparison. In the first row, we show images acquired when light coming from the object has 5 times more power than the ambient light. (a) Single-pixel image, (b) sequential complementary single-pixel image, and (c) simultaneous complementary image. In the second row, we present the same experiment but with images acquired when the light coming from the object has the same power as the ambient light. Note that even though second and third columns have very similar quality, the simultaneous complementary image is acquired in half the time.

this case, we used 20% of the total number of patterns. By doing this, we can achieve frame rates of 25 Hz (second row of Fig. 5). It can be seen that even though the quality decreases, the objects can still be clearly identified. The benefit of this approach is that no post-processing is needed, so the image can be displayed in real time. This is a good approach for using the camera in a microscopy-like setup, where the user wants to search for zones of interest in a sample in a fast way. It is possible to perform real-time visualization at 25 Hz, and once the region of interest is in the field of view of the system, perform a loss-less acquisition at 5 Hz. There also exist several ways of getting high frame rates without losing so much quality in the process. In the third row of Fig. 5, we show several frames acquired with the same subset of patterns (a 20% of the total set) but recovered by means of CS. Even though the acquisition speed is 25 Hz, the post-processing time needed to use the CS algorithm hinders the live visualization of the scene.

Lastly, in Fig. 6 we show a comparison between the quality obtained by CS techniques whether performing balanced or non-balanced detection. In order to quantify the faithfulness of the undersampled images, we use the correlation coefficient of each undersampled acquisition with the lossless image (see Methods). The comparison between balanced and non-balanced imaging can be seen in Fig. 6(a). Even though both curves have the same behavior, it is clear that the balanced scheme improves the fidelity of the recovered images. For example, if one looks at the vertical dotted line in the graph, for the same compression ratio, the non-balanced acquisition presents lower correlation coefficients. Watching to the horizontal dotted line, it can be seen that in order to acquire the same correlation coefficient, the non-balanced approach needs to perform a higher number of measurements. In Fig. 6(b) we show both the balanced and non-balanced reference images and three of the undersampled images recovered by CS.

Here we have shown a CS approach, but several adaptive techniques can be used, with the benefits of both increased frame rate and also real time visualization, as they need negligible post-processing times^{1,21}. As balanced detection works on the technical side of the camera (i.e. in the experimental setup), those adaptive approaches can directly benefit from this approach, doubling their frame rates and increasing the SNR of their measurements. In the Supplementary Information, we show the video comparison of the Pac-Man scene in the three cases explained here. With reduced resolutions of 32×32 pixels, we could achieve frame rates of 20 Hz without undersampling. With measurement ratios about 20%, which are common in the literature, we could perform live video at frame rates of 100 Hz, enabling the capture of low-resolution live biological processes.

Discussion

We have implemented balanced photodetection in a single-pixel architecture for the first time. In conjunction with proper collecting optics and complementary single-pixel imaging techniques, we have studied the benefits of the setup when imaging in presence of ambient light. Similar techniques based on single-pixel detection




Figure 5. Balanced single-pixel video imaging. We show multiple frames of a live video taken with our single-pixel camera. In the top row, we show the frames acquired without using subsampling techniques. As the images consist of 64×64 pixels, the 4096 measurements required to take a frame were performed in roughly 200 ms. In the middle row we show frames of the same scene acquired by subsampling the scene. In this case, we used 820 patterns, which correspond to a 20% compression ratio. Each frame was acquired in 40 ms. In the lower row we show a CS reconstruction of the frames obtained before. In this case, the quality drop is almost negligible and we still can record a video at 25 Hz. The numbers shown correspond to the frame number inside the video.

have tackled the same problem with good results²². However, the presented method does not require neither post-processing of the signal nor any *a priori* knowledge of the nature of the ambient light to work. We anticipate that this immunity to parasitic light will be of paramount interest in low-light-level scenarios or when an object is hidden inside a scattering medium, where single-pixel approaches have already started to be used⁴⁻⁶. The use of complementary measurements allows us to increase the signal-to-noise ratio of our measurements. Although similar proposals have been reported previously^{1,9,20,23,32,33}, the implementation shown here increases the frame rate of the devices by a factor of two. We have been able to obtain 64×64 images at a frame rate of 5 Hz without using advanced undersampling techniques. As the proposed technique only requires technical adaptation of the experimental setups, novel digital approaches, ranging from CS to adaptive techniques^{1,20,21}, can benefit from both the SNR and speed boost presented here. Lastly, this kind of setup can be easily adapted to different spectral zones. For instance, inexpensive single-pixel detectors are already available in the infrared region. Therefore, we expect that SPI techniques will be able to provide full real-time imaging feedback for both medicine and industrial applications in the near future.

Methods

Single-pixel and complementary single-pixel imaging. Consider an object represented by a N -pixel image. This image can be arranged in a $N \times 1$ column vector \vec{x} . We can express this vector in an orthonormal basis of functions $\Psi = \{\vec{\Psi}_l\}_{l=1}^N$. We can represent this in matrix form as $\vec{x} = \Psi \cdot \vec{\alpha}$, where Ψ is a $N \times N$ matrix that has the vectors $\{\vec{\Psi}_l\}$ as columns and $\vec{\alpha}$ is the $N \times 1$ vector with the expansion coefficients of \vec{x} in our chosen basis. To acquire an image, SPI techniques carry out those N projections using an SLM, and recover the image of the object by multiplexing the acquired information. In our experiments, we work with the Walsh-Hadamard basis. The basis is conformed by orthogonal discrete square waves, with values either $+1$ or -1 . As our SLM is binary and works by reflection, positive and negative reflection values cannot be readily implemented. To address this, two approaches can be used. The first one consists of shifting and rescaling the Walsh-Hadamard matrices so all their entries consist of either 1 or 0 values⁷. The second one, which we call complementary sensing, consists of generating a pair of matrices, related to the Walsh-Hadamard matrix by a subtraction operation. It has been reported that the complementary scheme improves the signal-to-noise ratio of the measurements and its better suited to use CS strategies²³.

The process goes as follows. We have a Walsh-Hadamard matrix, H , whose entries are either 1 or -1 . We create the complementary pair $H^\pm = (E \pm H)/2$, where E represents a matrix with all entries equal to 1. By doing this,




Figure 6. Balanced detection compressive imaging. (a) comparison between the correlation coefficient of several images with a reference lossless image in both balanced and non-balanced approaches. Vertical and horizontal dotted lines are included to ease the comparison between both approaches. For the same compression ratio, the balanced approach provides higher fidelity. In order to get the same quality, the non-balanced approach needs to perform more acquisitions. In (b) we show both reference images for balanced and non-balanced imaging, and different undersampled images (with 10%, 50% and 90% compression ratios).

we have one matrix H^+ where the original 1 entries preserve their value, but the -1 entries become zero valued. In the other matrix H^- , unity-valued entries will become zero valued, and -1 entries will have a value of 1. It is trivial to prove that all three matrices are related by $H = H^+ - H^-$. As the SPI measuring process is linear, one can acquire the projection of the object under the Walsh-Hadamard basis by measuring the projections under H^+ and H^- and calculating the difference between them.

This procedure reduces the noise introduced by light coming from outside the scene under study. Consider an object, X , expressed as a square matrix. Under homogeneous illumination (P_0), and considering the parasite light as homogeneous additive noise (N) added after object illumination, the light distribution in the DMD plane will be given by $F = X \cdot P_0 + N \cdot E$. Here the products are done element-wise. The measured intensities will be $I^\pm = P_0 \cdot X \cdot H^\pm + N \cdot E \cdot H^\pm$. Given the nature of the Hadamard matrices, $E \cdot H = 0$, and by using this property it is easy to prove that the subtraction of the two signals is given by $\Delta I = I^+ - I^- = P_0 \cdot X \cdot (H^+ - H^-) + \frac{N}{2} - \frac{N}{2} = P_0 \cdot X \cdot H$.

Compressive Sensing. CS provides a method to acquire a picture with $M < N$ measurements, with quality determined by the measurement ratio $\frac{M}{N}$. By choosing a measurement basis $\{\bar{\Phi}_m\}_{m=1}^M$, we can express the measurement process as

$$\bar{y} = \Phi \cdot \bar{x}, \quad (1)$$

where \bar{y} is a $M \times 1$ vector containing the measured projections and Φ is a $M \times N$ matrix called sensing matrix. As the transformation from \bar{x} to \bar{y} entails a dimension reduction, there is loss of information in the process. Given that $M < N$, there exist infinite \bar{x}' that satisfy $\Phi \cdot \bar{x}' = \bar{y}$. CS theory demonstrates that it is possible to recover an approximation to \bar{x} ¹⁷.

To recover the image \bar{x} one needs to solve the underdetermined matrix relation obtained after the measurement process. There are several methods to solve this problem, such as basis pursuit or *Dantzig Selector*³⁴. When working with images, it is also possible to use a model based on the gradient sparsity. Once the discrete gradient of the image is estimated, it is possible to minimize the total variation (TV), which works as a merit function of the gradient, and recover the approximated object. In our experiments, we use the Min-TV with equality constraints algorithm³⁵, that solves the following problem

$$\min \text{TV}(\bar{x}') \text{ subject to } \Phi \cdot \bar{x}' = \bar{y}. \quad (2)$$

Image comparison by correlation coefficient. In order to compare image acquisitions, we have used the correlation coefficient of the undersampled images with a reference image. This coefficient ranges from zero to one, depending on the resemblance of both images. Our reference image is always the image acquired without undersampling (i.e. measuring at the Nyquist-Shannon criterion). The correlation coefficient is calculated with the *corr2* Matlab function

$$r = \frac{\sum_m \sum_n (A_{mn} - \bar{A}) \cdot (B_{mn} - \bar{B})}{\sqrt{\sum_m \sum_n (A_{mn} - \bar{A})^2 \cdot (\sum_m \sum_n (B_{mn} - \bar{B})^2)}} \quad (3)$$

where A and B are the image matrices with indexes m and n. \bar{A} , \bar{B} represent the mean of the elements in A and B.

References

- Radwell, N. *et al.* Single-pixel infrared and visible microscope. *Optica* **1**, 285–289 (2014).
- Watts, C. M. *et al.* Terahertz compressive imaging with metamaterial spatial light modulators. *Nat. Photonics* **8**, 605–609 (2014).
- Greenberg, J., Krishnamurthy, K. & Brady, D. Compressive single-pixel snapshot x-ray diffraction imaging. *Opt. Lett.* **39**, 111–114 (2014).
- Morris, P. A., Aspden, R. S., Bell, J. E. C., Boyd, R. W. & Padgett, M. J. Imaging with a small number of photons. *Nat. Commun.* **6**, 5913 (2015).
- Tajahuerce, E. *et al.* Image transmission through dynamic scattering media by single-pixel photodetection. *Opt. Express* **22**, 16945–16955 (2014).
- Durán, V. *et al.* Compressive imaging in scattering media. *Opt. Express* **23**, 14424–14433 (2015).
- Studer, V. *et al.* Compressive fluorescence microscopy for biological and hyperspectral imaging. *Proc. Natl. Acad. Sci. USA* **109**, E1679–E1687 (2012).
- Howland, G. A., Lum, D. J., Ware, M. R. & Howell, J. C. Photon counting compressive depth mapping. *Opt. Express* **21**, 23822–23837 (2013).
- Soldevila, F. *et al.* Single-pixel polarimetric imaging spectrometer by compressive sensing. *Appl. Phys. B* **113**, 551–558 (2013).
- Clemente, P. *et al.* Compressive holography with a single-pixel detector. *Opt. Lett.* **38**, 2524–2527 (2013).
- Sun, M.-J. *et al.* Single-pixel 3D imaging with time-based depth resolution. 1–10, At <http://arxiv.org/abs/1603.00726> (2016).
- Ducros, M., Goulam Houssen, Y., Bradley, J., de Sars, V. & Charpak, S. Encoded multisite two-photon microscopy. *Proc. Natl. Acad. Sci. USA* **110**, 13138–13143 (2013).
- Liang, J., Gao, L., Li, C. & Wang, L. V. Spatially Fourier-encoded photoacoustic microscopy using a digital micromirror device. *Opt. Lett.* **39**, 430–433 (2014).
- Huynh, N. *et al.* Single-pixel optical camera for video rate ultrasonic imaging. *Optica* **3**, 26 (2016).
- Howland, G. A. & Howell, J. C. Efficient High-Dimensional Entanglement Imaging with a Compressive-Sensing Double-Pixel Camera. *Phys. Rev. X* **3**, 011013 (2013).
- Duarte, M. F. *et al.* Single-Pixel Imaging via Compressive Sampling. *IEEE Signal Process. Mag.* **25**, 83–91 (2008).
- Candès, E. In *Proceedings of the International Congress of Mathematicians Madrid, August 22–30, 2006* 1433–1452, doi: 10.4171/022-3/69 (European Mathematical Society Publishing House, 2006).
- Romberg, J. Imaging via Compressive Sampling. *IEEE Signal Process. Mag.* **25**, 14–20 (2008).
- Sankaranarayanan, A. C., Studer, C. & Baraniuk, R. G. CS-MUVI: Video compressive sensing for spatial-multiplexing cameras. In *2012 IEEE International Conference on Computational Photography (ICCP)* 1–10, doi: 10.1109/ICCP.2012.6215212 (IEEE, 2012).
- Edgar, M. P. *et al.* Simultaneous real-time visible and infrared video with single-pixel detectors. *Sci. Rep.* **5**, 10669 (2015).
- Soldevila, F., Salvador-Balaguer, E., Clemente, P., Tajahuerce, E. & Lancis, J. High-resolution adaptive imaging with a single photodiode. *Sci. Rep.* **5**, 14300 (2015).
- Zhang, Z., Ma, X. & Zhong, J. Single-pixel imaging by means of Fourier spectrum acquisition. *Nat. Commun.* **6**, 6225 (2015).
- Yu, W.-K. *et al.* Complementary compressive imaging for the telescopic system. *Sci. Rep.* **4**, 5834 (2014).
- Ferri, F., Magatti, D., Lugiato, L. A. & Gatti, A. Differential Ghost Imaging. *Phys. Rev. Lett.* **104**, 253603 (2010).
- Sun, B. *et al.* Normalized ghost imaging. *Opt. Express* **20**, 16892–16901 (2012).
- Li, G., Wang, W., Wang, Y., Yang, W. & Liu, L. Single-pixel camera with one graphene photodetector. *Opt. Express* **24**, 400–408 (2016).
- Carleton, H. R. & Maloney, W. T. A Balanced Optical Heterodyne Detector. *Appl. Opt.* **7**, 1241–1243 (1968).
- Vollmann, J., Profunser, D. M. & Dual, J. Sensitivity improvement of a pump-probe set-up for thin film and microstructure metrology. *Ultrasonics* **40**, 757–763 (2002).
- Yang, C., Perelman, L. T., Wax, A., Dasari, R. R. & Feld, M. S. Feasibility of field-based light scattering spectroscopy. *J. Biomed. Opt.* **5**, 138–143 (2000).
- Hobbs, P. C. D. Ultrasensitive laser measurements without tears. *Appl. Opt.* **36**, 903–920 (1997).
- Hee, M. R., Swanson, E. A., Izatt, J. A., Jacobson, J. M. & Fujimoto, J. G. Femtosecond transillumination optical coherence tomography. *Opt. Lett.* **18**, 950–952 (1993).
- Yu, W.-K., Yao, X.-R., Liu, X.-F., Li, L.-Z. & Zhai, G.-J. Three-dimensional single-pixel compressive reflectivity imaging based on complementary modulation. *Appl. Opt.* **54**, 363–367 (2015).
- Lan, R.-M., Liu, X.-F., Yao, X.-R., Yu, W.-K. & Zhai, G.-J. Single-pixel complementary compressive sampling spectrometer. *Opt. Commun.* 1–5, doi: 10.1016/j.optcom.2016.01.016 (2016).
- Candès, E. J. & Romberg, J. 11-magic: Recovery of Sparse Signals via Convex Programming. 1–19, At <http://statweb.stanford.edu/~candes/11magic/downloads/11magic.pdf> Date of access: 24/02/2016 (2005).
- Candès, E. J., Romberg, J. & Tao, T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory* **52**, 489–509 (2006).

Acknowledgements

This work was supported by MINECO through project FIS2013-40666-P, and Generalitat Valenciana projects PROMETEO/2012/021 and ISIC/2012/013. Also financial support by Universitat Jaume I (project P1-1B2012-55) is acknowledged. F.S. was supported by Universitat Jaume I (PREDOC/2013/32). N.U.-P. acknowledges partial support from the National Institutes of Health (P41EB015903).

Author Contributions

J.L., E.T., P.A. and N.U.-P. conceived the experiment. F.S. and P.C. conducted the experiments and wrote the manuscript. All the authors discussed the results and revised the manuscript.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Soldevila, F. *et al.* Computational imaging with a balanced detector. *Sci. Rep.* **6**, 29181; doi: 10.1038/srep29181 (2016).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>



Phase imaging by spatial wavefront sampling

F. SOLDEVILA,^{1,*} V. DURÁN,^{2,3} P. CLEMENTE,^{1,4} J. LANCIS,¹ AND E. TAJAHUERCE¹

¹GROC-UJI, Institute of New Imaging Technologies (INIT), Universitat Jaume I, E12071 Castelló, Spain

²Université Grenoble Alpes, LIPHY, F-38000 Grenoble, France

³CNRS, LIPHY, F-38000 Grenoble, France

⁴Servei Central d'Instrumentació Científica (SCIC), Universitat Jaume I, E12071 Castelló, Spain

*Corresponding author: fsoldevi@uji.es

Received 13 November 2017; revised 26 December 2017; accepted 27 December 2017 (Doc. ID 313294); published 2 February 2018

Phase-imaging techniques extract the optical path length information of a scene, whereas wavefront sensors provide the shape of an optical wavefront. Since these two applications have different technical requirements, they have developed their own specific technologies. Here we show how to perform phase imaging combining wavefront sampling using a reconfigurable spatial light modulator with a beam position detector. The result is a time-multiplexed detection scheme, capable of being shortened considerably by compressive sensing. This robust referenceless method does not require the phase-unwrapping algorithms demanded by conventional interferometry, and its lenslet-free nature removes trade-offs usually found in Shack–Hartmann sensors. © 2018 Optical Society of America under the terms of the OSA Open Access Publishing Agreement

OCIS codes: (110.1758) Computational imaging; (120.5050) Phase measurement; (230.6120) Spatial light modulators; (110.7348) Wavefront encoding.

<https://doi.org/10.1364/OPTICA.5.000164>

1. INTRODUCTION

Even though the physical nature of light has been fully understood for more than a century, there are still no available detectors capable of directly imaging both the amplitude and phase of a wavefront [Fig. 1(a)]. Information about those two quantities is of capital interest when trying to perform biomedical imaging [1,2], aberration measurement and correction in visual optics [3,4], and three-dimensional imaging [5], among other applications. The limitation in performing optical measurements ultimately arises from the extremely fast oscillations of the optical fields, which current electronics are unable to resolve. As detectors only capture light irradiance, several approaches have been proposed over time to tackle the phase problem. Historically, Gabor suggested in 1949 the first quantitative technique, which used interferometric information to recover the complex optical field [6], establishing the basis of modern holography and paving the way for applications such as quantitative phase microscopy [7]. In parallel, phase information plays a fundamental role in adaptive optics, a technique that intends to measure and correct optical aberrations in real time [8]. Although adaptive optics was initially conceived for circumventing atmospheric turbulences in astronomy, its simple operation principle has found applications in other areas, such as visual optics [3,4], microscopy [9,10], and biomedical imaging [11,12].

Two main groups of techniques have emerged to tackle the problem of recovering the complex amplitude of the light field. The first one includes interferometric approaches, which measure the interference between light coming from the object and a

reference beam [Fig. 1(b)]. Although they are extremely powerful for conducting precise phase measurements, their high sensitivity to environmental perturbations (such as mechanical vibrations and changes in temperature) and the need of a reference beam (not always attainable) hinder the implementation of portable and compact interferometric imaging systems in many applications. As an alternative, a second group of techniques has emerged, whose objective is to recover the same information without the need of a reference beam. These non-interferometric approaches rely on several assumptions about the object beam and use mathematical algorithms to infer the wavefront information [13]. Examples of these kinds of techniques are Fourier ptychography [14], coherent diffractive imaging [15,16], phase imaging based on the transport-of-intensity equation [17], and phase imaging with randomized illumination [18]. By eliminating the need for a reference beam, simpler and more robust devices can be designed. Furthermore, computational approaches also enable us to extend the physical capabilities of imaging systems, providing increased field of view [19], optical sectioning [20], or superresolution [21,22]. However, the recovery algorithms used in those techniques usually entail high post-processing times or multiple data acquisitions to recover one image.

Among the non-interferometric techniques, Shack–Hartmann (SH) wavefront sensors [23] are currently the most frequently employed for measuring optical aberrations. SH wavefront sensors combine a lenslet array with a pixelated detector, like a charge-coupled device (CCD) or a CMOS camera. By placing the detector at the focal plane of the lenslet array, the positions of




Fig. 1. Complex amplitude retrieval using spatial wavefront sampling compared to current imaging and wavefront sensing techniques. (a) In a conventional imaging system, the phase information carried by the light coming from an object (for example, a biological sample) is completely lost, as the object image is formed onto a detector that simply measures the light irradiance. (b) Digital holography captures the interference between the light coming from the object and a reference beam, allowing one to retrieve the complex amplitude of the object from irradiance measurements. (c) Shack–Hartmann sensors measure the displacements of different foci generated by an array of lenslets. From those displacements, the wavefront impinging onto the lenslet plane can be reconstructed, providing information about the phase variations introduced by the object. (d) Our technique uses a sequence of illumination amplitude patterns and a single focusing lens. The object’s complex amplitude can be obtained from the statistical properties of the irradiance distribution measured in the Fourier plane of the lens. The same operation principle can be applied if, instead of the configuration shown in the figure, the light coming from an object is modulated by a set of patterns generated on a spatial light modulator (SLM).

the foci generated by each lenslet can be linked to the phase of the wavefront in each region [Fig. 1(c)]. In other words, each lenslet samples the phase information of a region of the wavefront. This method provides a very compact system that is easy to use and align and that requires no complex post-processing stages. For these reasons, SH sensors are used in manifold scientific fields [4,24–27]. However, the number of lenslets, their focal length, and their diameter limit the spatial resolution, dynamic range, and sensibility of SH wavefront sensing. Although multiple solutions have been proposed to manage the several trade-offs between the above magnitudes [28–36], manufacturing processes still place a boundary on the size and curvature of available lenslets, constraining the attainable spatial resolution of commercial SH sensors. Quadriwave lateral shearing interferometry [37], also known as a modified Hartmann mask technique, makes it possible to increase the SH spatial resolution by a factor four [38]. Alternatively, phase imaging can be performed using pyramid wavefront sensors, which employ a four-sided refractive pyramid to split the Fourier plane into four parts, each one generating a phase gradient image on a pixelated detector [39]. Further developments replace the pyramid by a quadri-foil lens [40] or a liquid crystal display [41]. Despite the benefits of this approach in terms of spatial resolution, the illumination numerical aperture determines the detectable dynamic range, so in practice only relatively smooth phase gradients can be precisely recorded [42].

Here, we perform phase imaging using an alternative non-interferometric approach to measure the complex amplitude of a wavefront. We overcome the inherent limitations in spatial resolution, optical efficiency, and dynamic range that are found in SH wavefront sensing. We sample the wavefront by using a set of binary amplitude masks generated by a high-speed SLM. A single focusing lens forms a time-dependent light distribution on its focal plane, where an irradiance detector is placed. Measuring the changes of the total irradiance and the centroid position of that distribution, both the amplitude and phase of the wavefront can be recovered [Fig. 1(d)]. One advantage of this

time-multiplexed detection scheme is that it can be implemented using compressive sensing (CS), allowing us to go beyond the Shannon–Nyquist condition, with the subsequent reduction in acquisition time [43,44]. Unlike previous phase-retrieval methods based on SH wavefront sensors or wavefront modulation [45–49], our approach is lenslet-free and relies neither on any kind of iterative algorithm nor interferometric measurements. Moreover, the use of phase-unwrapping algorithms is not required, such as those employed in some interferometric techniques due to their limited phase unambiguity range [50]. As a proof of concept, we perform aberration sensing as well as phase imaging of a low-contrast sample with variable optical thickness. The corresponding results are compared, respectively, with Shack–Hartmann wavefront sensing and phase-shifting interferometry.

2. OPERATION PRINCIPLE

The basic idea of our technique is to use the relationship between the amplitude and phase of a wavefront and the statistical moments of its Fourier irradiance distribution. This idea is also at the heart of current SH sensors, as they sample a wavefront by means of an array of lenslets and measure the centroid position (the first-order statistical moment) of the Fourier distribution created by each lenslet, allowing one to derive the local slope of the wavefront. By placing a pixelated detector (typically a CCD) at the focal plane of the lenslet array, SH sensors fully map the local slopes of the wavefront. After that measurement, numerical integration provides the phase of the wavefront at each lenslet position [51] [see Fig. 2(a)].

Several tradeoffs arise from the SH configuration. First, the spatial resolution of the measurement is determined by the total number of lenslets and their size. In order to increase the spatial resolution, arrays with higher density of lenslets can be built. However, reducing the size of each lenslet also decreases the amount of light at each region of the detector. If the




Fig. 2. Operation principle of different wavefront sensing approaches. (a) Shack–Hartmann wavefront sensing. A wavefront coming from an object passes through an array of lenslets, which produce a distribution of focal spots on the detector. The position of each spot is linked to the local slope of the wavefront on every lenslet. Numerical integration of the slope data provides the phase of the wavefront. (b) Raster scanning wavefront sensing. A small amplitude aperture is moved over the wavefront plane and, for each consecutive scanning position, the Fourier distribution of the emerging light is recorded. The total irradiance of that distribution at each aperture position can be used to recover the amplitude image of the object. Additionally, the centroid relative location of each Fourier distribution provides the local slope of the wavefront at each scanning position, which allows one to obtain the phase image of the object. (c) Spatial wavefront sampling. Instead of using a small scanning aperture, the wavefront coming from the object is sampled by a set of amplitude masks (reconstruction basis). Now, from the irradiance of each Fourier distribution generated on the detector plane, one can measure the mathematical projection of the object amplitude into the reconstruction basis. By demultiplexing that information, the object amplitude is spatially resolved. In the same way, demultiplexing the data concerning the centroid locations provides the slope map of the wavefront and then, after numerical integration, the object phase image.

signal-to-noise ratio of the measurement is low, the accuracy of the centroid position detection is greatly reduced, which hinders a reliable wavefront reconstruction. Even if the detector can work in low-light-level scenarios, building large arrays of lenslets with a diameter below 100 μm and an accurate curvature is technologically challenging [52]. Moreover, centroid displacement and wavefront slope are related via the lenslet focal length, so, for a given slope, the higher the focal length is, the larger the centroid displacements that are measured. This brings about a second trade-off, since a strongly aberrated wavefront may produce centroid displacements larger than the size of the detection area allocated to every lenslet on the sensor. The result is crosstalk between nearby spots that produces errors in the reconstructed phase, limiting the attainable dynamic range of the sensor. This problem can be circumvented (without sacrificing sensitivity) by a number of techniques that track the real spot location or infer it by using computational techniques [29]. Instead of trying to expand the dynamic range of a sensor by increasing its hardware and/or software complexity, an alternative approach is to implement a reconfigurable (adaptive) SH sensor that includes an array of diffractive lenslets programmed onto a SLM [28,33]. The lenslet characteristics can be then chosen to match the requirements of a specific application optimally, albeit the tradeoff between

sensitivity and dynamic range still remains. Another aspect of SH sensors that is usually not considered is the fact that the total amount of light arriving at each detector region provides a measurement of the light power at the corresponding lenslet position. Using that information, one can map the irradiance of the light coming from an object. However, due to the relatively poor spatial resolution of SH sensors (usually around a thousand lenslets), they barely can compete with interferometric systems to measure the complex amplitude of an object with high spatial frequency content.

Another approach for measuring the local slope of a wavefront is the use of a small scanning aperture and a single lens rather than an array of lenslets [Fig. 2(b)]. For each position of the aperture, only light from a region of the wavefront will be focused by the lens. This light will generate a focal spot on the Fourier plane of the lens where the detector is placed. For a small enough aperture, the wavefront can be locally approximated by a plane wave, and it can be demonstrated (see Methods) that the position of the centroid of the distribution is related to the gradient of the phase inside the aperture by the equation

$$\vec{\Delta} = (\Delta x, \Delta y) = \frac{\lambda f}{2\pi} \vec{\nabla}\phi, \quad (1)$$

where $\vec{\nabla} = (\partial/\partial_x, \partial/\partial_y)$ represents the gradient operator in two dimensions, λ is the light wavelength, and f is the focal length of the focusing lens. One can see the combination of this small aperture and the single lens as one of the lenslets of a SH array. By scanning the wavefront with the aperture, the phase can be easily recovered by numerical integration. Furthermore, calculating the zero-order moment of the light distribution (i.e., its total irradiance) provides a measurement of the wavefront amplitude in the position of the aperture, leading to the recovery of an amplitude image. Although this approach would solve the spatial resolution problem of SH sensors, given that it is easier to generate small apertures than small lenslets, it still presents a trade-off between the spatial resolution and the amount of light at each area of the detector plane, unless the scanning is performed with a laser combined with a galvanometric mirror.

Our technique uses 2D structured illumination patterns as an alternative to a small aperture. As can be seen in Fig. 2(c), instead of just selecting a small region of the wavefront, a time-variable mask (a sequence of programmed patterns) simultaneously selects several regions of the wavefront. This procedure resembles those used in structured illumination microscopy or in single-pixel imaging [44,53,54]. Different single-pixel approaches have been recently proposed to obtain phase information using two-beam [55,56], and common path [57] interferometry. However, our approach is non-interferometric and uses a simpler geometric formulation similar to the one found in Shack–Hartmann sensors. Also, it can be extended to work with partially coherent light, and the detector in use makes it relatively easy to operate in different regions, such as infrared or terahertz (THz). It must be noted that, as in single-pixel imaging, the technique presented here can work in two different configurations: the object (or wavefront) under study can be imaged onto the SLM, or the codified patterns can be projected onto the object plane [58–61]. The key factor that has to be accomplished is that both the mask and the object wavefront overlap in one plane. Even though the irradiance distribution generated by the focusing lens results from the contribution of light coming from different areas of the sampling plane, the amplitude and phase information can be still retrieved, provided that the sequence of sampling patterns is known (see Methods). Now, the amount of light arriving at the detector for every pattern has been greatly increased in comparison to the scanning-aperture scheme. This fact is commonly known as the multiplex or Fellgett’s advantage, and has been extensively exploited to increase the signal-to-noise ratio of optical measurements over the last decades [54,62–64]. In our proposal, the set of masks used are a shifted and rescaled version of the Walsh–Hadamard functions [65], since this election provides several benefits. First, Walsh–Hadamard functions are binary, so they are a good choice when working with binary amplitude modulators, as is the case of digital micromirror devices (DMDs), which reach very high refresh rates (around 22 kHz). Second, the sampling patterns have the same number of absorbing and transmitting (or reflecting) areas, no matter the size of those areas. This fact is important, because once the window size on the modulator has been fixed, increasing or decreasing the spatial resolution of the masks does not change the total amount of light transmitted or reflected by the device. Then, irrespective of the chosen spatial resolution, each measurement always works with the same number of incident photons, unlike the wavefront sensing approaches described above. Furthermore, since in every measurement only a single spatial

light distribution is detected, the crosstalk errors that limit the dynamic range in a SH sensor are not present here. However, there is also a drawback. When operating with Hadamard functions, the number of photons used per measurement decreases by a factor two, since only half of the DMD mirrors are used to illuminate the sample. Nevertheless, this approach makes it possible to use compressive sensing measurements, considerably increasing the acquisition speed.

In order to recover the amplitude and phase information of a wavefront, our method requires the measurement of the zero- and first-order statistical moments of the spatial distribution of light, i.e., its total irradiance and the position of its centroid. The classical approach to measure those quantities is to place a pixelated sensor in the Fourier plane of the lens to acquire a digital image, and then calculate them computationally. However, other experimental approaches can be explored. Here we have chosen a lateral position detector instead of a camera. The benefit of using this detector lies in its capability to provide the information about the power of the beam and the position of its centroid at speeds in the order of several kilohertz (kHz) without the need of computational procedures. Despite digital cameras offering better sensitivity and accuracy in the detection of the centroid, we have opted for a simpler design to demonstrate the feasibility of our approach.

3. RESULTS

A. Experimental Verification

To demonstrate our method, we present the experimental device shown in Fig. 3(a). Light coming from a laser (Oxxius slim-532) emitting at 532 nm is collimated with a lens (L_1) and impinges onto a DMD (DLP Discovery 4100 V-7000 from ViALUX). By using a $4f$ system formed by lenses L_2 and L_3 , light is projected onto an object, which modifies the wavefront phase. After going through it, light passes through a condensing lens (CL), with a focal length of 150 mm. In its Fourier plane, the lateral position detector (Thorlabs PDP90A) measures the irradiance of the beam and the position of its centroid. To retrieve the amplitude and




Fig. 3. Experimental verification of the proposed technique. (a) Schematics of the system. Captions: LS, laser source; L_1 , L_2 , and L_3 , lenses; OBJ, object; CL, condensing lens; PD, position detector. (b) Image of the detector used in our experiments. It includes four electrodes connected to a metallic surface, whose voltage measurements provide both the irradiance of the light beam and the position of its centroid. (c) Experimental results for a plate simulating a coma aberration. As the sample is transparent, the amplitude image provides no information about the object. In the phase image, the object information is clearly recovered. (d) 3D view of the recovered phase.

phase information of the object, a full set of Hadamard patterns is projected. The size of this set depends on the pixel size of the image one wants to obtain. Then, for an $N \times N$ image, N^2 patterns must be sent (see Methods). This sequential acquisition represents a limitation of our technique. Whereas a SH sensor captures all the information in a single shot, our technique relies on the sequential sampling of the wavefront to work.

As a first example, we show the results for a photoresist plate representing a typical coma aberration [Figs. 3(c) and 3(d)]. In this case, the spatial resolution was set to 64×64 pixels with a pixel pitch of $82.08 \mu\text{m}$. For this size, the total acquisition time was roughly 1 s. This is not caused by the refresh rate of our DMD (that would need roughly 200 ms to project 64^2 patterns), but by the limited bandwidth of our detector [Fig. 3(b)]. Faster lateral position sensing detectors, which are commercially available, would speed up the acquisition process. A comparison of this result with those obtained by other wavefront sensing techniques (using a SH sensor and through an interferometric measurement) can be found in Supplement 1.

B. Comparison with Shack–Hartmann Wavefront Sensing

To test the capabilities of the proposed technique, we have compared it with a commercial SH wavefront sensor (Thorlabs WFS150-5C) in several scenarios. As a first object, we employ a spherical lens [Fig. 4(a)] with a focal length of 500 mm aligned with the optical axis. This simple choice facilitates the comparison of both methods quantitatively, since we expect to recover the phase of a spherical wavefront. The object plane is imaged onto the lenslet array plane of the SH wavefront sensor by using a $4f$ system. The focal spots produced by the lenslet array can be seen in Fig. 4(b). Using the displacements of each focal spot with respect to a reference position (previously determined by a calibration measurement), the 3D view of the wavefront can

be recovered by direct integration. To ease the visualization, we also present a wrapped view of the wavefront phase [Fig. 4(c)].

Using the optical setup shown in Fig. 3(a), we reconstruct the wavefront phase by using spatial wavefront sampling. A small region of the full measured 64×64 map of centroid displacements is shown in Fig. 4(d). The modulo- 2π phase reconstructed from those displacements can be observed in Fig. 4(e). The total phase change over the aperture is 9.67λ , a value that differs in 0.06λ from that obtained using the SH sensor.

As a second experiment, we move the lens away from the optical axis, thus producing a larger phase gradient. Given the technical specifications of our SH sensor, the maximum measurable total phase change over the aperture of the sensor is around 100λ . When the phase gradient introduced by the off-axis lens produces a higher phase variation, the focal spots move so much in one direction that crosstalk appears and the sensor is not able to provide reliable results.

In the top row of Fig. 5, we show the result obtained by using the SH sensor at the limit of its dynamic range (here we show a wrapped view of the phase to ease visualization). This limit can be considerably overcome using the setup shown in Fig. 3(a). As can be seen in the bottom row of Fig. 5, spatial wavefront sampling allows us to measure a total gradient phase of 217λ . The above improvement in the dynamic range is due to the relative increase in the size of the detector area attained by our technique. In a SH sensor, the detector is divided into a number of regions equal to the number of lenslets in the array. Each region, containing a small number of pixels, is used to map the position of each focal spot and sets a limit to the maximum spot displacement. As was mentioned before, walking off beyond that limit causes several focal spots to share the same region of the detector, leading to significant errors in the reconstructed phase. Since in our technique we sequentially detect the light distribution as a whole, the full size of the sensor can be used in every consecutive




Fig. 4. Reconstruction of a spherical wavefront and comparison with a commercial Shack–Hartmann sensor. (a) Photograph of the lens used as a phase object. (b) Spot map generated by the SH lenslet array. By comparing the position of each spot to a reference value, the spot displacements can be calculated. From those data, a 3D view of the phase can be recovered. (c) Wrapped phase of the lens. (d) Small region of the full displacement map for our system. Those displacements can be related to the gradient of the phase. Numerical integration of those gradients provides the phase of the object, and after decomposition in the Zernike basis, a high-resolution 3D view of the phase can be displayed (plot on the right). (e) Wrapped phase of the lens obtained using the proposed technique.




Fig. 5. High dynamic-range measurement and comparison with a commercial SH sensor. In the top row, we show the wrapped phase measured with the aid of a SH sensor when the lens shown in Fig. 4(a) is placed in an off-axis position. The maximum phase gradient measurable by the commercial sensor is 100λ . A zoomed region, marked in red, is shown in the right part of the figure. In the bottom row, we show the results for the same object obtained with the system shown in Fig. 3(a). In this case, the lens has been displaced a bigger distance from the optical axis. The phase gradient measured is 217λ . A zoomed region, marked in blue, is shown in the right part of the figure. Given the small size of both figures, some Moiré artifacts appear in the images on the left that are not part of the fringe pattern of a spherical phase gradient (see Supplement 1).

measurement. As a result, the crosstalk penalty typically found in a SH sensor is not present anymore. In the same way as in the design stage of SH sensors, fine tuning of the focal length of the condenser lens used in our setup, as well as a proper selection of the size of the detector, will provide bigger or smaller dynamic ranges and sensitivities in the phase detection.

C. Use of Compressive Sensing

Given the similarities between our technique and single-pixel imaging, some advanced undersampling approaches can be used to improve both the acquisition speed and the light efficiency of our method.

CS provides a recovery framework for signals sampled below the limit imposed by the Shannon–Nyquist theorem [43]. For a sequential signal sampling, this implies a reduction in measurement time. Once the signal has been sampled, an optimization algorithm provides an estimation of the original signal. This estimation is usually based on sparsity constraints and provides a high-fidelity reconstruction (see Methods).

In our experiments, we have used standard CS algorithms (l_1 -magic package [66]) to recover the coma aberration presented in Fig. 3. The results are shown in Fig. 6.

As can be seen, for relatively smooth aberrations, estimations with fidelity higher than 90% can be recovered by using 10% of the total number of measurements established by the Shannon–Nyquist criterion. Figure 6 includes two insets that show individual reconstructions for compression ratios corresponding to correlation coefficients over 0.9 and 0.99, respectively (points A and B in the fidelity plot). As in our commercial Shack–Hartmann sensor, the recovered wavefronts are expressed in the Zernike basis, and high-resolution images are generated from that expansion (see Supplement 1 for details). By taking this into




Fig. 6. Evolution of reconstruction quality when CS is used for wavefront recovery. For different values of the compression ratio (k/N^2 , k being the number of projected patterns), we present the correlation coefficient between the CS estimation and the recovery without using undersampling. Here, the images have a resolution of 64×64 pixels, so $N^2 = 4096$. The points represent the median of the correlation coefficient for 100 different realizations of the algorithm (each realization corresponds to a different subset of random measurements).

account, high-resolution images with acquisition times of 20 ms could be achieved operating at the full frame rate of the DMD (~ 20 kHz), allowing us to perform real-time aberrometry. This could be very useful in ophthalmic scenarios, where the patient needs to stand still while the eye aberrations are measured. Also, using patterned illumination could be beneficial in biological scenarios, where photodamage thresholds can limit the amount of light that can be used at each region of the sample when working on a reflection configuration [67,68]. By using this CS-based procedure, one would be able to increase the measured signal level without causing photodamage.

D. Complex Amplitude Retrieval and Comparison with Phase-Shifting Interferometry

Our technique offers the possibility of reconstructing not only the phase distribution of a wavefront coming from an object but also its amplitude. For the set of illumination patterns, the light power measured by the detector provides the projections of the wavefront amplitude into the basis of Hadamard functions. By measuring all the successive projections, the object amplitude image can be recovered offline, following the operation principle of single-pixel imaging [44] (see Methods). A simple example with an object composed of an amplitude mask attached to the lens used in the previous section can be seen in Supplement 1.

In principle, the complex amplitude of a wavefront can be retrieved by means of a SH sensor, but with a relatively poor spatial resolution (given by the lenslet diameter and the lens fill factor). As a consequence, spatial features smaller than a few hundreds of micrometers are lost in the recovery process. Even though high-resolution SH sensors are being developed and can be used to perform phase imaging, the trade-offs between dynamic range and sensitivity are still present [69]. In our technique, the limit in the spatial resolution is ultimately given by the modulator pixel size (typically $\sim 10 \mu\text{m}$) and the magnification of the projecting system. In our setup, this supposes an increment in the resolution of one order of magnitude with respect to a typical SH sensor.

To demonstrate this improvement, we use as a sample a thin layer of a photoresist material. We placed the photoresist over a transparent plate, creating an object with different regions with




Fig. 7. Comparison with phase-shifting interferometry. (a) Photograph of the photoresist layer used for the experiment. The black square represents the region that will be imaged, consisting of zones with and without photoresist material. (b) Amplitude and phase images obtained with the proposed technique. Due to the absorption properties of the object, the amplitude image presents poor quality. However, in the phase image, fine details of the sample can be observed. (c) 3D representation of the obtained phase. (d) Phase image obtained with phase-shifting interferometry for the same region of the object. In the inset, we show the region between the two small holes present in one of the photoresist bars. (e) Physical profile of the photoresist strip between the two holes obtained with an optical profilometer (Sensofar Plμ 2300).

and without material, as can be seen in Fig. 7(a). Due to the difference of refractive index between the photoresist and the air, an impinging wavefront acquires a spatial phase distribution when it passes through the object. By using the setup shown in Fig. 3(a), the amplitude and phase information of the sample are recovered [Fig. 7(b)]. A 3D visualization of the phase information is shown in Fig. 7(c). The images have a spatial resolution of 128×128 pixels with a pixel pitch of $41.04 \mu\text{m}$. Since both the photoresist and the transparent plate present a similar absorption in the green region of the spectra, the amplitude image presents low quality. However, fine details of the sample are recovered in the phase image, such as the small holes (with a diameter of around $80 \mu\text{m}$) present in the photoresist stripes produced by air bubbles in the manufacturing process. These small spatial features cannot be retrieved by our SH sensor, as the diameter of each lenslet is $150 \mu\text{m}$. Additionally, due to the low number of lenslets ($39 \times 31 = 1209$), the SH sensor would only provide an image clearly “pixelated.” An additional recovery of the same object, using CS techniques, can be seen in Supplement 1.

In order to check the validity of our results, we resort to phase-shifting interferometry. This technique uses a pixelated detector, which makes it possible to exploit the high spatial resolution of a commercial CCD (in our experiment, the pixel size is $6.5 \mu\text{m}$). The recovered phase is shown in Fig. 7(d). Unlike our wavefront sensing technique, phase-shifting interferometry only provides modulo- 2π mappings [see the vertical scale in Fig. 6(d)], requiring the use of unwrapping algorithms to reconstruct a continuous phase like the one shown in Fig. 7(b). We focus our attention on the small region between the two air bubbles in one of the photoresist stripes. After zooming in, it can be seen that the phase difference between the top of the strip and the substrate is 22.7 radians. In our measurements with structured illumination, this difference is 23 radians, showing a very good agreement with the previous value. We also performed a thickness measurement of the region of interest with a mechanical profilometer

(Dektak 6, Veeco). By using the strip height provided by the profilometer ($3.8 \mu\text{m}$), the refractive index of the photoresist layer can be estimated from the optical path length, ($L = \lambda \Delta\phi / 2\pi \Delta n$). In both phase-shifting holography and our technique, the index estimation is 1.51 , which is in good agreement with the nominal value of the photoresist material.

E. Discussion

We have introduced a wavefront sensor based on patterned illumination produced by an amplitude SLM. Instead of resorting to an array of lenslets, as in SH wavefront sensing, the spatial information is captured by illuminating the object with binary amplitude masks generated with a DMD. The sensor is a simple position-sensitive photodetector. This provides several benefits. First, we eliminate the spatial resolution constraints of traditional SH wavefront sensors. Second, by using the full size of the detector in each measurement, the dynamic range is greatly extended and the crosstalk problem is eliminated. Last, the use of Hadamard functions makes the method robust when changing the spatial dimension (number of pixels) to be obtained. The total number of photons used by the system will be determined only by the window size used on the DMD. This is due to the fact that Hadamard patterns always have half of their pixels on each binary value, regardless of the function dimension. Then, increasing or decreasing their spatial dimension does not necessarily change the total bright area displayed by the DMD. This is not the case when working with SH wavefront sensing, as increasing spatial resolution entails decreasing the size of the lenslets, and this reduces the number of photons arriving to each part of the sensor. This, in combination with the multiplex advantage, makes us think that the technique is well suited to work at low-light-level scenarios, where similar approaches have already been proposed to obtain amplitude information [70]. Additionally, we have exploited the high spatial resolution offered by commercial DMDs to perform phase imaging with a pixel size of tens of micrometers.

To this end, we have imaged a photoresist sample and calculated its local refraction index. These results have been demonstrated to be comparable with those provided by phase-shifting interferometry, a well-established interferometric technique.

A distinctive feature of our technique is the fact that the spatial resolution is fixed by the SLM and the projecting optics, while the other optical elements included in the sensor, the focusing lens and the light detector, set the attainable sensitivity and dynamic range. By calculating several statistical moments of the light distribution at the detector plane (its irradiance and the position of its centroid), the method provides both amplitude and phase information. This enables the technique to produce results that are closer to the conventional phase-imaging approaches, but still retaining some of the benefits of the SH approach. These benefits are the referenceless nature of SH wavefront sensing, and the fact that the recovery process does not require phase-unwrapping or iterative algorithms. In addition, the easy implementation of our system also offers the possibility of being used as an add-on module of a conventional microscope [20,71], thus allowing us to carry out quantitative phase imaging with sub-micrometric spatial resolutions.

Due to the temporal multiplexing nature of the technique, there is a trade-off between spatial resolution and image acquisition time. Here, we have shown results with resolutions between 64×64 and 128×128 pixels, which take seconds to be acquired. Using faster, already available detectors will allow image capturing at video frame rates. However, the greater the resolution, the slower the acquisition will be. This drawback can be considerably alleviated using CS [72] or adaptive approaches [73,74]. The feasibility of applying CS has been shown in Section 3, where a typical ocular aberration has been reconstructed with high compression ratios. In particular, acquisitions times of around 20 ms could be achieved operating at the full frame rate of the DMD, which would open up the door to perform real-time aberrometry with our system. Last, given the technological challenge of manufacturing both arrays of lenslets and pixelated sensors that work outside the visible spectrum, the technique proposed here is a good candidate to operate in regions such as IR and THz, where similar approaches have already been used to obtain amplitude images [54,75].

4. METHODS

A. Phase Measurement from Centroid Position

Let us assume a light distribution, at position $z = 0$, with the form: $U_1(x, y) = A(x, y)e^{i\varphi(x, y)}$. If a thin lens, with a focal length f , is placed at the position $z = f$, the complex amplitude distribution at the position $z = 2f$, using Fresnel propagation, will be $U_2(x, y) = \frac{1}{j\lambda f} \tilde{U}_1\left(\frac{x}{\lambda f}, \frac{y}{\lambda f}\right)$, where \tilde{U}_1 denotes the Fourier transform of U_1 . The irradiance of this light distribution will be given by $I_2 = |U_2(x, y)|^2 = \frac{1}{(\lambda f)^2} |\tilde{U}_1\left(\frac{x}{\lambda f}, \frac{y}{\lambda f}\right)|^2$. Our technique is based on the calculus of several statistical moments of I_2 . In particular, we will use the energy, S , and the centroid position, $\Delta = (\Delta x; \Delta y)$. Those quantities will be given by

$$S = \iint_{-\infty}^{\infty} I_2(x, y) dx dy, \quad (2)$$

$$\vec{\Delta} = \begin{pmatrix} \Delta x \\ \Delta y \end{pmatrix} = \frac{1}{S} \begin{pmatrix} \iint_{-\infty}^{\infty} x \cdot I_2(x, y) dx dy \\ \iint_{-\infty}^{\infty} y \cdot I_2(x, y) dx dy \end{pmatrix}. \quad (3)$$

By using Parseval's theorem, it is easy to prove that $S = |U_2(x, y)|^2 = |U_1(x, y)|^2$, and the first measurement provides the energy of the wavefront. Using the Moment's theorem of Fourier transformations and some algebra (see Supplement 1), it can be proved that the position of the centroid at the measurement plane and the phase of the object at the origin position are related by

$$\vec{\Delta} = (\Delta x; \Delta y) = \frac{\lambda f}{2\pi S} \iint_{-\infty}^{\infty} A^2(x, y) \vec{\nabla} \varphi(x, y) dx dy. \quad (4)$$

For a small square aperture with lateral size L , placed at position (a, b) , the amplitude of the field can be described as

$$A(x, y) = K \cdot \text{rect}\left(\frac{x-a}{L}\right) \cdot \text{rect}\left(\frac{y-b}{L}\right), \quad (5)$$

where K is a constant. Introducing this into Eq. (4), we get

$$\begin{aligned} \Delta x &= \frac{\lambda f K^2}{2\pi S} \int_{y=b-L/2}^{y=b+L/2} \int_{x=a-L/2}^{x=a+L/2} \partial_x [\varphi(x, y)] dx dy, \\ \Delta y &= \frac{\lambda f K^2}{2\pi S} \int_{y=b-L/2}^{y=b+L/2} \int_{x=a-L/2}^{x=a+L/2} \partial_y [\varphi(x, y)] dx dy. \end{aligned} \quad (6)$$

If the square aperture is small enough, the gradients of the phase can be considered constant over the integration domain. Then, the centroid positions results,

$$\vec{\Delta} = (\Delta x; \Delta y) = \frac{K^2 \lambda f}{2\pi S} L^2 \vec{\nabla}_{a,b} \varphi(x, y), \quad (7)$$

where $\vec{\nabla}_{a,b} = (\partial/\partial x^{\hat{a}}; \partial/\partial y^{\hat{b}})|_{x=a, y=b}$ represents the bidimensional gradient evaluated at point $x = a, y = b$. By combining Eqs. (5) and (2), it is easy to prove that $S = K^2 L^2$, and thus the centroid of the distribution and the phase of the object are related by

$$\vec{\Delta} = (\Delta x; \Delta y) = \frac{\lambda f}{2\pi} \vec{\nabla}_{a,b} \varphi(x, y). \quad (8)$$

It is clear that, for a small region of the wavefront, the centroid of the resulting distribution provides information about the gradient of the phase in that region. If the small square aperture is displaced, one can measure the correspondent centroid positions and then estimate the gradient map of the wavefront. After that, numerical integration provides the wavefront at the original position. If multiple square masks are used at the same time, it is also possible to relate the measurements with the phase at each one of the mask positions. Then, it is possible to understand Hadamard illumination as a particular case of this procedure, and the phase recovery can be performed while illuminating the full scene with a set of Hadamard patterns (see Supplement 1).

B. Object Recovery by Multiplexing with Hadamard Patterns

Every measurement can be mathematically described by the equation

$$\vec{y} = M\vec{x}, \quad (9)$$

where the object to be measured is described by the object vector \vec{x} , the measurements made are represented by the vector \vec{y} , and the measurement process is carried out by the sensing operator, represented by the matrix M . Once the measurements have been done, one just needs to solve the algebraic problem presented

by Eq. (9) to recover the object. The structure of the sensing operator will be determined by both the nature of the object under study and the system used to recover an image. For a traditional camera setup, matrix M is just the identity matrix, and each element of \vec{x} is proportional to the energy of a given area of the object. By placing a detector at the image plane, the energies of each zone are measured and the object is recovered. Usually this is done with a sensor array, like a CCD camera, and all the measurements are done at the same time. However, there are scenarios, such as confocal microscopy, where this process is conducted by a bucket detector, usually a photomultiplier tube. In this case, the measurement process is sequential. It is possible to exploit Eq. (9) to increase the capabilities of an optical setup. In our approach, instead of just using the identity matrix as the sensing operator, each row of the matrix M contains a Hadamard function. Now, the measurement process consists of making the superposition between the object and each one of the Hadamard functions. After that, the energy of that superposition is measured by the detector. This process can be performed by projecting the functions onto the object with a SLM, as is done in our phase-imaging system. When using this approach, each measurement contains light coming from several regions of the object. This increases the amount of light at each measurement, and thus the signal-to-noise ratio. As Hadamard functions conform an orthonormal basis, Eq. (9) can be easily solved.

Once this way of measuring has been introduced, it is easy to see that it is possible to recover an image of different physical parameters, provided that the measurement process can be described by Eq. (9). It can be noted that our centroid measurements can be arranged in matrix form as

$$\vec{\Delta} = \frac{\lambda f}{2\pi} M \vec{\nabla} \varphi. \quad (10)$$

Here, $\vec{\Delta}$ is the vector containing the measured positions of the centroids, and now the object is the spatial gradient distribution, expressed as a vector, $\vec{\nabla} \varphi$ (each element of this vector is a gradient for a given illumination pattern). Again, in our experiments we used the Hadamard functions as rows of the sensing matrix. For each Hadamard function projected onto the object, the energy of the distribution and the position of its centroid are measured with the position sensing detector. Then, Eq. (10) is solved and the gradient of the phase is obtained. After that, numerical integration provides the phase of the wavefront, $\varphi(x, y)$.

C. Compressive Sensing

Solving the problem given by Eq. (9) is easy when sampling the object at the Shannon–Nyquist ratio. However, for $k < N$ measurements, the equation system presents infinite solutions. The fundamental idea of CS is to use the fact that most of the signals found in nature are sparse in some basis of functions, i.e., they have a representation in some basis where most of the expansion coefficients are zero, so only a small number of them contain relevant information. This fact can be used to provide a solution for the above-mentioned underdetermined equation system. We can express the measurements as $\vec{y} = M\vec{x} = M\Psi\vec{\alpha} = \Phi\vec{\alpha}$, where we have represented the object in a given basis of functions ($\vec{x} = \Psi\vec{\alpha}$). Then, CS algorithms provide a solution to the l_1 -norm minimization problem,

$$\vec{\alpha} = \arg \min \|\vec{\alpha}'\|_1 \text{ such that } \Phi\vec{\alpha}' = \vec{y}. \quad (11)$$

In order to find a good solution, some premises need to be fulfilled [76]. First, the number of measurements cannot be arbitrarily low. Depending on the sparsity of the object (which depends on the chosen recovery basis), a higher or lower number of measurements is needed to obtain good results. Second, the basis pair (measurement and recovery) need to be incoherent. In practice, it is usually easy to find pairs of basis that fulfill this principle. Usually, the measurement basis is chosen for its ease of generation in a SLM, and once one has been picked up, the user selects a recovery basis that fulfills the incoherence constraint and where the object is sparse. Last, a sensing strategy needs to be defined. Initially, it was proposed that randomly choosing a subset of elements in the measurement basis was enough to obtain good results [43]. Soon after that, other strategies were defined to reduce the number of measurements maintaining the image quality. For example, it is known that natural scenes tend to have a power spectrum centered around low frequencies, so mixed sampling strategies have been used with very good results [63].

In our experiments, M represents the shifted and rescaled Hadamard basis (the elements of which are either 0 or 1), easily generated on a DMD, and Ψ is chosen to be the Haar wavelet basis, where most images tend to be sparse. The measurements have been performed following the mixed approach: fixing the number of samples, k , we chose the lowest frequency $k/2$ Hadamard patterns, and then we select $k/2$ random Hadamard patterns from the remaining elements of the complete basis.

D. Phase-Shifting Holography Measurements

For the measurements using phase-shifting digital holography, we used an interferometer in a Mach–Zehnder configuration. A collimated laser (Oxxius slim-532) was divided by a beam splitter into the object and the reference beam. The first one illuminated the object while the second traveled directly towards the camera (Allied Stingray F-145). The object, a layer of photoresist spin coated onto a transparent plate, was imaged onto the camera by an optical system in a $4f$ configuration. The phase shifts were generated by shifting a grating codified onto a DMD (DLP Discovery 4100) and filtering the first diffraction order in the Fourier plane of a $4f$ optical configuration located before the object plane [77]. The camera recorded four interferograms with a phase shift interval equal to $\pi/2$ and a standard phase-shifting algebraic operation [78] was used to measure the amplitude and phase at the object plane.

Funding. Generalitat Valenciana (PROMETEO/2016/079); Ministerio de Economía y Competitividad (MINECO) (FIS2015-72872-EXP, FIS2016-75618-R); Universitat Jaume I (UJI) (P1-1B2015-35, PREDOC/2013/32).

Acknowledgment. We thank Salvador Bará for the photoresist plates simulating optical aberrations used in our experiments. Author contributions: F. S. and E. T. conceived the idea; F. S. performed the experiments with the aid of P. C. and V. D.; F. S. and V. D. wrote the paper; all the authors discussed the results and revised the paper; J. L. and E. T. coordinated the project.

See [Supplement 1](#) for supporting content.

REFERENCES

- F. Zernike, "How I discovered phase contrast," *Science* **121**, 345–349 (1955).
- G. Popescu, "The power of imaging with phase, not power," *Phys. Today* **70**(5), 34–40 (2017).
- J. Liang, B. Grimm, S. Goelz, and J. F. Bille, "Objective measurement of wave aberrations of the human eye with the use of a Hartmann–Shack wave-front sensor," *J. Opt. Soc. Am. A* **11**, 1949–1957 (1994).
- P. Artal, "Optics of the eye and its impact in vision: a tutorial," *Adv. Opt. Photon.* **6**, 340–367 (2014).
- G. Nehmetallah and P. P. Banerjee, "Applications of digital and analog holography in three-dimensional imaging," *Adv. Opt. Photon.* **4**, 472–553 (2012).
- D. Gabor, "Microscopy by reconstructed wave-fronts," *Proc. R. Soc. A* **197**, 454–487 (1949).
- M. Mir, B. Bhaduri, R. Wang, R. Zhu, and G. Popescu, "Quantitative phase imaging," in *Progress in Optics* (Elsevier, 2012), Vol. **57**, pp. 133–217.
- R. Tyson, *Principles of Adaptive Optics*, Series in Optics and Optoelectronics, 3rd ed. (CRC Press, 2010), Vol. 20102628.
- M. J. Booth, "Adaptive optical microscopy: the ongoing quest for a perfect image," *Light Sci. Appl.* **3**, e165 (2014).
- N. Ji, "Adaptive optical fluorescence microscopy," *Nat. Methods* **14**, 374–380 (2017).
- K. Wang, D. E. Milkie, A. Saxena, P. Engerer, T. Misgeld, M. E. Bronner, J. Mumm, and E. Betzig, "Rapid adaptive optical recovery of optimal resolution over large volumes," *Nat. Methods* **11**, 625–628 (2014).
- A. Kumar, W. Drexler, and R. A. Leitgeb, "Subaperture correlation based digital adaptive optics for full field optical coherence tomography," *Opt. Express* **21**, 10850–10866 (2013).
- Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase retrieval with application to optical imaging: a contemporary overview," *IEEE Signal Process. Mag.* **32**, 87–109 (2015).
- G. Zheng, R. Horstmeyer, and C. Yang, "Wide-field, high-resolution Fourier ptychographic microscopy," *Nat. Photonics* **7**, 739–745 (2013).
- H. N. Chapman and K. A. Nugent, "Coherent lensless x-ray imaging," *Nat. Photonics* **4**, 833–839 (2010).
- R. Horisaki, Y. Ogura, M. Aino, and J. Tanida, "Single-shot phase imaging with a coded aperture," *Opt. Lett.* **39**, 6466–6469 (2014).
- N. Streibl, "Phase imaging by the transport equation of intensity," *Opt. Commun.* **49**, 6–10 (1984).
- R. Horisaki, R. Egami, and J. Tanida, "Single-shot phase imaging with randomized light (SPIRAL)," *Opt. Express* **24**, 3765–3773 (2016).
- L. Tian, X. Li, K. Ramchandran, and L. Waller, "Multiplexed coded illumination for Fourier ptychography with an LED array microscope," *Biomed. Opt. Express* **5**, 2376–2389 (2014).
- L. Tian and L. Waller, "3D intensity and phase imaging from light field measurements in an LED array microscope," *Optica* **2**, 104–111 (2015).
- J. J. Field, K. A. Wernsing, S. R. Domingue, A. M. Allende Motz, K. F. DeLuca, D. H. Levi, J. G. DeLuca, M. D. Young, J. A. Squier, and R. A. Bartels, "Superresolved multiphoton microscopy with spatial frequency-modulated imaging," *Proc. Natl. Acad. Sci. USA* **113**, 6605–6610 (2016).
- P. Gao, G. Pedrini, and W. Osten, "Phase retrieval with resolution enhancement by using structured illumination," *Opt. Lett.* **38**, 5204–5207 (2013).
- B. C. Platt and R. Shack, "History and principles of Shack–Hartmann wavefront sensing," *J. Refract. Surg.* **17**, S573–S577 (2001).
- B. Stoklasa, L. Motka, J. Rehacek, Z. Hradil, and L. L. Sánchez-Soto, "Wavefront sensing reveals optical coherence," *Nat. Commun.* **5**, 3275 (2014).
- R. W. Wilson, "SLODAR: measuring optical turbulence altitude with a Shack–Hartmann wavefront sensor," *Mon. Not. R. Astron. Soc.* **337**, 103–108 (2002).
- D. Dayton, J. Gonglewski, B. Pierson, and B. Spielbusch, "Atmospheric structure function measurements with a Shack–Hartmann wave-front sensor," *Opt. Lett.* **17**, 1737–1739 (1992).
- X. Cui, J. Ren, G. J. Tearney, and C. Yang, "Wavefront image sensor chip," *Opt. Express* **18**, 16685–16701 (2010).
- L. Seifert, J. Liesener, and H. J. Tiziani, "The adaptive Shack–Hartmann sensor," *Opt. Commun.* **216**, 313–319 (2003).
- Y. Saita, H. Shinto, and T. Nomura, "Holographic Shack–Hartmann wavefront sensor based on the correlation peak displacement detection method for wavefront sensing with large dynamic range," *Optica* **2**, 411–415 (2015).
- S. Thomas, T. Fusco, A. Tokovinin, M. Nicolle, V. Michau, and G. Rousset, "Comparison of centroid computation algorithms in a Shack–Hartmann sensor," *Mon. Not. R. Astron. Soc.* **371**, 323–336 (2006).
- J. Pfund, N. Lindlein, and J. Schwider, "Dynamic range expansion of a Shack–Hartmann sensor by use of a modified unwrapping algorithm," *Opt. Lett.* **23**, 995–997 (1998).
- S. Wang, P. Yang, B. Xu, L. Dong, and M. Ao, "Shack–Hartmann wavefront sensing based on binary-aberration-mode filtering," *Opt. Express* **23**, 5052–5063 (2015).
- R. Martínez-Cuenca, V. Durán, V. Climent, E. Tajahuerce, S. Bará, J. Ares, J. Arines, M. Martínez-Corral, and J. Lancis, "Reconfigurable Shack–Hartmann sensor without moving elements," *Opt. Lett.* **35**, 1338–1340 (2010).
- T. Godin, M. Fromager, E. Cagniot, M. Brunel, and K. Ait-Ameur, "Reconstruction-free sensitive wavefront sensor based on continuous position sensitive detectors," *Appl. Opt.* **52**, 8310–8317 (2013).
- R. Navarro and E. Moreno-Barriuso, "Laser ray-tracing method for optical testing," *Opt. Lett.* **24**, 951–953 (1999).
- S. Olivier, V. Laude, and J. Huignard, "Liquid-crystal Hartmann wavefront scanner," *Appl. Opt.* **39**, 3838–3846 (2000).
- J.-C. Chanteloup, "Multiple-wave lateral shearing interferometry for wave-front sensing," *Appl. Opt.* **44**, 1559–1571 (2005).
- "Phasics," <http://phasicscorp.com>.
- I. Iglesias, "Pyramid phase microscopy," *Opt. Lett.* **36**, 3636–3638 (2011).
- A. B. Parthasarathy, K. K. Chu, T. N. Ford, and J. Mertz, "Quantitative phase imaging using a partitioned detection aperture," *Opt. Lett.* **37**, 4062–4064 (2012).
- I. Iglesias and F. Vargas-Martin, "Quantitative phase microscopy of transparent samples using a liquid crystal display," *J. Biomed. Opt.* **18**, 026015 (2013).
- H. Lu, J. Chung, X. Ou, and C. Yang, "Quantitative phase imaging and complex field reconstruction by pupil modulation differential phase contrast," *Opt. Express* **24**, 25345–25361 (2016).
- E. J. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.* **25**, 21–30 (2008).
- M. F. F. Duarte, M. A. A. Davenport, D. Takhar, J. N. N. Laska, K. F. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.* **25**, 83–91 (2008).
- H. Gong, O. Soloviev, D. Wilding, P. Pozzi, M. Verhaegen, and G. Vdovin, "Holographic imaging with a Shack–Hartmann wavefront sensor," *Opt. Express* **24**, 13729–13737 (2016).
- F. Zhang, G. Pedrini, and W. Osten, "Phase retrieval of arbitrary complex-valued fields through aperture-plane modulation," *Phys. Rev. A* **75**, 043805 (2007).
- P. Gao, G. Pedrini, C. Zuo, and W. Osten, "Phase retrieval using spatially modulated illumination," *Opt. Lett.* **39**, 3615–3618 (2014).
- S. Chowdhury and J. Izatt, "Structured illumination quantitative phase microscopy for enhanced resolution amplitude and phase imaging," *Biomed. Opt. Express* **4**, 1795–1805 (2013).
- S. Chowdhury and J. Izatt, "Structured illumination diffraction phase microscopy for broadband, subdiffraction resolution, quantitative phase imaging," *Opt. Lett.* **39**, 1015–1018 (2014).
- G. Vdovin, H. Gong, O. Soloviev, P. Pozzi, and M. Verhaegen, "Lensless coherent imaging by sampling of the optical field with digital micromirror device," *J. Opt.* **17**, 122001 (2015).
- W. H. Southwell, "Wave-front estimation from wave-front slope measurements," *J. Opt. Soc. Am.* **70**, 998–1006 (1980).
- R. Kasztelanica, A. Filipkowski, D. Pysz, R. Stepien, A. J. Waddie, M. R. Taghizadeh, and R. Buczynski, "High resolution Shack–Hartmann sensor based on array of nanostructured GRIN lenses," *Opt. Express* **25**, 1680–1691 (2017).
- M. Saxena, G. Eluru, and S. S. Gorthi, "Structured illumination microscopy," *Adv. Opt. Photon.* **7**, 241–275 (2015).
- C. M. Watts, D. Shrekenhamer, J. Montoya, G. Lipworth, J. Hunt, T. Sleasman, S. Krishna, D. R. Smith, and W. J. Padilla, "Terahertz compressive imaging with metamaterial spatial light modulators," *Nat. Photonics* **8**, 605–609 (2014).
- P. Clemente, V. Durán, E. Tajahuerce, P. Andrés, V. Climent, and J. Lancis, "Compressive holography with a single-pixel detector," *Opt. Lett.* **38**, 2524–2527 (2013).

56. L. Martínez-León, P. Clemente, Y. Mori, V. Climent, J. Lancis, and E. Tajahuerce, "Single-pixel digital holography with phase-encoded illumination," *Opt. Express* **25**, 4975–4984 (2017).
57. P. A. Stockton, J. J. Field, and R. A. Bartels, "Single pixel quantitative phase imaging with spatial frequency projections," *Methods* (2017), in press.
58. F. Soldevila, E. Irlles, V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis, "Single-pixel polarimetric imaging spectrometer by compressive sensing," *Appl. Phys. B* **113**, 551–558 (2013).
59. F. Soldevila, P. Clemente, E. Tajahuerce, N. Uribe-Patarroyo, P. Andrés, and J. Lancis, "Computational imaging with a balanced detector," *Sci. Rep.* **6**, 29181 (2016).
60. E. Tajahuerce, V. Durán, P. Clemente, E. Irlles, F. Soldevila, P. Andrés, and J. Lancis, "Image transmission through dynamic scattering media by single-pixel photodetection," *Opt. Express* **22**, 16945–16955 (2014).
61. V. Durán, F. Soldevila, E. Irlles, P. Clemente, E. Tajahuerce, P. Andrés, and J. Lancis, "Compressive imaging in scattering media," *Opt. Express* **23**, 14424–14433 (2015).
62. P. R. Griffiths, H. J. Sloane, and R. W. Hannah, "Interferometers vs monochromators: separating the optical and digital advantages," *Appl. Spectrosc.* **31**, 485–495 (1977).
63. V. Studer, J. Bobin, M. Chahid, H. S. Mousavi, E. Candes, and M. Dahan, "Compressive fluorescence microscopy for biological and hyperspectral imaging," *Proc. Natl. Acad. Sci. USA* **109**, E1679–E1687 (2012).
64. M. Ducros, Y. Goulam Houssen, J. Bradley, V. de Sars, and S. Charpak, "Encoded multisite two-photon microscopy," *Proc. Natl. Acad. Sci. USA* **110**, 13138–13143 (2013).
65. W. Pratt, J. Kane, and H. Andrews, "Hadamard transform image coding," *Proc. IEEE* **57**, 58–68 (1969).
66. E. Candès and J. Romberg, "11-magic: recovery of sparse signals via convex programming," https://www.cs.bham.ac.uk/~axk/Sakinah/inspiring_readings/11magic.pdf.
67. B. Lochocki, A. Gambín, S. Manzanera, E. Irlles, E. Tajahuerce, J. Lancis, and P. Artal, "Single pixel camera ophthalmoscope," *Optica* **3**, 1056–1059 (2016).
68. Y. Wu, P. Ye, I. O. Mirza, G. R. Arce, and D. W. Prather, "Experimental demonstration of an optical-sectioning compressive sensing microscope (CSM)," *Opt. Express* **18**, 24565–24578 (2010).
69. H. Gong, T. E. Agbana, P. Pozzi, O. Soloviev, M. Verhaegen, and G. Vdovin, "Optical path difference microscopy with a Shack–Hartmann wavefront sensor," *Opt. Lett.* **42**, 2122–2125 (2017).
70. P. A. Morris, R. S. Aspden, J. E. C. Bell, R. W. Boyd, and M. J. Padgett, "Imaging with a small number of photons," *Nat. Commun.* **6**, 5913 (2015).
71. Z. Wang, L. Millet, M. Mir, H. Ding, S. Unarunotai, J. Rogers, M. U. Gillette, and G. Popescu, "Spatial light interference microscopy (SLIM)," *Opt. Express* **19**, 1016–1026 (2011).
72. E. Candès, "Compressive sampling," in *Proceedings of the International Congress of Mathematicians* (European Mathematical Society, 2006), pp. 1433–1452.
73. F. Soldevila, E. Salvador-Balaguer, P. Clemente, E. Tajahuerce, and J. Lancis, "High-resolution adaptive imaging with a single photodiode," *Sci. Rep.* **5**, 14300 (2015).
74. H. Dai, G. Gu, W. He, L. Ye, T. Mao, and Q. Chen, "Adaptive compressed photon counting 3D imaging based on wavelet trees and depth map sparse representation," *Opt. Express* **24**, 26080–26096 (2016).
75. N. Radwell, K. J. Mitchell, G. Gibson, M. Edgar, R. Bowman, and M. J. Padgett, "Single-pixel infrared and visible microscope," *Optica* **1**, 285–289 (2014).
76. J. Romberg, "Imaging via compressive sampling," *IEEE Signal Process. Mag.* **25**, 14–20 (2008).
77. D. B. Conkey, A. M. Caravaca-Aguirre, and R. Piestun, "High-speed scattering medium characterization with application to focusing light through turbid media," *Opt. Express* **20**, 1733–1740 (2012).
78. I. Yamaguchi and T. Zhang, "Phase-shifting digital holography," *Opt. Lett.* **22**, 1268–1270 (1997).

Bibliography

1. Darwin, C. Journal of Researches into the Geology and Natural History of the various countries visited by H.M.S. Beagle, under the command of Captain Fitz-Roy, R.N., from 1832 to 1836. *The Quarterly Review* **65**, 194–234 (1839).
2. O'Connor, J. *Male medium tree finch* https://commons.wikimedia.org/wiki/File:MTF_male.jpg.
3. Buccafusca, O. *Results from OPN's third annual After Image Photo Contest 2008*. https://www.osa-opn.org/home/articles/volume_19/issue_11/departments/none/and_the_winners_are_results_from_opn's_third_a/ (2018).
4. Masur. *Saccharomyces cerevisiae cells in DIC microscopy 2010*. https://commons.wikimedia.org/wiki/File:S_cerevisiae_under_DIC_microscopy.jpg.
5. Youn, S. H. *et al.* Optical performance test and validation of microcameras in multiscale, gigapixel imagers. *Optics Express* **22**, 3712. ISSN: 1094-4087 (2014).
6. Wang, N. & ElMasry, G. in *Hyperspectral Imaging for Food Quality Analysis and Control* 295–320 (Elsevier, 2010). ISBN: 9780123747532. doi:10.1016/B978-0-12-374753-2.10009-7. <http://linkinghub.elsevier.com/retrieve/pii/B9780123747532100097>.
7. Billinghurst, M., Clark, A. & Lee, G. A Survey of Augmented Reality. *Foundations and Trends® in Human–Computer Interaction* **8**, 73–272. ISSN: 1551-3955 (2015).
8. Tabernero, J., Piers, P., Benito, A., Redondo, M. & Artal, P. Predicting the Optical Performance of Eyes Implanted with IOLs to Correct Spherical Aberration. *Investigative Ophthalmology & Visual Science* **47**, 4651. ISSN: 1552-5783 (2006).
9. Qi, J., He, C. & Elson, D. S. Real time complete Stokes polarimetric imager based on a linear polarizer array camera for tissue polarimetric imaging. *Biomedical Optics Express* **8**, 4933. ISSN: 2156-7085 (2017).
10. Deisseroth, K. Optogenetics. *Nature Methods* **8**, 26–29. ISSN: 1548-7091 (2011).
11. Bates, M., Huang, B., Dempsey, G. T. & Zhuang, X. Multicolor Super-Resolution Imaging with Photo-Switchable Fluorescent Probes. *Science* **317**, 1749–1753. ISSN: 0036-8075 (2007).
12. Wong, M. D., Dorr, A. E., Walls, J. R., Lerch, J. P. & Henkelman, R. M. A novel 3D mouse embryo atlas based on micro-CT. *Development* **139**, 3248–3256. ISSN: 0950-1991 (2012).
13. Wenner, M. The most transparent research. *Nature medicine* **15**, 1106–1109. ISSN: 1078-8956 (2009).
14. Duarte, M. F. *et al.* Single-Pixel Imaging via Compressive Sampling. *IEEE Signal Processing Magazine* **25**, 83–91. ISSN: 1053-5888 (2008).
15. Sun, B. *et al.* 3D Computational Imaging with Single-Pixel Detectors. *Science* **340**, 844–847. ISSN: 0036-8075 (2013).

16. Bechhoefer, J. What is superresolution microscopy? *American Journal of Physics* **83**, 22–29. ISSN: 0002-9505 (2015).
17. Romberg, J. *et al.* Imaging via Compressive Sampling. *IEEE Signal Processing Magazine* **25**, 14–20. ISSN: 1053-5888 (2008).
18. Chen, H, Xi, N., Song, B. & Lai, K. Single pixel infrared camera using a carbon nanotube photodetector. *Sensors, 2011 IEEE*. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6127226 (2011).
19. Radwell, N. *et al.* Single-pixel infrared and visible microscope. *Optica* **1**, 285–289 (2014).
20. Li, G., Wang, W., Wang, Y., Yang, W. & Liu, L. Single-pixel camera with one graphene photodetector. *Optics Express* **24**, 400. ISSN: 1094-4087 (2016).
21. Gibson, G. M. *et al.* Real-time imaging of methane gas leaks using a single-pixel camera. *Optics Express* **25**, 2998. ISSN: 1094-4087 (2017).
22. Hobbs, M. J., Grainger, M. P., Zhu, C., Tan, C. H. & Willmott, J. R. Quantitative thermal imaging using single-pixel Si APD and MEMS mirror. *Optics Express* **26**, 3188. ISSN: 1094-4087 (2018).
23. Chan, W. L. *et al.* A single-pixel terahertz imaging system based on compressed sensing. *Applied Physics Letters* **93**, 121105. ISSN: 0003-6951 (2008).
24. Li, C., Grant, J., Wang, J. & Cumming, D. R. S. A Nipkow disk integrated with Fresnel lenses for terahertz single pixel imaging. *Optics Express* **21**, 24452. ISSN: 1094-4087 (2013).
25. Watts, C. M. *et al.* Terahertz compressive imaging with metamaterial spatial light modulators. *Nature Photonics* **8**, 605–609. ISSN: 1749-4885 (2014).
26. Edgar, M. P. *et al.* Simultaneous real-time visible and infrared video with single-pixel detectors. *Scientific Reports* **5**, 10669. ISSN: 2045-2322 (2015).
27. Stantchev, R. I. *et al.* Noninvasive, near-field terahertz imaging of hidden objects using a single-pixel detector. *Science Advances* **2**, e1600190–e1600190. ISSN: 2375-2548 (2016).
28. Dicke, R. H. Scatter-Hole Cameras for X-Rays and Gamma Rays. *The Astrophysical Journal* **153**, L101. ISSN: 0004-637X (1968).
29. Greenberg, J., Krishnamurthy, K. & Brady, D. Compressive single-pixel snapshot x-ray diffraction imaging. *Optics Letters* **39**, 111–114. ISSN: 0146-9592 (2014).
30. Chen, S. X-ray ‘ghost images’ could cut radiation doses. *Science* **359**, 1452–1452. ISSN: 0036-8075 (2018).
31. Zhang, A.-X., He, Y.-H., Wu, L.-A., Chen, L.-M. & Wang, B.-B. Tabletop x-ray ghost imaging with ultra-low radiation. *Optica* **5**, 374. ISSN: 2334-2536 (2018).
32. Lochocki, B. *et al.* Single pixel camera ophthalmoscope. *Optica* **3**, 1056. ISSN: 2334-2536 (2016).
33. Phillips, D. B. *et al.* Adaptive foveated single-pixel imaging with dynamic supersampling. *Science Advances* **3**, e1601782. ISSN: 2375-2548 (2017).
34. Morris, P. A., Aspden, R. S., Bell, J. E. C., Boyd, R. W. & Padgett, M. J. Imaging with a small number of photons. *Nature Communications* **6**, 5913. ISSN: 2041-1723 (2015).
35. Tajahuerce, E. *et al.* Image transmission through dynamic scattering media by single-pixel photodetection. *Optics Express* **22**, 16945. ISSN: 1094-4087 (2014).

36. Durán, V. *et al.* Compressive imaging in scattering media. *Optics Express* **23**, 14424–14433. ISSN: 1094-4087 (2015).
37. Rodríguez, A., Clemente, P., Tajahuerce, E. & Lancis, J. Dual-mode optical microscope based on single-pixel imaging. *Optics and Lasers in Engineering* **82**, 87–94. ISSN: 01438166 (2016).
38. Clemente, P., Durán, V., Torres-Company, V., Tajahuerce, E. & Lancis, J. Optical encryption based on computational ghost imaging. *Optics Letters* **35**, 2391. ISSN: 0146-9592 (2010).
39. Zhang, Z., Jiao, S., Yao, M., Li, X. & Zhong, J. Secured single-pixel broadcast imaging. *Optics Express* **26**, 14578. ISSN: 1094-4087 (2018).
40. Li, Z. *et al.* Efficient single-pixel multispectral imaging via non-mechanical spatio-spectral modulation. *Scientific Reports* **7**, 41435. ISSN: 2045-2322 (2017).
41. Rousset, F. *et al.* Time-resolved multispectral imaging based on an adaptive single-pixel camera. *Optics Express* **26**, 10550. ISSN: 1094-4087 (2018).
42. Durán, V., Clemente, P., Fernández-Alonso, M., Tajahuerce, E. & Lancis, J. Single-pixel polarimetric imaging. *Optics Letters* **37**, 824. ISSN: 0146-9592 (2012).
43. Welsh, S. S., Edgar, M. P., Bowman, R., Sun, B. & Padgett, M. J. Near video-rate linear Stokes imaging with single-pixel detectors. *Journal of Optics* **17**, 025705. ISSN: 2040-8978 (2015).
44. Salvador-Balaguer, E., Clemente, P., Tajahuerce, E., Pla, F. & Lancis, J. Full-color stereoscopic imaging with a single-pixel photodetector. *Journal of Display Technology* **12**, 1–1. ISSN: 1551-319X (2015).
45. Salvador-Balaguer, E. *et al.* Low-cost single-pixel 3D imaging by using an LED array. *Optics Express* **26**, 15623. ISSN: 1094-4087 (2018).
46. Zhang, Z. *et al.* Simultaneous spatial, spectral, and 3D compressive imaging via efficient Fourier single-pixel measurements. *Optica* **5**, 315. ISSN: 2334-2536 (2018).
47. Martínez-León, L. *et al.* Single-pixel digital holography with phase-encoded illumination. *Optics Express* **25**, 4975. ISSN: 1094-4087 (2017).
48. Shibuya, K. *et al.* Scan-less hyperspectral dual-comb single-pixel-imaging in both amplitude and phase. *Optics Express* **25**, 21947. ISSN: 1094-4087 (2017).
49. González, H. *et al.* High sampling rate single-pixel digital holography system employing a DMD and phase-encoded patterns. *Optics Express* **26**, 20342. ISSN: 1094-4087 (2018).
50. Soldevila, F. *et al.* Single-pixel polarimetric imaging spectrometer by compressive sensing. *Applied Physics B* **113**, 551–558. ISSN: 0946-2171 (2013).
51. Soldevila, F. *et al.* in *Multidimensional Imaging* (eds Javidi, B., Tajahuerce, E. & Andrés, P.) 371–390 (Wiley, Chichester, UK, 2014). ISBN: 978-1-118-44983-7. doi:10.1002/9781118705766.ch16. <http://eu.wiley.com/WileyCDA/WileyTitle/productCd-1118449835.html><http://doi.wiley.com/10.1002/9781118705766.ch16>.
52. Soldevila, F., Durán, V., Clemente, P., Lancis, J. & Tajahuerce, E. Phase imaging by spatial wavefront sampling. *Optica* **5**, 164. ISSN: 2334-2536 (2018).
53. Soldevila, F., Salvador-Balaguer, E., Clemente, P., Tajahuerce, E. & Lancis, J. High-resolution adaptive imaging with a single photodiode. *Scientific Reports* **5**, 14300. ISSN: 2045-2322 (2015).

54. Soldevila, F. *et al.* Computational imaging with a balanced detector. *Scientific Reports* **6**, 29181. ISSN: 2045-2322 (2016).
55. Soldevila, F., Clemente, P., Tajahuerce, E. & Lancis, J. in *Encyclopedia of Modern Optics* 193–198 (Elsevier, 2018). ISBN: 9780128035818. doi:10.1016/B978-0-12-803581-8.09378-4. <https://linkinghub.elsevier.com/retrieve/pii/B9780128035818093784>.
56. Hansen, P. C. *Rank-Deficient and Discrete Ill-Posed Problems* ISBN: 978-0-89871-403-6. doi:10.1137/1.9780898719697. <http://epubs.siam.org/doi/book/10.1137/1.9780898719697> (Society for Industrial and Applied Mathematics, 1998).
57. Candes, E. *et al.* Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics* **59**, 1207–1223. ISSN: 0010-3640 (2006).
58. Donoho, D. Compressed sensing. *IEEE Transactions on Information Theory* **52**, 1289–1306. ISSN: 0018-9448 (2006).
59. Candès, E. in *Proceedings of the International Congress of Mathematicians Madrid, August 22–30, 2006* 1433–1452 (European Mathematical Society Publishing House, Zurich, Switzerland, 2006). ISBN: 978-3-03719-022-7. doi:10.4171/022-3/69. <http://www.ems-ph.org/doi/10.4171/022-3/69>.
60. Candès, E. & Romberg, J. *l1-magic : Recovery of Sparse Signals via Convex Programming* 2005. <http://statweb.stanford.edu/~candes/l1magic/downloads/l1magic.pdf> (2016).
61. Sankaranarayanan, A. C., Studer, C. & Baraniuk, R. G. CS-MUVI: Video compressive sensing for spatial-multiplexing cameras in 2012 IEEE International Conference on Computational Photography (ICCP) (IEEE, 2012), 1–10. ISBN: 978-1-4673-1662-0. doi:10.1109/ICCPHOT.2012.6215212. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6215212>.
62. Li, Y., Sankaranarayanan, A. C., Xu, L., Baraniuk, R. & Kelly, K. F. Realization of hybrid compressive imaging strategies. *Journal of the Optical Society of America. A, Optics, image science, and vision* **31**, 1716–20. ISSN: 1520-8532 (2014).
63. Candès, E. J. & Recht, B. Exact Matrix Completion via Convex Optimization. *Foundations of Computational Mathematics* **9**, 717–772. ISSN: 1615-3375 (2009).
64. Cai, J.-F., Candes, E. J. & Shen, Z. A Singular Value Thresholding Algorithm for Matrix Completion, 1–26 (2008).
65. Candès, E. J. *et al.* Robust principal component analysis? *Journal of the ACM* **58**, 1–37. ISSN: 00045411 (2011).
66. Scotté, C. *et al.* Assessment of Compressive Raman versus Hyperspectral Raman for Microcalcification Chemical Imaging. *Analytical Chemistry* **90**, 7197–7203. ISSN: 0003-2700 (2018).
67. August, Y., Vachman, C., Rivenson, Y. & Stern, A. Compressive hyperspectral imaging by random separable projections in both the spatial and the spectral domains. *Applied Optics* **52**, D46. ISSN: 1559-128X (2013).
68. Yu, W.-K., Yao, X.-R., Liu, X.-F., Li, L.-Z. & Zhai, G.-J. Three-dimensional single-pixel compressive reflectivity imaging based on complementary modulation. *Applied Optics* **54**, 363–367. ISSN: 0003-6935 (2015).

69. Jauregui-Sánchez, Y., Clemente, P., Latorre-Carmona, P., Tajahuerce, E. & Lancis, J. Signal-to-noise ratio of single-pixel cameras based on photodiodes. *Applied Optics* **57**, B67. ISSN: 1559-128X (2018).
70. Podoleanu, A. G. Unbalanced versus balanced operation in an optical coherence tomography system. *Applied Optics* **39**, 173. ISSN: 0003-6935 (2000).
71. Vollmann, J., Profunser, D. M. & Dual, J. Sensitivity improvement of a pump-probe set-up for thin film and microstructure metrology. *Ultrasonics* **40**, 757–763. ISSN: 0041624X (2002).
72. Kuo, W.-C., Lai, C.-M., Huang, Y.-S., Chang, C.-Y. & Kuo, Y.-M. Balanced detection for spectral domain optical coherence tomography. *Optics Express* **21**, 19280. ISSN: 1094-4087 (2013).
73. Teleanu, E. L., Durán, V. & Torres-Company, V. Electro-optic dual-comb interferometer for high-speed vibrometry. *Optics Express* **25**, 16427. ISSN: 1094-4087 (2017).
74. Orth, A., Tomaszewski, M. J., Ghosh, R. N. & Schonbrun, E. Gigapixel multispectral microscopy. *Optica* **2**, 654. ISSN: 2334-2536 (2015).
75. Coulson, K. L. Effects of Reflection Properties of Natural Surfaces in Aerial Reconnaissance. *Applied Optics* **5**, 905. ISSN: 0003-6935 (1966).
76. Simmons, D. F., Fortgang, C. M. & Holtkamp, D. B. Using multispectral imaging to measure temperature profiles and emissivity of large thermionic dispenser cathodes. *Review of Scientific Instruments* **76**, 044901. ISSN: 0034-6748 (2005).
77. Hagen, N. & Kudenov, M. W. Review of snapshot spectral imaging technologies. *Optical Engineering* **52**, 090901. ISSN: 0091-3286 (2013).
78. August, Y. & Stern, A. Compressive sensing spectrometry based on liquid crystal devices. *Optics letters* **38**, 4996 (2013).
79. Amiot, C., Ryczkowski, P., Friberg, A. T., Dudley, J. M. & Genty, G. Supercontinuum spectral-domain ghost imaging. *Optics Letters* **43**, 5025. ISSN: 0146-9592 (2018).
80. Groner, W. *et al.* Orthogonal polarization spectral imaging: A new method for study of the microcirculation. *Nature Medicine* **5**, 1209–1212. ISSN: 1078-8956 (1999).
81. Zhao, Y., Zhang, L. & Pan, Q. Spectropolarimetric imaging for pathological analysis of skin. *Applied Optics* **48**, D236. ISSN: 0003-6935 (2009).
82. Shell, J. R. *Polarimetric Remote Sensing in the Visible to Near Infrared* PhD thesis (Rochester Institute of Technology, 2005). <https://scholarworks.rit.edu/theses/8197>.
83. Zhu, Z., Woodcock, C. E., Rogan, J. & Kellndorfer, J. Assessment of spectral, polarimetric, temporal, and spatial dimensions for urban and peri-urban land cover classification using Landsat and SAR data. *Remote Sensing of Environment* **117**, 72–82. ISSN: 00344257 (2012).
84. Rose, G., Rodgers, A., Burden, T. G. & Database, D. T Cells Cause Lung Damage in Emphysema. *PLoS Medicine* **1**, e25. ISSN: 1549-1676 (2004).
85. Doñate-Buendia, C. *et al.* Fabrication by Laser Irradiation in a Continuous Flow Jet of Carbon Quantum Dots for Fluorescence Imaging. *ACS Omega* **3**, 2735–2742. ISSN: 2470-1343 (2018).

86. Opilik, L., Schmid, T. & Zenobi, R. Modern Raman Imaging: Vibrational Spectroscopy on the Micrometer and Nanometer Scales. *Annual Review of Analytical Chemistry* **6**, 379–398. ISSN: 1936-1327 (2013).
87. Evans, C. L. & Xie, X. S. Coherent Anti-Stokes Raman Scattering Microscopy: Chemical Imaging for Biology and Medicine. *Annual Review of Analytical Chemistry* **1**, 883–909. ISSN: 1936-1327 (2008).
88. Berto, P., Scotté, C., Galland, F., Rigneault, H. & de Aguiar, H. B. Programmable single-pixel-based broadband stimulated Raman scattering. *Optics Letters* **42**, 1696. ISSN: 0146-9592 (2017).
89. Decker, J. a. Experimental Realization of the Multiplex Advantage with a Hadamard-Transform Spectrometer. *Applied Optics* **10**, 510. ISSN: 0003-6935 (1971).
90. Deverse, R. A., Hammaker, R. M. & Fateley, W. G. Realization of the Hadamard Multiplex Advantage Using a Programmable Optical Mask in a Dispersive Flat-Field Near-Infrared Spectrometer. *Applied Spectroscopy* **54**, 1751–1758. ISSN: 0003-7028 (2000).
91. Lin, H., Liao, C.-S., Wang, P., Kong, N. & Cheng, J.-X. Spectroscopic stimulated Raman scattering imaging of highly dynamic specimens through matrix completion. *Light: Science & Applications* **7**, 17179. ISSN: 2047-7538 (2018).
92. Soldevila, F., Dong, J., Tajahuerce, E., Gigan, S. & de Aguiar, H. B. Fast compressive chemical bio-imaging via matrix completion. *to be published* (2018).
93. Xu, Y., Yin, W., Wen, Z. & Zhang, Y. An alternating direction algorithm for matrix completion with nonnegative factors. *Frontiers of Mathematics in China* **7**, 365–384. ISSN: 1673-3452 (2012).
94. Clemente, P., Durán, V., Tajahuerce, E., Torres-Company, V. & Lancis, J. Single-pixel digital ghost holography. *Physical Review A* **86**, 041803. ISSN: 1050-2947 (2012).
95. Clemente, P. *et al.* Compressive holography with a single-pixel detector. *Optics Letters* **38**, 2524. ISSN: 0146-9592 (2013).
96. Horisaki, R., Matsui, H. & Tanida, J. Single-pixel compressive diffractive imaging with structured illumination. *Applied Optics* **56**, 4085. ISSN: 0003-6935 (2017).
97. Saita, Y., Shinto, H. & Nomura, T. Holographic Shack–Hartmann wavefront sensor based on the correlation peak displacement detection method for wavefront sensing with large dynamic range. *Optica* **2**, 411. ISSN: 2334-2536 (2015).
98. Pfund, J, Lindlein, N & Schwider, J. Dynamic range expansion of a Shack–Hartmann sensor by use of a modified unwrapping algorithm. *Optics Letters* **23**, 995. ISSN: 0146-9592 (1998).
99. Seifert, L., Liesener, J. & Tiziani, H. The adaptive Shack–Hartmann sensor. *Optics Communications* **216**, 313–319. ISSN: 00304018 (2003).
100. Martínez-Cuenca, R. *et al.* Reconfigurable Shack–Hartmann sensor without moving elements. *Optics Letters* **35**, 1338. ISSN: 0146-9592 (2010).
101. Zhang, Z., Wang, X., Zheng, G. & Zhong, J. Fast Fourier single-pixel imaging via binary illumination. *Scientific Reports* **7**, 12029. ISSN: 2045-2322 (2017).
102. Augustin, S., Szollmann, T., Jung, P. & Huebers, H.-W. Breaking Imaging Limits using Dithering Masks in 0.35 Terahertz Single-Pixel Imaging, 0–4 (2017).

103. Zhang, Y., Suo, J., Wang, Y. & Dai, Q. Doubling the pixel count limitation of single-pixel imaging via sinusoidal amplitude modulation. *Optics Express* **26**, 6929. ISSN: 1094-4087 (2018).
104. Conkey, D. B., Caravaca-Aguirre, A. M. & Piestun, R. High-speed scattering medium characterization with application to focusing light through turbid media. *Optics Express* **20**, 1733. ISSN: 1094-4087 (2012).
105. Goorden, S. A., Bertolotti, J. & Mosk, A. P. Superpixel-based spatial amplitude and phase modulation using a digital micromirror device. *Optics Express* **22**, 17999. ISSN: 1094-4087 (2014).
106. Drémeau, A. *et al.* Reference-less measurement of the transmission matrix of a highly scattering material using a DMD and phase retrieval techniques. *Optics Express* **23**, 11898. ISSN: 1094-4087 (2015).