

Detecting Availability to Facilitate Social Communication

Nils Einecke, Thomas Weisswange

2022

Preprint:

This is an accepted article published in ICRA Workshop: Exploring the Roles of Robots for Embodied Mediation. The final authenticated version is available online at: <https://doi.org/No DOI available for this workshop>

Detecting Availability to Facilitate Social Communication

Nils Einecke

Honda Research Institute Europe GmbH
Offenbach, Germany
nils.einecke@honda-ri.de

Thomas H. Weisswange

Honda Research Institute Europe GmbH
Offenbach, Germany
thomas.weisswange@honda-ri.de

Abstract—Modern communication has shifted strongly towards the digital domain. Technologies like chat applications and messengers provide means to connect with anyone at anytime. However, being always reachable can create stress and provides potential for social pressure due to expected timely responses. In contrast, the classical communication types via telephone is based on dedicated time slots for an ongoing communication, but suffers from the need to schedule or find a mutual opportunity to initiate this connection. Many Messenger applications allow the setting of a status to provide information about inappropriate times for initiating a communication, but constant updating of these states is often not feasible. We propose a framework for automatic detection of communication-availability and for signaling this to dedicated interaction partners using a video-mediated communication embodiment. In contrast to typical automatic status setting, our approach does not require a person to use or be close to its device, as it is not relying on device usage but rather on recognizing everyday activities and situations. Our algorithmic framework analyses situations using mainly image and audio features that are perceived through the communication device and indicates opportunities to start a communication to a possible partner. Within this framework, we compare a number of different machine learning techniques for classifying availability in a simple, real-world test dataset. The best approach is able to predict availability with 83% accuracy and presence of communication partners with an accuracy of 95%.

Index Terms—availability recognition, presence detection, video-mediated communication,

I. INTRODUCTION

Communication is one of the most important activities for social participation of human beings. However, increased relocation flexibility and spatial fragmentation of peer groups in modern societies required the development of technical means to mediate communication. Telephone calls and, more recently, video-mediated communication applications provide such means to connect with others independent of spatial closeness. These technologies allow communication based on joint appointments for a fixed time or by one side requesting a connection without knowing about the current interest and availability of the partner. Tee et al. [20] interviewed remote families on their communication behavior and found that many people would like to communicate more but feel like they need a concrete reason to initiate a communication and they try to take care to only share “relevant” information. In addition, timing of a communication request is an important issue as receiving requests at inopportune moments can be

very disruptive and the initiating party often considers this in their decision making [1], [7]. Several studies [5], [7], [14] found both callers and callees think that callers should take the callee’s situation into account when deciding whether or not to call. The advent of smart phones and general availability of internet connections has enabled more ubiquitous communication opportunities. The transition to using messenger and chat applications provided opportunities to reduce the importance of information relevance through easing the cost of communication, e.g. with respect to required time. Yet, recent studies have shown that being always reachable creates stress [14], [18]. Certain features like the *last seen* or *online* status of messengers additionally creates social pressure of reading and answering messages in a timely manner [5].

Beyond these “common” communication means, there is a large number of research on how to mitigate social connections or “connectedness” between distant partners. Hassenzahl et al. [12] provide an extensive overview of different approaches and highlight the importance of interactions beyond exchanging information. One example is the Messaging Kettle [4], which allowed sharing abstract and textual messages around a joint activity to overcome physical and social distances. Many similar proposals in the community provide means to enhance sharing between remote people without the need for explicit communication, often called “ambient awareness” [10].

Aiming to combine the positive effects of these different communication means, we want to work towards technologies that can mediate communication opportunities between people. Prior work provided ideas and products which allowed communication partners to manually set their current status for communication and further advancements that tried to infer such status automatically based on patterns of computer usage [18]. The method proposed here, will go beyond using data from direct interaction with the device used for communication towards considering regular activities of daily life. Additionally, the context information will help providing explicit opportunities for communication instead of an implicit state display. We will use the opportunities provided by the specific embodiment of semi-stationary video-mediated communication devices like Komp¹, ViewClix² or Facebook

¹www.noisolation.com/komp

²www.viewclix.com

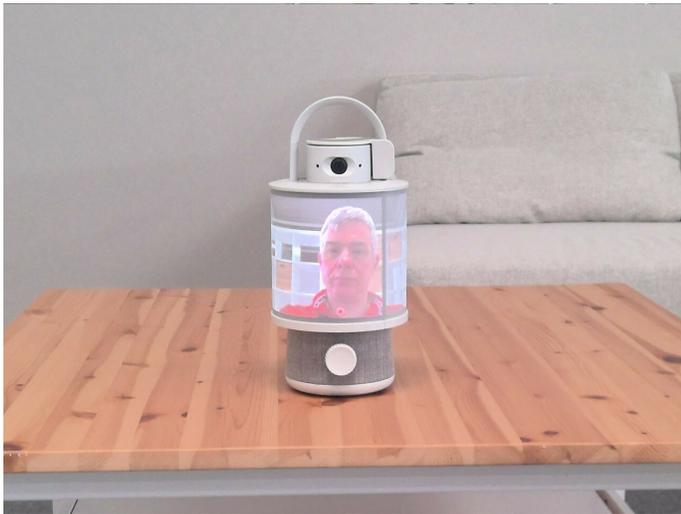


Fig. 1. Prototype semi-stationary video-mediated communication device Clara.

Portal³. These devices have a reduced functionality (with respect to a smartphone or tablet) by concentrating on the communication aspect and allow to be set up with only a small number of communication peers. At the same time, they allow insights into more complex and everyday user states through their positioning outside the common 'computer rooms' and through containing a camera and microphone. We use a prototype communication device called Clara with a lantern-like shape, which is depicted in Fig. 1. The lantern shape supports the semi-stationary nature with a carrying handle in a form-follows-function principle. A wide-angle camera, multiple microphones, and a curved display allow to move around the device without disturbing the tele-presence impression. Based on these sensors, we develop a framework that can detect situations where the user would most likely be interested in accepting a communication request – we will refer to this as communication availability. This state will then be communicated to devices at selected remote users, who will receive an ambient signal that allows them to use that opportunity or ignore it dependent on their own situation. We assume that such a passive communication initiation will leverage the social implications that are reported for classical communication means like phone calls or messenger services [19].

In the remainder of the paper, we will discuss prior work on communication mediation, in particular from the point of view of detecting communication availability or status. We will then propose a first proof-of-concept framework for automatic detection of a user's availability using Clara and evaluate multiple algorithms for their classification quality on a first set of test recordings. We will close with discussing future steps to develop the proposed method further.

³portal.facebook.com

II. RELATED WORK

Communication apps like WhatsApp, Discord or Teams provide only simple means for sharing contextual information. In most apps, the user status is restricted to a few alternatives like "online", "offline", "busy" or "away". In a work by Cho et al. [5], the predominant messenger app in south Korea, Kakao, was extended to enable a more detailed setting of status information. This included states like "having lunch", "showering", "in the lab" or "reading a book." Results of a study showed a significant release in stress but also a concern for privacy violations. Most participants liked the detailed setting but would want to share this state only with family or very close friends. The most frequently used group described the unavailability for communication. However setting a state manually was considered very tedious, and many participants did not frequently update or use the more detailed status types.

There are several approaches in the literature [2], [9], [13], [15], [17] that demonstrate that simple cues can be used to predict unavailability with a high accuracy. In most cases, the user activity with the device and speech were highly relevant cues. For example, Pielot et al. [17] used information like *time since last call* and *device posture*. They found that using only mobile phone internal sensor data is enough to predict a call engagement with an accuracy of 83.2%. However, just predicting the acceptance of a call neglects the social implications that might lead people to answer a call even though they would rather like not to be disturbed. The authors in [2] found that users might initiate communication also when they are likely to disturb the peer but would then keep the interaction rather short. To evaluate salient features for initiating face-to-face communication, they introduced the "Lilsys" system, a dedicated sensor array which included speech, calendar, door, phone and motion sensors. Lilsys used a simple decision tree based on correlations with the target state from a prior *wizard-of-oz* study [13]. One important finding was that detecting speech is a very relevant cue for signaling unavailability. The authors also found that some users did not want to reveal their status in detail, which provided an argument for an agglomerated score instead of detailed status information. Like all prior work, Lilsys focuses on signaling unavailability instead of proposing communication opportunities.

A recent approach by Chou [6] extracted context factors from simple sensory data. These context factors included *social awareness*, *environment awareness* and *location awareness*. The authors employed machine learning for estimating if a callee is willing to interrupt its current activity to pick up a call. Using all context factors they were able to predict potential interruptions with 85% accuracy. They showed that time and location are a very important cues but also that the social relation to the caller played a role, which emphasizes the importance of considering dedicated context sharing.

In this work, we evaluate the potential embodied mediation of suggesting communication opportunities instead of providing passive information about when to not call. We assume that the embodied device for mediation has a semi-

stationary character. This means, we assume the communication device to be rather put at a central place in the living environment instead of being carried around all the time. Hence, we concentrate on visual and auditory features for situation evaluation in everyday life and for suggesting of communication opportunities based upon this evaluation. In particular, we do not use device-specific activity as the semi-stationary devices are rarely operated besides communication events.

III. AVAILABILITY DETECTION FRAMEWORK

In order to infer the availability of a person, i.e. the willingness to engage in communication, we mainly rely on sensors available on the Clara device. These are the camera and microphone, which are an integral part of every video-mediated communication device. The two sensors are rich in information, can capture a wide variety of every-day activities and data is available as long as the device is running. Another potentially interesting set of information sources would be environment sensors like thermometer, door sensors or power meter. Although such sensors are also rich in information their concrete setting and availability vary strongly between different households and device and would therefore limit the generalization of the approach. Hence, we designed our framework to be able to include these as additional means but focus on video and audio data in the current paper.

It is possible to derive a large number of different indicators that could allow to infer the availability, and it is therefore vital to understand what availability means. Similar to [2], we consider different aspects. First, there is physical availability or presence. If a person is not located in the vicinity of the communication device, it will not be reachable, even in urgent cases. Second, there is social availability. If a person is already engaged in a social interaction it is usually regarded impolite to start a second communication activity like a phone call. Third, there is the mental availability. If a person needs time for itself or is not interested to have a communication, it will likely not be interested in receiving a call. The detection of these three aspects can be considered to increase in difficulty. In this paper, we will focus on the first two aspects. While our framework can also include mental availability, this requires more complex modeling, for example through first inferring specific activities from the sensory indicators and then use a personalized mapping to relate those to mental states of the user. However, we will provide some ideas for the latter in the discussion.

Physical and social availability correlate directly and indirectly with many environment states that can be measured to a high accuracy. We construct indicator functions for such states based on existing approaches, for example, to detect if a room is lit or if there is a person close to the camera. In a second step, we use the output of these indicators to learn a mapping to the availability of a person. Fig. 2 depicts the general structure of our availability detection framework. On the left side are the potential input sources. Camera and microphone sensors are available on the embodied mediation

device. The environment and wearables might provide additional data but are not available in all settings. From the data a set of indicators is extracted. The indicators have a wide range of complexity depending on the aspect of availability they indicate. Typically, indicators for physical availability are rather simple while those for social availability are more complex.

We distinguish five main types of indicators for availability recognition in embodied mediation devices:

- **Visual:** lights on/off, background movement, person detected, person recognized, location recognition
- **Auditory:** silence, voice detection, voice recognition, ambient sound detection
- **Activities:** discussion ongoing, cooking, reading a book, watching TV, cleaning, meditating
- **Technical:** mediation device reachable, local time, user smartphone within same wireless network, current energy consumption of household, door or window open/closed, household devices switched on/off
- **User Settings:** (un)available times, (un)interruptible activities

In the work presented here, we use a set of seven indicators: *LightsOn*, *MovementSeen*, *FaceSeen*, *SoundHeard*, *SpeechHeard*, *MusicHeard* and *Time*.

All indicators compute a simple score and yield three output elements. The first is a Boolean about whether this indicator is active or not. This is derived from empirical thresholds applied to the internally computed score. The second output is a confidence indicator for the Boolean value, which depends on the distance of the decision threshold to the score value. Finally, the indicators also output their internally computed score.

LightsOn is a simple descriptor that uses the camera image mean intensity as a score value in order to derive whether a room is lit by light (including sunlight) or not. The idea behind this cue is that dark or faintly lit rooms probably are not used and hence nobody is nearby that could engage in a communication.

MovementSeen uses the Gaussian mixture model background subtraction [21] implemented in opencv [3] for detecting if something has moved in the scene. Images from the camera are fed into the algorithm at a rate of 1Hz to detect small movement while being somewhat robust to noise. The score of this indicator is the number of foreground pixels.

FaceSeen employs the deep neural network version of face detection implemented in opencv [3]. We are using this indicator because if someone is looking towards the device it can be considered more likely the (s)he is willing to engage in communication. The number of detected faces is used as the score. Multiple detected faces can give some indication about ongoing social activities which indicate lowered availability.

SoundHeard is a simple auditory indicator for detecting the presence of people. Typically, household environments are very silent if nobody is nearby. Using a simple score composed of the mean amplitude as well as the maximum amplitude for

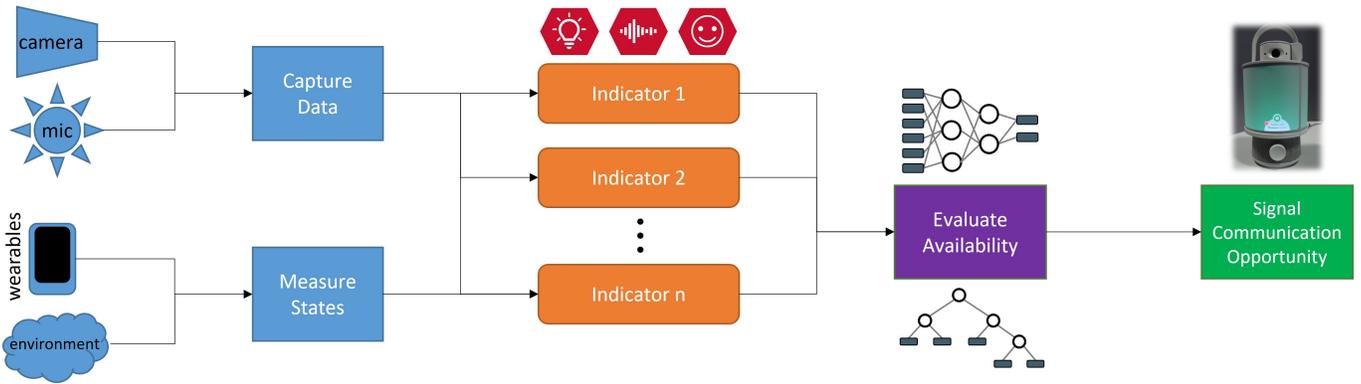


Fig. 2. Availability detection framework.

a time range of one second, *SoundHeard* constitutes a cheap to compute yet powerful feature.

The *SpeechHeard* and *MusicHeard* indicators analyze the sound data from the microphone to detect if somebody is speaking or music is played. We use the *ina* speech segmenter [8] to do this analysis. The two indicators are more elaborate versions of the *SoundHeard* indicator as they provide additional auditory context. In contrast to the other indicators, these two do not have a meaningful simple internal score.

Finally, we employ the technical indicator *Time* which simply reports the current weekday encoded as integer [0..6] and time of day based on the system clock. This indicator has no meaningful Boolean and confidence and thus always outputs true and 1.0, respectively

The output of the indicators is used to estimate the user’s current availability. A common approach for communication devices is to apply an empirical set of rules for deriving user availability. In contrast, machine-learning approaches are clearly advantageous when it comes to complex situation inference and to user-adapted inference. It is also to be expected that machine-learning-based inference achieves a much higher accuracy. In order to confirm this, we performed a proof-of-concept evaluation of our framework using standard machine-learning techniques for the availability inference.

In a final step, we want to use the information about availability to support the initiation of a communication. In a first implementation, we use the feature of the Clara device to set a primary contact to determine a single communication partner to receive this information. In future work, we plan to feed the system with information of all users to only recommend communication to pairs with shared availability. We adapted the user interface of Clara to introduce an explicit communication recommendation state. The detection of availability will trigger setting up a connection to the corresponding Clara at the partner’s location. For a classical phone or video call system the other Clara would signal an incoming communication request, which requires the user to actively respond to it (accept, reject, ignore). Our system instead only triggers a visual signal, a light green glow with slow cyclic brightness modulation, and a message ”Start a

Conversation with <peer>”. This is depicted on Fig. 2 on the far right side. Pressing the button on the device will then open the communication channel as in a classical interface, but our aim is that the availability state will only be noticed if the remote user was actively interested in a communication and looking at the device.

IV. EXPERIMENTS

The target of our experiments is to confirm our framework in a proof-of-concept fashion and to find the most suitable machine-learning approach for inferring availability from the indicator data. In our experimental setup, we use the embodied communication prototype Clara depicted in Fig. 1. It has carrying lantern like shape, which underlines its semi-stationary nature. This prototype is equipped with a 180° 8MP camera, a stereo microphone, a 320° display and a simple turn-push knob. For evaluating our availability detection framework, we used the camera and stereo microphone to capture image and sound data that is used to compute the indications described in section III.

As a baseline, we use an inference that always outputs true for availability and a simple hand-crafted manual inference. The pseudocode of the manual inference is shown at Algorithm 1. We apply two empirical thresholds 0.1 and 0.5 for extracting a final presence and availability state Booleans, respectively.

For gathering training data, we asked participants to put the Clara device (Fig. 1) next to their home-office work place. In this setting, we recorded indicator data (Boolean, confidence, score) at a rate of 1 Hz. Altogether, our training data sums up to 27 days of indicator data with roughly 0.9 million data points. All participants were asked to write a dairy and note down times when they were close to the device and times when they would have been willing to engage in a communication. The dairy entries were transformed in simple machine-readable csv format containing start and end times as well as Booleans for presence and availability.

In order to find the most suitable machine learning approach for availability inference, we compared a set of five standard binary classifiers from the sklearn toolbox [16]: Logistic Regression (LR), Random Forest (RF), Multi-Layer Perceptron

Algorithm 1 Manual Availability Inference

```
indicator_set  $\leftarrow$  {LightsOn, MovementSeen, FaceSeen,
SoundHeard, SpeechHeard, MusicHeard}
procedure AVAILABILITY(indicator_set)
  av  $\leftarrow$  0.5
  if  $\neg$  LightOn then
    av  $\leftarrow$  0.1
    return av
  end if
  if FaceSeen then
    av  $\leftarrow$  av + 0.3
  end if
  if SpeechHeard then
    av  $\leftarrow$  av - 0.5
  end if
  for indicator  $\in$  {MovementSeen, SoundHeard, Music-
Heard} do
    if indicator then
      av  $\leftarrow$  av + 0.1
    end if
  end for
  CLIP(av, 0, 1)
  return av
end procedure
```

(MLP), Support Vector Machine (SVM) and Naive Bayes (NB). All classifiers were run with the default parameter setting except for an increased iteration number for LR and SVM. Furthermore, all classifiers were trained separately for the prediction of the *available* and the *present* state. Tables I and II show the results for *present* and *available*, respectively.

The results demonstrate that presence in front of the device can be detected with an accuracy of 95% and availability can be detected with an accuracy of 83%. The classifiers LR, RF and SVM show similar performance while MLP and NB are clearly outperformed. With some minor margin RF is superior to the other classifiers in this setting. However, the model files of LR and SVM (a few KB) are substantially smaller than for RF (hundreds of MB), which lends itself better for sharing and managing of multiple user-specific models. Hence, we conclude that LR, RF and SVM are all suitable availability predictors, which can be selected depending on other constraints.

Another interesting insight is of course which indicator contributes to what extent to the performance. To analyze this, we extracted the RF feature weightings. Table III shows the ten highest weighted features. Interestingly, *LightsOn*, *SoundHeard* and *MovementSeen* score have the highest weights. This makes sense, as these are good indicators that someone is close by. On the other hand, face detection seems to be much less relevant. The problem arises from the narrow orientation range at which faces can be detected. Only if a person is directly facing the device, this will be registered. Another important feature is time. This is also reasonable since there are regular time patterns like having a lunch. Last but not least,

TABLE I

CLASSIFIER RESULTS FOR DETECTING PRESENCE OF A PERSON NEAR THE COMMUNICATION DEVICE. AS A BASELINE THE RESULT FOR A MANUALLY IMPLEMENTED RULE-BASED INFERENCE AND AN INFERENCE RETURNING ALWAYS *true* IS SHOWN.

algorithm	accuracy	f1	precision	recall
LR	0.95	0.97	0.95	0.99
RF	0.95	0.97	0.96	0.98
MLP	0.91	0.94	0.93	0.96
SVM	0.94	0.96	0.95	0.98
NB	0.92	0.95	0.95	0.95
Manual	0.88	0.93	0.88	0.99
True	0.82	0.90	0.82	1.00

TABLE II

CLASSIFIER RESULTS FOR DETECTING WILLINGNESS AND AVAILABILITY OF A PERSON TO ENGAGE IN A COMMUNICATION. AS A BASELINE THE RESULT FOR A MANUALLY IMPLEMENTED RULE-BASED INFERENCE AND AN INFERENCE RETURNING ALWAYS *true* IS SHOWN.

algorithm	accuracy	f1	precision	recall
LR	0.81	0.82	0.79	0.87
RF	0.83	0.84	0.81	0.89
MLP	0.72	0.72	0.75	0.79
SVM	0.81	0.82	0.79	0.88
NB	0.68	0.76	0.64	0.98
Manual	0.80	0.82	0.77	0.88
True	0.57	0.71	0.57	1.00

SpeechHeard seems to be a good indication for availability. As already pointed out earlier, a person is likely to be not wanting to engage in a communication when it already talks to someone. Altogether, this analysis shows that among the set of features the simple ones are ranking highest, i.e. they contribute significantly to the performance.

V. CONCLUSION

In this work, we presented a framework for mediation of social communication initiation through detecting availability on a dedicated embodiment. Our approach is based on a wide range of basic context indicators which are combined into an agglomerated availability estimation using machine learning. In a first set of experiments, we identified the best performing approaches, which allow our system to detect physical and social availability with an accuracy of 83%. We provided a basic user interface to allow communication of opportunities to connect in contrast to many standard approaches that focus on signalling unavailability. With a basic system for availability detection, we can next start investigating the user acceptance of this approach.

The system implementation presented here is supposed to demonstrate the feasibility of the concept. In future work, we want to incorporate more indicators, which could also address aspects of mental availability, especially *activity* indicators that are able to detect typical household activities like "watching TV" or "doing laundry". There are many existing approaches for vision based activity recognition, that could be used for this [11]. Another planned component will consider data from two connected embodiments and improve communication recommendations using shared availability.

TABLE III
FEATURE RANKING.

rank	features	weight
1	LightsOn score	0.149
2	SoundHeard score (mean)	0.130
3	MovementSeen score	0.119
4	SoundHeard score (peak)	0.089
5	Time (hour)	0.084
6	SpeechHeard Boolean	0.083
7	Time (minute)	0.075
8	SoundHeard confidence	0.070
9	Time (weekday)	0.048
10	LightsOn confidence	0.038

We also plan to extend the experimental data to a larger set of users and move to data from other rooms, as the homeoffice setup used here has some obvious limitations. When doing this, it would also be interesting to investigate the impact of personalized adaptation of the availability detection. For this, one would need to setup a (semi-) online learning framework, to regularly collect data over some time and retrain the system on this data. This could use a direct means for a user to provide partial ground truth or alternatively the success rate of recommended communication intentions could be used to label such data. As far as our experiments showed there seems to not be a performance issue with frequent retraining. We would even recommend to apply retraining to the three best classifiers (LR, RF and SVM) each time and select the best. Combining such personalized availability detectors of multiple persons could then even produce targeted availability states that differ based on the proposed communication recipient.

Based on the proof-of-concept system described here and future plans on advancing implementations of the proposed concept, we hope to be able to improve communication initiation mediation between remote people beyond what is currently possible.

ACKNOWLEDGMENT

We thank Joel Schwartz, Gustavo Huber and Nick Mead for supporting the development of the telepresence lantern and Mark Dunn for technical and software support.

REFERENCES

- [1] D. Avrahami, D. Gergle, S. E. Hudson, and S. Kiesler. Improving the match between callers and receivers: A study on the effect of contextual information on cell phone interruptions. *Behaviour & Information Technology*, 26(3):247–259, May 2007.
- [2] J. Begole, N. E. Matsakis, and J. C. Tang. Lilsys: Sensing unavailability. In J. D. Herbsleb and G. M. Olson, editors, *CSCW*, pages 511–514. ACM, 2004.
- [3] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.
- [4] M. Brereton, A. Soro, K. Vaisutis, and P. Roe. The Messaging Kettle: Prototyping Connection over a Distance between Adult Children and Older Parents. In *33rd ACM Conference on Human Factors in Computing Systems (CHI 2015)*, pages 713–716, Seoul, Republic of Korea, 2015. ACM Press.

- [5] H. Cho, J. Oh, J. Kim, and S.-J. Lee. I share, you care: Private status sharing and sender-controlled notifications in mobile instant messaging. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCWI):1–25, 2020.
- [6] H.-M. Chou, D. Zhang, L. Zhou, and Y. Kang. Caem: Context-aware call management for mobile phones. In *2017 IEEE 3rd International Conference on Collaboration and Internet Computing (CIC)*, pages 399–405, 2017.
- [7] E. S. De Guzman, M. Sharmin, and B. P. Bailey. Should i call now? understanding what context is considered when deciding whether to initiate remote communication via mobile devices. In *Proceedings of Graphics Interface 2007*, pages 143–150, 2007.
- [8] D. Doukhan, J. Carrive, F. Vallet, A. Larcher, and S. Meignier. An open-source speaker gender detection framework for monitoring gender equality. In *Acoustics Speech and Signal Processing (ICASSP), 2018 IEEE International Conference on*, pages 5214–5218, 2018.
- [9] J. Fogarty, J. Lai, and J. Christensen. Presence versus availability: the design and evaluation of a context-aware communication client. *International Journal of Human-Computer Studies*, 61(3):299–317, 2004.
- [10] B. Gaver. Provocative Awareness. *Computer Supported Cooperative Work*, 11(3-4):475–493, 2002.
- [11] F. Gu, M. H. Chung, M. Chignell, S. Valaee, B. Zhou, and X. Liu. A Survey on Deep Learning for Human Activity Recognition. *ACM Computing Surveys*, 54(8):1–34, oct 2021.
- [12] M. Hassenzahl, S. Heidecker, K. Eckoldt, S. Diefenbach, and U. Hillmann. All You Need is Love: Current Strategies of Mediating Intimate Relationships through Technology. *ACM Transactions on Computer-Human Interaction*, 19(4):1–19, December 2012.
- [13] S. Hudson, J. Fogarty, C. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. Lee, and J. Yang. Predicting human interruptibility with sensors: a Wizard of Oz feasibility study. In *SIGCHI Conference on Human Factors in Computing Systems (CHI 2003)*, page 264, Ft. Lauderdale, Florida, USA, 2003. Association for Computing Machinery (ACM).
- [14] O. Inbar, G. Joost, F. Hemmert, T. Porat, and N. Tractinsky. Tactful calling: investigating asymmetric social dilemmas in mobile communications. *Behaviour & Information Technology*, 33(12):1317–1332, 2014.
- [15] M. Muhlenbrock, O. Brdiczka, D. Snowdon, and J.-L. Meunier. Learning to detect user activity and availability from a variety of sensor data. In *Second IEEE Annual Conference on Pervasive Computing and Communications, 2004. Proceedings of the*, pages 13–22, 2004.
- [16] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [17] M. Pielot. Large-scale evaluation of call-availability prediction. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 933–937. Association for Computing Machinery, 2014.
- [18] M. Pielot, R. De Oliveira, H. Kwak, and N. Oliver. Didn't you see my message? predicting attentiveness to mobile instant messages. In *Proceedings of the 32nd ACM Conference on Human Factors in Computing Systems (CHI '14)*, pages 3319–3328, 2014.
- [19] A. Salovaara, A. Lindqvist, T. Hasu, and J. Häkkinä. The phone rings but the user doesn't answer: Unavailability in mobile communication. In *13th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI 2011)*, pages 503–512, Stockholm, Sweden, 2011. ACM.
- [20] K. Tee, A. B. Brush, and K. M. Inkpen. Exploring communication and sharing between extended families. *International Journal of Human-Computer Studies*, 67(2):128–138, February 2009.
- [21] Z. Zivkovic and F. Van Der Heijden. Efficient adaptive density estimation per image pixel for the task of background subtraction. *Pattern recognition letters*, 27(7):773–780, 2006.