INFORMATICA, 2023, Vol. 34, No. 3, 665–677

DOI: https://doi.org/10.15388/23-INFOR526

Causal Inference Applied to Explaining the Appearance of Shadow Phenomena in an Image

Jairo Ivan VÉLEZ BEDOYA[1,2], Manuel Andres GONZÁLEZ BEDIA[2], Luis Fernando CASTILLO OSSA[1,3], Jeferson ARANGO LÓPEZ[1], Fernando MOREIRA[4,*]

[1] *C. 65 No. 26-10 Universidad de Caldas, Manizales-Caldas, Colombia*
[2] *C. de Mariano Esquillor Gómez, s/n, 50018 Zaragoza, España Univesidad de Zaragoza, Zaragoza, Spain*
[3] *Carrera 27 No. 64-60 Univesidad Nacional de Colombia, sede Manizales, Manizales-Caldas, Colombia*
[4] *R. Dr. António Bernardino de Almeida 541, 4200-072 Porto, Portugal REMIT, IJP, Universidade Portucalense and IEETA, Porto, Portmugal*
e-mail: jairo.velez@ucaldas.edu.co, mgbedia@unizar.es, luis.castillo@ucaldas.edu.co, jeferson.arango@ucaldas.edu.co, afmoreira@upt.pt

Received: April 2023; accepted: August 2023

Abstract. Due to the complexity and lack of transparency of recent advances in artificial intelligence, Explainable AI (XAI) emerged as a solution to enable the development of causal image-based models. This study examines shadow detection across several fields, including computer vision and visual effects. Three-fold approaches were used to construct a diverse dataset, integrate structural causal models with shadow detection, and apply interventions simultaneously for detection and inferences. While confounding factors have only a minimal impact on cause identification, this study illustrates how shadow detection enhances understanding of both causal inference and confounding variables.

Key words: causality, causal inference, causal discovery, structural causal model, shadow detection, XAI.

1. Introduction

When considering any image, beyond seeing it as a container of objects, among other things, it is natural for a human being to give it meaning or to infer the explanation of some event of interest captured in it, but how can such an inference be reached through artificial intelligence? Causal inference can be applied in many areas of science and technology, such as economics, epidemiology, image processing, and autonomous driving, which are areas where it is crucial to make accurate decisions. Currently, there are widely studied methods that, through correlation, recognize and classify objects using datasets

*Corresponding author.

such as (Deng *et al.*, 2009) which has sufficient size and information to ensure high accuracy in such tasks (Zeiler and Fergus, 2014). However, in the last decade, as pointed out by Saeed and Omlin (2021), explainable AI (XAI) has been proposed to respond to the need raised by important contributions in artificial intelligence, which have led to an increasing complexity of algorithms and lack of transparency of models, and to advance the adoption of AI in critical domains. Then, to obtain the explanation we are looking for about an event captured in an image, we would have to consider causal relationships that can either be inferred through expert knowledge (Martin, 2018) or intervene such data sets through experimentation, as indicated in He and Geng (2008), taking into account that, in probabilistic language, not having a way to distinguish between giving value to a variable and observing it, prevents modelling cause and effect relationships (Perry, 2003). Thus, taking modelling as an essential step to achieve causal inference, Xin *et al.* (2022) discusses the role of causal inference to improve the interpretability and robustness of machine learning methods, and highlights opportunities in the development of machine learning models with causal capacity adapted for the analysis of mobility considering images or sequential data. In the punctual case on causal inference applied to images, Lopez-Paz *et al.* (2017) propose to use neural causality coefficients (NCCs) that are calculated by applying convolutional neural networks (CNNs) to the pixels of an image, so that the appearance of causality between variables suggests that there is a causal link between the real-world entities themselves, Lebeda *et al.* (2015) have proposed a statistical approach – transfer entropy – to discover and quantify the relationship between camera motion and the motion of a tracked object to predict the location of the tracked object, Fire and Zhu (2013) presented a Bayesian grammar (C-AOG) model for human-perceived causal relationships that can be learned from a video, and Pickup *et al.* (2014) use the causality method, supplemented with computer vision and machine learning techniques, to determine whether a video is playing forward or backward by observing the "arrow of time" in a temporal sequence.

This paper consists of four sections, in Section 1 we state the motivation of the study, present some antecedents that have made important contributions in the area of causal inference applied to images and define our contribution as a starting point to address a problem area already detected by several authors. In Section 2, we present the method used to generate the data, define the causal model and validate it with the NOTEARS algorithm, and then query the model by means of interventions. In Section 3, we analyse the results obtained in the applied causal discovery and causal inference processes and, finally, in Section 4, we conclude that the graphical representation of a causal model makes it simpler to understand the problem, although for the validation using NOTEARS we had to make restrictions based on expert knowledge. Likewise, we recognize the importance of the structure of the dataset for causal inference in contrast to the structure of a machine learning dataset and, finally, thanks to the interventions and queries of the causal model, it was possible to deduce, with a high level of certainty, the cause of the shadow projection.

1.1. *Related Work*

Regarding the explainability of events or phenomena captured in an image or video, taking modelling as an essential step to achieve causal inference, Xin *et al.* (2022) discusses the role of causal inference to improve the interpretability and robustness of machine learning methods, and highlights opportunities in the development of machine learning models with causal capability adapted for mobility analysis considering images or sequential data. In the punctual case on causal inference applied to image analysis, Lopez-Paz *et al.* (2017) propose to use neural causality coefficients (NCCs) that are calculated by applying convolutional neural networks (CNNs) to the pixels of an image, so that the appearance of causality between variables suggests that there is a causal link between the real-world entities themselves, Lebeda *et al.* (2015) have proposed a statistical approach – transfer entropy – to discover and quantify the relationship between camera motion and the motion of a tracked object to predict the location of the tracked object, Fire and Zhu (2013) presented a Bayesian grammar (C-AOG) model for human-perceived causal relationships that can be learned from a video, and Pickup *et al.* (2014) use the causality method, supplemented with computer vision and machine learning techniques, to determine whether a video is playing forward or backward by observing the "arrow of time" in a temporal sequence.

1.2. *Our Contribution*

By taking into account the underlying causes of shadow formation, causal inference can provide more accurate predictions and improve the overall realism of virtual environments. By virtue of this, given the relevance of this topic and the need for experimentation on specific cases that would potentially be contributory to evolving fields such as 3D graphics where shadow detection is an area where causal inference can be applied to improve accuracy and efficiency in this process, as opposed to traditional techniques such as ray tracing which is computationally expensive in terms of handling complex scenes with many objects (Levoy, 1990), this challenge is equally recognized in novel solutions such as that of Li *et al.* (2018) which computes derivatives of scalar functions on a rendered image with respect to arbitrary scene parameters such as camera location, scene geometry, materials and lighting parameters using an edge sampling algorithm. Our study provides a benchmark for addressing one aspect of this problem using causal inference to detect and deduce the cause of a shadow cast on the surface of a 3D scene.

2. Materials and Method

Our objective was to explain, by means of causal inference, the appearance of a shadow cast on the surface defined on the lower face of a 3D scene in which, in addition to being illuminated, a sphere is present. For this we established 3 steps; in the first one we generated 1×10^6 scenes (images) in which 4 features can be observed: light, sphere, surface and shadow. There are numerous libraries that facilitate both statistical and causal estimation and inference operations, therefore, we chose to use pgmpy (Ankan and Panda, 2015), causalinference (Wong, 2016) and causalnex (Beaumont *et al.*, 2021) for the development of the experiment, thus, in the second step we define a structural causal model (SCM) taking into account the relationships between these features, also identifying the independence of the features, and then, through Causal Discovery, structurally validate the model by applying the NOTEARS algorithm devloped by Zheng *et al.* (2018) which differs significantly from other search approaches in the discrete space of directed acyclic graphs (DAG). In the third, through conditional probability distributions (CPDs), it was possible to calculate the probabilities of each variation in the observations and thus validate the hypothesis that the appearance of the shadow is due to the presence of the sphere in an illuminated scenario with a defined surface, all by means of "What if..." queries to the model, as indicated in level 2 of the causal scale (DOING), which is reached when one can predict the effect(s) of deliberate alterations of the environment and choose among these alterations to produce a desired outcome (Pearl and Mackenzie, 2018). Then, we intervened the model by removing the sphere and looked at, for a 95% confidence interval, the average treatment effect (ATE) and p-value to have an estimate of how far our hypothesis was from being null. Finally, we compare the performance of our experiment with two conventional shadow detection techniques.

2.1. *Data*

As already stated, we considered 4 observable features for the confirmation of the dataset, Table 1 shows the label used for each of them and the values they can take, with 1 indicating that the feature is present in the scene and 0, otherwise.

Figure 1 shows the dataset based on a hypothetical scenario similar to that used by Pearl and Mackenzie (2018) to demonstrate the importance of probabilities emphasizing that a causal model involves more than drawing arrows, for behind these are probabilities.

Thus, by means of Algorithm 1, the data of 1×10^6 scenes were synthetically generated, of which all features are perceived in 99% (0), while in 70% of the remaining 1% illumination and sphere are observed, but neither surface nor shadow (1); in half of the images (15%) of the remaining 30% no feature can be observed (2), while, in the second half, only illumination and surface are observed, but neither sphere nor shadow (3).

Table 1
Labelling of the observed features.

Feature	Label	Posible values
Light	A	[1, 0]
Sphere	B	[1, 0]
Surface	C	[1, 0]
Shadow	Y	[1, 0]

	A	B	C	Y	Count
0	1	1	1	1	990000
1	1	1	0	0	7000
2	0	0	0	0	1500
3	1	0	1	0	1500

Fig. 1. Observations considered in the data set.

Algorithm 1: Data generation

$DS \leftarrow createDataset()$;
$DS.columns \leftarrow [{}'A', {}'B', {}'C', {}'Y']$;
$DS.size \leftarrow 1000000$;
$idxAllColumnsSample \leftarrow DS.sample(990000)$;
foreach *idx in idxAllColumnsSample* **do**
  foreach *col in DS.columns* **do**
    $DS[idx][col] \leftarrow 1$;
$Filter \leftarrow DS[A] = 0 \ \& \ DS[B] = 0 \ \& \ DS[C] = 0 \ \& \ DS[Y] = 0$;
$sublistLightSphere \leftarrow DS.sublist(Filter)$;
$idxLightSphereColumnsSample \leftarrow sublistLightSphere.sample(7000)$;
foreach *idx in idxLightSphereColumnsSample* **do**
  $Ds[idx][A] \leftarrow 1$;
  $Ds[idx][B] \leftarrow 1$;
$Filter \leftarrow DS[A] = 0 \ \& \ ds[B] = 0 \ \& \ ds[c] = 0 \ \& \ ds[y] = 0$;
$sublistLightSurface \leftarrow DS.sublist(Filter)$;
$idxLightSurfaceColumnsSample \leftarrow sublistLightSurface.sample(1500)$;
foreach *idx in idxLightSurfaceColumnsSample* **do**
  $DS[idx][A] \leftarrow 1$;
  $DS[idx][C] \leftarrow 1$;

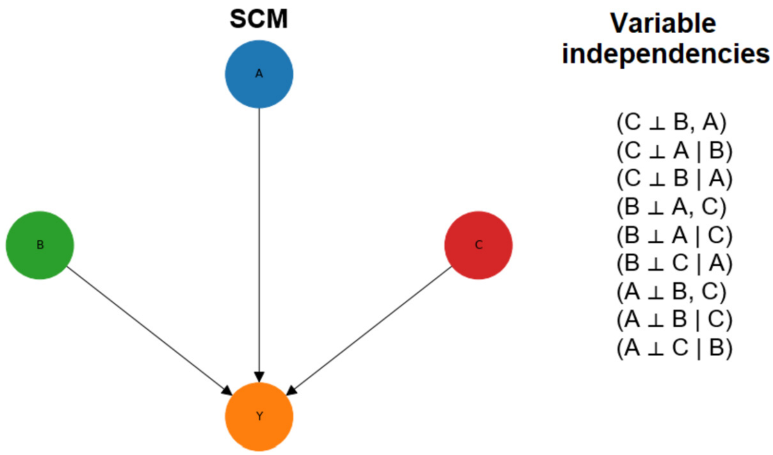

Fig. 2. SCM and variables independencies.

2.2. *Structured Causal Model SCM*

According to Pearl *et al.* (2019), an SCM is a way of describing the relevant features of the world and how they interact with each other. Specifically, a structural causal model describes how nature assigns values to variables of interest. Causal inference usually requires specialized knowledge and untestable hypotheses about the causal network linking treatment, outcome, and other variables. By summarizing knowledge and hypotheses in an intuitive way, graphs help clarify conceptual issues and improve communication between researchers (Hernán and Robins, 2020). Accordingly, we model an SCM in which each node corresponds to each observable feature and all edges point to a single collider node -Y-. Figure 2 shows the SCM and the independence of the variables that comprise it.

2.3. *Causal Inference*

Once the model is built, we calculate the conditional probability distributions (CPD) which are defined for a set of discrete and mutually dependent random variables to show the conditional probabilities of a single variable with respect to the others (Murphy, 2012). These are calculated by applying the chain rule as illustrated in (1), where X_i is an event and N is the number of events considered in the model, hence $0 < i < N$.

$$P(x_n) = P(x_1)P(x_2|x_1)P(x_3|x_1, x_2)\ldots P(x_n|x_1, x_2, x_3\ldots x_{(n-1)}). \tag{1}$$

As shown in Fig. 3, we calculate the probability of each possible value of each variable knowing the values taken by the other variables.

Then, to strengthen our hypothesis, we asked the model what would happen if no sphere had been detected, in other words, we intervened the model by not detecting the sphere in order to obtain the probability of detecting the shadow. To provide clarity on

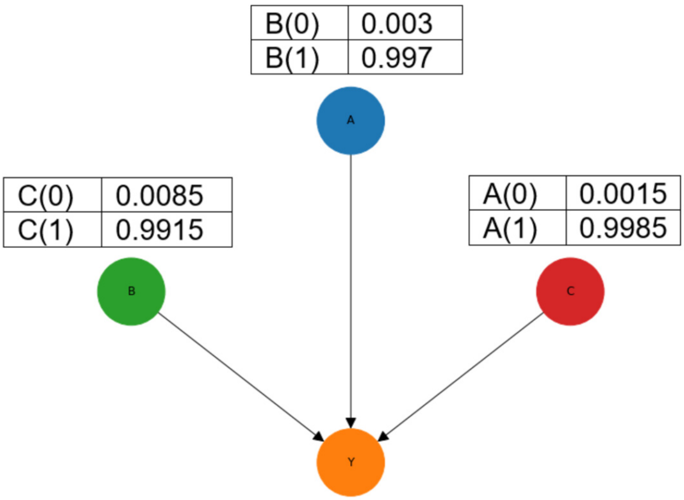

A	A(0)	A(0)	A(0)	A(0)	A(1)	A(1)	A(1)	A(1)
B	B(0)	B(0)	B(1)	B(1)	B(0)	B(0)	B(1)	B(1)
C	C(0)	C(1)	C(0)	C(1)	C(0)	C(1)	C(0)	C(1)
Y(0)	1.0	0.5	0.5	0.5	0.5	1.0	1.0	0.0
Y(1)	0.0	0.5	0.5	0.5	0.5	0.0	0.0	1.0

Fig. 3. CPD for each variable of the model.

Table 2
Role of SCM variables in the
Causal Inference process.

Label	Variable
Treatment	B
Confounders	[A, C]
Outcome	Y

what role the SCM variables play in the causal inference process we follow, among others, Chiappa and Isaac (2019), Guo *et al.* (2020) and Dague and Lahey (2019) who explain that a causal inference process is determined by a treatment, a set of confounders and an outcome. Table 2 shows the role of each MCS variable in the causal inference process.

Subsequently, considering this intervention, we calculated for the whole set of cases (N) the treatment effect to measure the average difference between the cases in which the treatment was applied (Y1) and the control cases (Y2) by applying (2).

$$ATE = \frac{1}{N}\sum_{i=1}^{N}\big(Y_1(i) - Y_0(i)\big). \tag{2}$$

Finally, to contrast the treatment results and thus obtain the estimate of how far the hypothesis was from being null, i.e. that there was no relationship between the sphere and the shadow, for a 95% confidence interval, we calculated a table of z-scores by applying (3) where X_i is each outcome, μ is the mean and σ is the standard deviation, and then calculated the bilateral tailed p-value and a significance level $\alpha = 0.05$.

$$Z = \frac{X_i - \mu}{\sigma}. \tag{3}$$

2.4. *Shadow Detection*

Within the realm of shadow detection in images, prominent methods include adaptive thresholding, threshold segmentation (Bradski, 2000), and clustering-based segmentation (Felzenszwalb) (Van der Walt *et al.*, 2014). We extended the latter by integrating a causal inference module. This advancement not only facilitated shadow detection but also enabled the identification of their root causes within specific scenarios. This comprehensive approach demanded parameters encompassing a comprehensive record of perceptible elements within the image (light, sphere, surface, and shadow), alongside the causal inference model.

Illustrating this process, Algorithm 2 outlines the intricate interplay between shadow detection and causal inference within the shadow phenomenon.

Algorithm 2: Algorithm for detection and causal inference of shadow phenomena

Data: Image, scale, sigma, min_size, light, sphere, shadow, surface, CI_model
Result: Segmented regions

Initialize an empty image segment map *segments*;
Initialize an empty priority queue pq for merging;
for *each pixel p in Image* **do**
    Create a new segment for p and add it to pq;
end
while *pq is not empty* **do**
    Merge the two smallest segments from pq;
    if *merge does not violate min_size constraint* **then**
        Add the merged segment to *segments*;
        Add the merged segment to pq;
    end
end
return CI_model.estimates(light, sphere, shadow, surface)

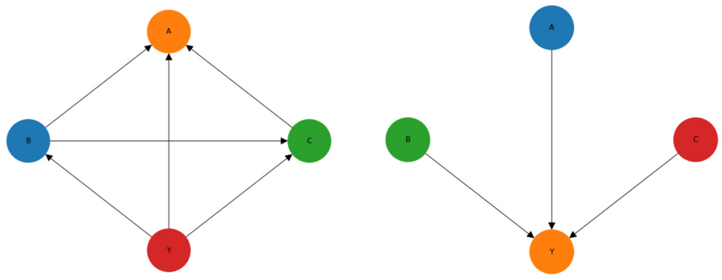

(a) Causal Discovery using NOTEARS without restrictions

(b) Causal Discovery using NOTEARS with restrictions

Fig. 4. NOTEARS causal discovery models with and without restrictions.

Table 3
Causal inference from intervention $P(Y \mid \mathrm{do}(B = 0))$.

Outcome	Probability	ATE	z	$P > \lvert z \rvert$	Confidence interval
Y(0)	0.995	0.993	11874	0.00001	95%
Y(1)	0.005				

3. Results

3.1. *Causal Model*

The structural causal model (SCM) was designed based on expert knowledge as Hernán and Robins (2020) pointed out, but validated in two attempts by means of causal Discovery using NOTEARS. In the first attempt, the algorithm took almost 5 minutes to generate the model shown in Fig. 4(a), which we consider quite long for the size of the dataset, resulting in a model that was not very coherent according to the expert knowledge. On the other hand, in the second attempt, we added a constraint to the algorithm to consider that A, B and C are independent as already shown in Fig. 2. The algorithm, as can be seen in Fig. 4(b), generated the model in less than 10 seconds and with the expected consistency.

3.2. *Causal Inference*

From the conditional probability distribution (CPD) it was possible to query the model under the hypothesis formulated. In Table 3, it can be seen that by eliminating the sphere there would be a 99.5% probability that no shadow would be cast; furthermore, it can be seen that the hypothesis gains strength by obtaining a p-value of less than 0.05 (the default threshold value) indicating that the null hypothesis is false, and a positive average treatment effect (ATE) value suggests that $P(Y|B) > P(Y)$, indicating that the presence of the sphere increases the probability of detecting a shadow cast on the surface or, in other words, it was possible to infer that the sphere is the most likely cause of the shadow cast on the surface.

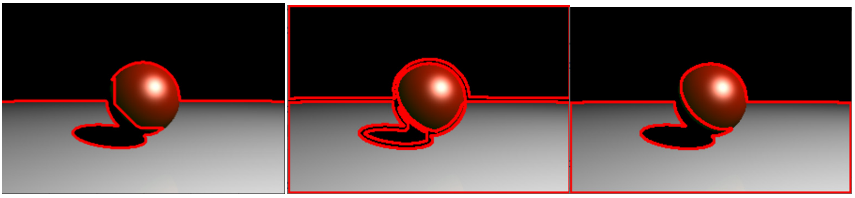

(a) Felzenszwalb with causal inference (b) Threshold segmentation (c) Adaptive threshold

Fig. 5. Shadow detection result.

3.3. *Shadow Detection and Causality*

To establish a contrast, we employed an identical image and introduced a confounding element by aligning the background colour with the shade's hue projected onto the surface. Subsequent execution encompassed the Felzenszwalb method integrated with the causal inference module, as well as the adaptive thresholding and threshold segmentation techniques. The outcomes of this comprehensive approach are visualized in Fig. 5.

In the context of shadow detection, the outcomes are evident. Among the approaches, the combination of the Felzenszwalb method with causal inference (refer to Fig. 5(a)) showcased the most promising results. It achieved an acceptable shadow detection accuracy of 51.5%. Following closely, the adaptive threshold method achieved an accuracy of 51.7% (refer to Fig. 5(c)), while the threshold segmentation method achieved 55.9% accuracy (refer to Fig. 5(b)).

It's important to emphasize that the presence of confounding factors significantly influenced the accuracy of the detection results. However, when considering the determination of the shadow's causality, the impact of confounding factors became negligible. Notably, only the Felzenszwalb method (refer to Fig. 5(a)) yielded a substantial result in this regard.

4. Conclusions and Future Work

We've shown how to employ causal inference to strengthen a hypothesis and, as a result, deduce the cause of a shadow phenomenon with high certainty. This is accomplished by utilizing interventions and inquiries within the causal model. We start with a set of photos from a 3D scenario in which four occurrences were examined as part of a structural causal model validated with the NOTEARS algorithm for causal detection. By contrasting their performance, we also demonstrated that adding a causal inference module to a shadow detection approach is feasible and advantageous. This opens the door for similar connections in other diverse and complex ways. A causal model's visual representation improves understanding of the problem and the roles that events play in its resolution. Despite testing the causal model with NOTEARS, there was some worry about the need to set limits based on expert knowledge. A dataset with a more intricate structure is required for causal inference when compared to typical datasets utilized for machine learning applications. Confounding factors had a considerable impact on the detection method's accuracy but not on the causal inference model. In the future, we hope to create a second version of this project. We intend to improve causal inference in this iteration by incorporating machine learning techniques. This combined approach will determine the origin of shadows sensed in complex graphical settings.

Data Availability Statement

The data presented in this study are available on request from the corresponding authors.

Acknowledgements

Jairo I. Velez: I thank the University of Zaragoza and Banco Santander (Universidad de Zaragoza-Santander Universidades) for providing me with one of the grants for Ibero-Americans in doctoral studies, thus making this research possible.

References

Ankan, A., Panda, A. (2015). pgmpy: probabilistic graphical models using python. In: *Proceedings of the 14th Python in Science Conference (SCIPY 2015)*. Citeseer.

Beaumont, P., Horsburgh, B., Pilgerstorfer, P., Droth, A., Oentaryo, R., Ler, S., Nguyen, H., Ferreira, G.A., Patel, Z., Leong, W. (2021). *CausalNex*. https://github.com/quantumblacklabs/causalnex.

Bradski, G. (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 120, 122–125.

Chiappa, S., Isaac, W.S. (2019). A causal Bayesian networks viewpoint on fairness. In: Kosta, E., Pierson, J., Slamanig, D., Fischer-Hübner, S., Krenn, S. (Eds.), *Privacy and Identity Management. Fairness, Accountability, and Transparency in the Age of Big Data. Privacy and Identity 2018*, *IFIP Advances in Information and Communication Technology*, Vol. 547. Springer, Cham. https://doi.org/10.1007/978-3-030-16744-8_1.

Dague, L., Lahey, J.N. (2019). Causal inference methods: lessons from applied microeconomics. *Journal of Public Administration Research and Theory*, 29(3), 511–529. https://doi.org/10.1093/jopart/muy067.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L. (2009). ImageNet: a large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. https://doi.org/10.1109/CVPR.2009.5206848.

Fire, A., Zhu, S.-C. (2013). Using causal induction in humans to learn and infer causality from video. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 35, 2297–2302. https://escholarship.org/uc/item/4ng247kx.

Guo, R., Cheng, L., Li, J., Hahn, P.R., Liu, H. (2020). A survey of learning causality with data: problems and methods. *ACM Computing Surveys*, 53(4), 1–37. https://doi.org/10.1145/3397269.

He, Y.B., Geng, Z. (2008). Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research*, 9, 2523–2547.

Hernán, M.A., Robins, J.M. (2020). *Causal Inference: What If*, I ed. Chapman & Hall/CRC, Boca Raton, pp. 312. 978-1420076165. https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/.

Lebeda, K., Hadfield, S., Bowden, R. (2015). Exploring causal relationships in visual object tracking. In: *2015 IEEE International Conference on Computer Vision (ICCV)*, Santiago, Chile, 2015, pp. 3065–3073. https://doi.org/10.1109/ICCV.2015.351.

Levoy, M. (1990). Efficient ray tracing of volume data. *ACM Transactions on Graphics (TOG)*, 9(3), 245–261. https://doi.org/10.1145/78964.78965.

Li, T.-M., Aittala, M., Durand, F., Lehtinen, J. (2018). Differentiable Monte Carlo ray tracing through edge sampling. *ACM Transactions on Graphics*, 37(6). 1–11. https://doi.org/10.1145/3272127.3275109.

Lopez-Paz, D., Nishihara, R., Chintala, S., Sch¨lkopf, B., Bottou, L. (2017). Discovering causal signals in images. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, USA, 2017, pp. 58–66. https://doi.org/10.1109/CVPR.2017.14.

Martin, O. (2018). *Bayesian Analysis with Python*, 2nd ed. Packt Publishing, Birmingham, Mumbai. 978-1-78588-380-4.

Murphy, K.P. (2012). *Machine Learning: A Probabilistic Perspective*, Vol. 16. The MIT Press, Massachusetts, pp. 1–1098. 978-0-262-01802-9.

Pearl, J., Mackenzie, D. (2018). *The Book of Why*, I ed. Basic Books, New York. 0465097618.

Pearl, J., Glymour, M., Jewell, N.P. (2019). *Causal Inference in Statistics A Primer*. John Wiley & Sons Ltd, West Sussex, pp. 159. 9781119186847.

Perry, M.J. (2003). Causality: models, reasoning, and inference. In: *The Encyclopedia of Ancient History*, Vol. 19, pp. 675–685. https://doi.org/10.1002/9781444338386.wbeah13034.

Pickup, L.C., Pan, Z., Wei, D., Shih, Y., Zhang, C., Zisserman, A., Scholkopf, B., Freeman, W.T. (2014). Seeing the arrow of time. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Columbus, OH, USA, 2014, pp. 2043–2050. https://doi.org/10.1109/CVPR.2014.262.

Saeed, W., Omlin, C. (2021). Explainable AI (XAI): a systematic meta-survey of current challenges and future opportunities. *Knowledge-Based Systems*, 263, 110273. https://doi.org/10.1016/j.knosys.2023.110273.

van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., Yu, T. (2014). scikit-image: image processing in Python. *PeerJ*, 2, e453. https://doi.org/10.7717/peerj.453.

Wong, L. (2016). *Causalinference*. https://causalinferenceinpython.org/.

Xin, Y., Tagasovska, N., Perez-Cruz, F., Raubal, M. (2022). Vision paper: causal inference for interpretable and robust machine learning in mobility analysis. In: *SIGSPATIAL'22: Proceedings of the 30th International Conference on Advances in Geographic Information Systems*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3557915.3561473.

Zeiler, M.D., Fergus, R. (2014). Visualizing and understanding convolutional networks. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (Eds.), *Computer Vision – ECCV 2014*, *Lecture Notes in Computer Science*, Vol. 8689. Springer, Cham, pp. 818–833. https://doi.org/10.1007/978-3-319-10590-1_53.

Zheng, X., Aragam, B., Ravikumar, P., Xing, E.P. (2018). Dags with NO TEARS: continuous optimization for structure learning. In: *NIPS'18: Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018 December, pp. 9472–9483.

J.I. Vélez Bedoya has a master's degree in software project management and development (2011) from the Autónoma University of Manizales in Colombia. He is currently a PhD candidate in systems engineering and informatics at the University of Zaragoza in Spain. He belongs to the Systems and Informatics Department at the University of Caldas, where he is a full time professor since 2017.

M.A. González Bedia is a postdoctoral researcher from University of Sussex (2009), he has a PhD in computer science from the University of Salamanca (Spain). He is a full time professor at the Department of Computer Science and Systems Engineering of the University of Zaragoza in Spain (since 2008) and also he is a researcher at the Institute of Engineering Research of Aragon (I3A) (since 2012), from 2019 he has been Science and Universities Advisor for the Ministry of Science, Innovation and Universities of the Spanish government, Deputy Director General of University Research Activity Ministry of Universities of the Spanish government (since 2020) and Director General, Commissioner for the New Language Economy at the Ministry of Economic Affairs and Digital Transformation of the Spanish government (since 2022).

L.F. Castillo Ossa has a PhD in computer science and automation from University of Salamanca (Spain) and was the dean of the Faculty of Engineering at the University of Caldas (Colombia) from 2014 until 2018 where he also works as professor in the Department of Systems and Informatics. He is the head of the Research Group Artificial Intelligence, senior researcher in MinCiencias. His main lines of research are related to multiagent systems and artificial cognitive systems. He works part-time as a professor in the National University of Colombia, Campus Manizales, professor of doctorate in Cognitive Science Autonomous University, Manizales. He belongs to the Systems and Informatics Department at the University of Caldas where he is a full time professor since 2011.

J. Arango López received his MSc in computational engineering from the University of Caldas in Colombia. He obtained his PhD in electronic science from the University of Cauca in Colombia and PhD in communication and information technologies from the University of Granada in Spain (2019). He is currently full-time professor at the University of Caldas (Colombia) and his research interests are pervasive games, semantic web, artificial intelligence and linked open data.

F. Moreira has a MSc in electronic engineering (1997) and PhD in electronic engineering (2003), both from the Faculty of Engineering of the University of Porto and Habilitation (2018). He has been a member of the Science and Technology Department at Portucalense University since 1992, currently as a full-time professor, and a visiting professor at the University of Porto Business School.