Smart Heating: Energy Savings Through Occupancy Sensing and Prediction

Marc Hüppin
Seminar Report
Ubiquitous Computing Seminar
ETH Zürich
Aeschstrasse 28b
5430 Wettingen, CH
hueppinm@ethz.ch

ABSTRACT
This paper will give an overview over the topic of smart heating, an important topic in saving energy. Aspects it will cover specifically are occupancy sensing and prediction, there will be descriptions and reviews of approaches of several papers in the field. It will highlight different models for thermal comfort and give a critical analysis of their quality. It will then give an overview over two current commercial products implementing smart heating approaches. Finally there is an analysis on future commercial potential, the current state of the field and what future research must look into in order to further advance smart heating and realize its potential.

General terms: Heating, Ubiquitous Computing

Keywords: heating control, home automation

INTRODUCTION
Energy saving is one of the biggest topics in the world today. Not only from a monetary and economical view but from an ecological view as well. Asides from using different, renewable energy sources and build devices that use less energy, there is the idea of saving energy by simply being more thoughtful of when and how we use them. For instance shutting off a TV when nobody is watching and not only setting it to stand by mode, but pulling the plug as well.
A lot of the energy is consumed by residential homes (21% of total energy usage in US) and heating and cooling homes makes up around 46% of the total consumption by these houses [1]. Saving money in heating can of course be done in many ways like improving insulation and installing heating systems that are more energy-efficient. The approaches discussed in this paper are trying to heat in a smarter way by only heating the house when necessary e.g. when the home is occupied. This problem can be viewed as an optimization problem. Maximize the comfort of residents with minimal heating costs.

Current Status and Limitations
Approaches already deployed in most houses today are the thermostats. A thermostat is a feedback-driven device. Given a setpoint it starts or shuts off heating and/or cooling devices to keep the temperature at the desired setpoint. The thermostats can be divided into two groups, the manual thermostat and the programmable thermostat.

Manual Thermostat With a manual thermostat one can simply choose a current setpoint. In order to save heating cost one would have to use a setback temperature when leaving the house and set it back to the desired temperature upon returning. This leads to two major drawbacks. Firstly one has to adjust the thermostat every time upon arrival and when leaving a home and secondly it will take time for the home to heat up again once the resident is home and has adjusted the setpoint. During this time there will most probably be a certain discomfort for the resident.

Programmable Thermostat The second option is the programmable thermostat, with this device the user has the possibility to enter a schedule. In this way one can ensure the home to be already warm upon arrival. This method though more advanced has its limitations as well. It relies on the user predicting its own occupancy patterns, when these change there is a need for reprogramming the schedule which is can be quite cumbersome with a lot of devices.
Numbers from [1] suggest that 55.06% of people in the U.S. do not use setbacks even when their given the possibilities to do so. This means that the issue is not only the potential in energy savings of the current solutions but also the usability. The approaches mentioned and discussed in this paper try to tackle these issues by automatically controlling the heating by sensing when a home is occupied and predicting when it will be in the future.

OCCUPANCY SENSING
One goal for implementing smart heating is to detect if a home is occupied or not. There are a lot of devices that can be used to decide whether occupancy currently occurs or not.

Motion detection

A first solution is using devices that detect motion. Below you can see a list with current technologies for occupancy sensing discussed in [2].

- Passive infrared occupancy sensors (PIR)
- Ultrasonic occupancy Sensors
- Audible sound/passive acoustic sensors
- Microwave Sensors
- Light barriers
- Video Cameras

These devices detect motion from which they decide whether a room or a building is currently occupied. Systems like the PIR, which need line-of-sight to a moving object to detect it, or acoustic sensors are prone to false-offs (declaring a room as unoccupied when someone is there) while other devices like the ultrasonic sensors can be easily set off accidentally like by moving outside of the window and therefore can produce false-ons (declaring a room as occupied when it is empty).

An approach often considered to improve performance is combining two or more of these systems. The authors of [5] use PIR sensors and magnetic reed switches to detect door openings and closing.
To further improve the performance, the authors built a Hidden Markov Model. The states are 'Active', 'Sleep' and 'Away', observable variables contain data from the sensors as well as the time of the day. Such a model needs to be trained to achieve better accuracy. The results from their Hidden Markov Model can be seen in Figure 1. Their algorithm was tested against a reactive algorithm. The reactive algorithm declares a home occupied whenever a sensor firing occurs, and changes to unoccupied after a certain number of minutes which was set at different numbers ranging from 5 to 120. The results show that their HMM has the highest accuracy at about 88%, but the wrongly classified 12% are about equally distributed as active when inactive and inactive as active, while the reactive algorithms errors are almost completely active as inactive. This means that while the house may be heated unnecessarily using the reactive algorithm, the house will always be warm when someone is present.

In range

Another way to sense occupancy is having residents carry devices that send out signals, based on these signals we can then decide whether or not a home is occupied.
One approach taken by the authors of [3] is placing active RFID tags on the residents keys. These tags send out a signals to a receiver to declare their presence. The reach of the signals is about 8 meters so they will only be detected by the receiver (installed at home) when the residents are at home as well. To make sure this approach works properly residents have to leave their keys in the entrance hall (close to the receiver) whenever they are home and take them with them whenever they leave the house.
GPS can be used to detect occupancy as well. In [4] the authors used GPS loggers that have to be carried by the users to monitor their location. Whenever the logged coordinates are within 100 meters of their home it is considered occupied.

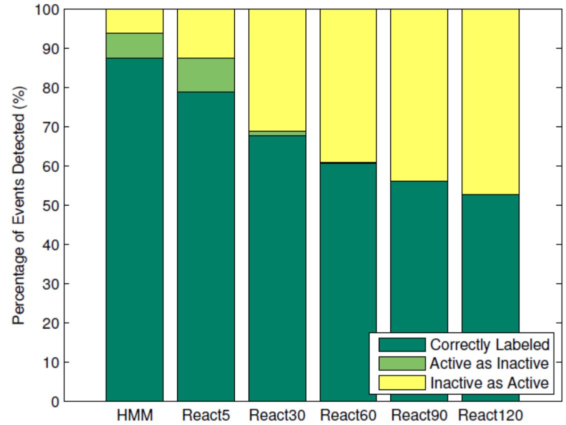

Figure 1: Results of occupancy sensing from [5]

(a) Motion Sensor (b) Door Sensor

Figure 2: Picture of the sensors taken from [5]

This can lead to false results, for instance when somebody is spending time at their neighbors. But most often a person would not be within 100 meters of their residence for longer periods of time without going home. Nowadays more and more phones have GPS capabilities as well, making them usable for these purposes too. This is an approach taken in [1].

Discussion

Having users carry around devices makes these system vulnerable to human errors, e.g. forgetting your GPS Logger or keys or keeping them in your pocket instead of the entrance hall. But in contrast to the motion sensors they are not vulnerable to other factors like pets or moving plants and are generally more reliable. A drawback from the in range systems is that they only work well on a big granularity or per house level. Motion sensors can be used to decide occupancy on a per room level as well. But besides being more error-prone users might not like having their houses equipped with motion sensors. As they may feel interference with their privacy or simply view them as an aesthetic issue. Having devices like these in Figure 2 all over a house may be irritating to some.

The way to go here in my opinion is using GPS and have phones act as loggers. The obvious advantage of this approach is that most people own a phone with GPS capabilities eliminating the need to buy new devices. A lot of people rely on their phones as well and therefore carry them all the time, which leads to less human error. Of course other systems need to exist and be explored as well. Especially for families where children may not have a phone with GPS but also for other purposes. When the heating system allows to control the temperature per room the sensing can be expanded by using motion sensors to detect occupancy per room.

OCCUPANCY PREDICTION

In order to maximize the comfort for residents it would be very helpful to be able to predict occupancy for future times. Occupancy prediction makes it possible to heat up homes to a desired temperature upon the return of residents. While schedules vary from one person to the next most of them have a certain regularity to them. The approaches described in this section try to exploit this regularity in order to make an accurate prediction whether a residency will be occupied at a certain time in the future.

Neural Network

A neural network is a computational model which is capable of learning and pattern recognition [7]. The authors from [6] proposed to use such a network to compute probabilities of occupancy for a future time. The computational model takes several inputs including the current time of the day and day of the week. Further it takes as input occupancy data from the three past days as well as the same weekdays from the past four weeks. Finally it considers occupancy in the past few hours. It then processes these inputs to derive a probability.

This system takes into account daily as well as weekly patterns. The choice to look at occupancy in the past three days seems a little bit curious. Occupancy patterns are often occurring over longer timespans and an untypical day in the past three days might effect the prediction in a negative way. Another drawback of this approach is that the suggested training takes 150 days, during which humans must themselves keep track of their occupancy in order to use the back propagation training and compare prediction to the actual events. A task which can not be asked of a typical user.

Linear Algebra

Other approaches use vectors having each vector entry representing a timeslot. The authors from [3] have a vector for each day and the values in the vector entries are either 0 for away or 1 for home. In order to predict occupancy the vector of the current day (up to the current time) is compared to the vectors of previous days. Then the K most similar days are considered. Similarity is decided using the the Hamming distance, which counts the number of elements in which the vectors differ. The probability of occupancy for a future timeslot is then calculated using the mean of the corresponding values in these K vectors. The authors of the paper found 5 to be a good choice for K. They do not give any further insight to this decision and for different circumstances different numbers might work better. To calculate probabilities at the beginning and at the end of the day these vectors are padded with four hours at either side. For further improvements of results they only consider weekdays for weekdays and weekend days for Saturdays and Sundays.

As can be seen in Figure 3 the results from this approach are satisfying. While the calculations take daily patterns heavily into account, it does not regard any weekly patterns. Meaning past Tuesdays do not have more effect on prediction occupancy on a Tuesday than past Wednesdays. Taking this into account might lead to even better results.

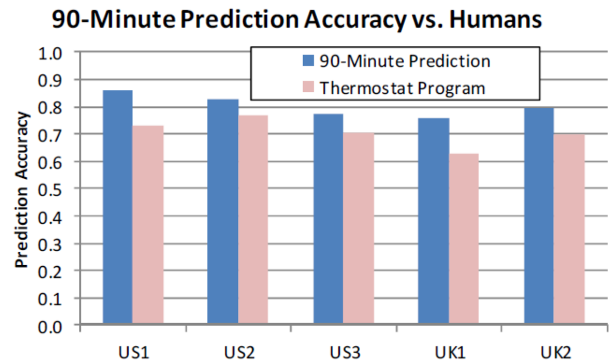

Figure 3: Results of the prediction algorithm used in [3]

A paper which does take weekly patterns into account is [4]. Their approach uses two vectors. One vector represents the whole week and contains an entry for every half hour of the week resulting in a length of 336. Each entry contains the probability of the user being away during this specific timeslot. Making the assumption that weekdays resemble each other they use a second vector that contains the probabilities for being away for a typical weekday. To calculate the actual probabilities for timeslots during weekends it relies solely on the week-long vector, while predictions during weekdays are influenced by the typical weekday vector as well. This approach takes into account weekly patterns as well as daily patterns in the form of the typical weekday. The approach was tested with 34 subjects. All of them had to fill out a schedule of being home or away for a week as well. Their evaluation shows that their proposed prediction algorithm gets better results than the residents predicting their own schedule, which would be used in a programmable thermostat. This is of course a relatively small sample size.

Drive Time

The advantage of using GPS sensors is that in contrast to other occupancy detection system it is not simply binary (home/away), but allows for detection of the position of residents outside their homes as well. The authors of [1] try to take this knowledge to improve heating systems. When receiving data of a location from resident it calculates the drive time (the time it would take the resident to get home) using MapQuest. The system then predicts the user will be home in this calculated time. If the resident is then for instance 30 minutes away, we can set the temperature at a low setpoint but high enough to have it at a desired temperature in 30 minutes.

This approach will obviously not produce accurate prediction results especially for users with small commute times. With increasing commute times the benfit of this approach will increase as well since this allows for lower setbacks. An advantage it has over the other approaches is, that it would not need any training. This approach could be improved using more data like current traffic congestions to get more accurate drive times. Another improvement could be to include other means of transport into the calculation. A user could tell the system whether he goes away by car or using public transportation.

Combination To improve prediction results the authors of [4] combined their solution with the idea from [1]. By setting the probability of being home to 0 for a future timeslot when the timeslot is closer to the current time than it would take the resident to drive home. This further improved their results.

	[1]	[4]	[3]	[6]
Time of the Day		✓	✓	✓
Occupancy on current day			✓	✓
Specific Weekday		✓		✓
Weekday/Weekend		✓	✓	
Similarity in occupancy between days			✓	
Drive-Home-Time	✓			

Table 1: Input factors for occupancy prediction in [1] Drive Time Algorithm, [4] Probabilistic Schedule, [3] PreHeat, [6] Neural Network

Discussion

The several models and approaches discussed in this section all predict occupancy based on different inputs. An overview over these factors can be seen in Table 1. The models proposed in [3] and [4] both show decent results. They focus on different kinds of regularity. While the method from [3] focuses on similarities between any two days, the model from [4] calculates probabilities for future occupancy on weekly patterns their approach is built on the belief that people follow a relatively stable weekly schedule.
Can we classify one as better than the other? I think this depends on the users. For residents with a relatively fixed weekly schedule the approach from [4] will probably deliver better results, but considering someone working a job with different shifts the calculation model from [3] could lead to a more accurate prediction.

Both of these models need training in order to deliver satisfying results. Would it be possible to use pre-trained systems? The idea here would be to do training for different demographic groups (i.e. students, families, retired people). I think the possibilities for further improvement in occupancy prediction lie in exploring different lifestyles and find models that match the best to a certain lifestyle. Also one has to look into combining models, or ideas of models for the best possible results.

CONTROLLING THE HEATING

When we gain knowledge of current and future occupancy the task is now to apply this knowledge to heat in a smarter way. The general idea is to not heat the house (using a setback) when no occupancy is sensed and have the house at a desired setpoint when occupancy is expected. A lot of the occupancy prediction models have occupancy probability as an output. We can now set a threshold on these probabilities. Meaning at which probability do we want to have a heated house. Here we have a trade off between heating a home unnecessarily and returning to a cold home at times. This threshold should be set by the user complying to the users preferences (e.g. save more money/never be cold).

COST MODELS

To evaluate these systems, we need to have cost models. The first cost we need to measure is heating cost. This can be measured in dollars and is relying on the type of heating system, the insulation of the home as well as the current cost of the fuel. Another option is to simply calculate the amount of energy used for heating.
The other cost we need to measure is the one of comfort. There has been rich research on thermal comfort of human beings wikipedia gives a good overview [10].

Standards

One of the standards for measuring or calculating human comfort at certain temperatures is ASHRAE (formerly American Society of Heating, Refrigerating and Air Conditioning Engineers) standard 55 [8].

This calculation model takes into account several factors to decide whether a human is comfortable or not:

- Air temperature
- Mean radiant temperature
- Air speed
- Humidity
- Metabolic rate
- Clothing level

Unfortunately it is not always possible to have information on all these values. Therefore the discussed papers used different approaches to describe the comfort or discomfort of the residents in their tested homes.

Discomfort in Minutes

The so called MissTime Metric was introduced in [5] and used in [3] as well. It expresses the discomfort as the amount of time a home was occupied while the temperature was not within 1C of the desired temperature per day.
This metric is easy to calculate but it does not take into account how far the actual temperature is off the setpoint.

Expressing Discomfort in Money

The authors from [6] tried to measure discomfort from an economical standpoint. They try to measure the loss in productivity when working in a room that is not heated to a desired temperature. To convert that measurement into dollars they use an hourly salary. They take into account the size of the difference between desired and actual temperature.
While the idea is certainly interesting and allows for a direct comparison between heating and comfort cost. The measurement does have some arbitrary factors in it. Do we take the hourly salary of the resident? Or average hourly salary in a certain country? While it is an interesting approach I think it is close to impossible to come up with a calculation that describes thermal discomfort in dollars in a way that is right for every circumstance. In general comparing heating cost directly to comfort is very difficult as it is a very subjective matter.

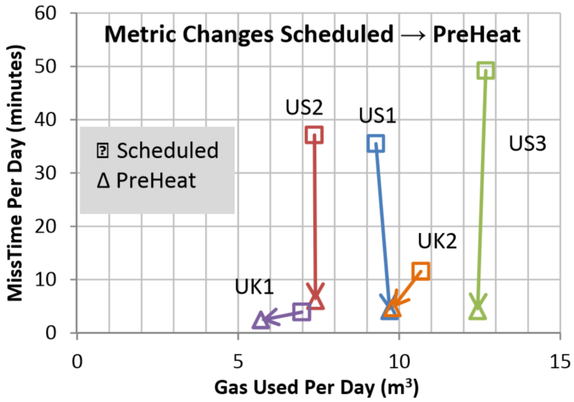

Figure 4: Graph from [3] displaying their results

Discussion

The task of finding a good model for measuring comfort is not easy. Most of the papers use simplified measurements to evaluate their results. The latter two metrics leave a lot of things out of consideration for instance the thermal discomfort is assumed to be the same for everybody at any activity level, clothing level and so on. I feel by using advance thermal comfort models like the standard 55 from ASHRAE more possibilities could arise. Currently the ideal temperature is chosen by the user and is (as long as it is not adjusted) static. Including thermal comfort calculation models the desired temperature could be further adjusted taking more external factors into account.

RESULTS

We will briefly look at the results in saving heating cost and/or improving comfort of residents that the authors from [3] and [5] displayed in their papers. Both solutions were deployed in a small number of houses (5 and 8 respectively).
Figure 4 displays the achieved results from [3]. They compare their solution to a scheduled heating scheme. In the two houses in the UK (working on a per room level) they achieve both energy savings and improvement in comfort. In the three US houses (per house level) they achieved similar energy consumption while improving comfort. It is obviously hard to make general assumptions from such a small sample size and should not be done. Also the comparison to a scheduled algorithm is not optimal, since it highly depends on the quality of the programmed schedule. They claim for the schedules to be carefully programmed but what does that really mean? Is there no possibility of programming a better schedule? Also the authors tested their system in their own houses, making it impossible to evaluate usability. To really be able to make a hypothesis on the quality of the product we need results from a lot more households.
The number of households that tested the system described in [5] is also too small to make general claims. Their results show lower energy usage than both a reactive algorithm and the always-on approach. The comfort measured in the MissTime Metric was higher than the one achieved by the reactive algorithm on average but not in every tested household. This highlights that the same approach might not be the right one for everybody. The authors also present a back-of-the-envelope calculation on nationwide energy savings in the US if their system was deployed in all homes. Calculations of this type are making too many assumptions for the number to have any real value.
What we can take from the results from both papers is that both of the approaches can work and can achieve great results. But it should not lead us to believe that these solutions are ready to be deployed. There needs to be a lot of further testing to allow us making such assumptions.

Figure 5: The nest thermostat (source: [14])

COMMERCIALLY AVAILABLE SOLUTIONS

There are already commercial solutions on the market that try to bring smart heating to residential homes. I would like to further look into two companies.
Nest [11], currently only selling their products in North America, was recently acquired by Google for $3.2 billion [12]. Their major product is the nest thermostat which is currently selling for $249 see Figure 5. The device has a simple interface that allows the user to change the setpoint. The idea is for the user to manually change the temperature, like with a manual thermostat. The nest thermostat will remember the settings. After 12 days the so-called Auto-Schedule should be ready and adjust the temperature automatically.
Another feature called "Auto-Away" uses occupancy sensing. The thermostat contains motion sensors, during the first few days it tries to learn when people are home or away. It uses the learned occupancy pattern and the motion sensors to detect when nobody is home. This is a similar idea to the one presented in [5].
The availability of a smart phone app allows users to control their thermostat when their away.
The approach from nest can be seen as a hybrid solution. It works automatically while still allowing the user control. In contrast to the programmable thermostat the user interaction is very intuitive and simple. Since the features also rely on weather data which it currently can only access for US and

Figure 6: Screenshot of the App from tado [13]

Canada locations it is not yet deployable with the same results in other countries.

Another solution which was and is developed in Germany is tado [13], aside from Germany it is also available in Switzerland, Austria and the UK. The current connector kit ships at 299 Euros, it consists of a box that is connected to the boiler (or replaces an existing thermostat), a gateway installed at the router and a temperature sensor. The approach their taking is similar to the one from [1]. The functionality is based on a smartphone app (See Figure 6), this app communicates with the gateway installed at home to transmit the current location from which the distance to the home can be calculated. The temperature is then adjusted according to the current distance. This means the house cools down the further a resident is away and starts heating up when he approximates his home again. This way the house will always be at the desired temperature when one returns.

Comparing the two solutions, we can see that they take different approaches to reach the same goal. Nest uses machine learning to derive an occupancy pattern and heat accordingly. It learns the human preferences in order to ensure comfort. An advantage that tado has over nest is that it does not require any user interaction to work. Nest though has a very short training time (12 days are advertised) and the interaction with the device is as stated before very simple. This is a domain where tado has yet to improve. Currently the physical interface of tado is limited to a button that one can press in order to declare occupancy. There is no way of setting new temperatures without the app or the web interface, which can be cumbersome for some people.
At the moment the two products are not competing directly since their respective markets do not overlap. Comparisons are of course still important since it wont be too long when they do.

COMMERCIAL POTENTIAL

Like stated in the introduction the potential energy savings with improved heating techniques are big. The numbers from the research papers ([3], [5] and [1]) have to be looked at critically, since they only include a very small sample size. Deployments and numbers from nest and tado are also suggesting heating cost savings up to 31%. But not all customers are happy with these devices, looking through the review section on amazon.com the main feelings and reviews are rather positive (average score is 4 stars) but there are also a good amount of customers that seem to be unhappy with the product.
This leads to conclusions that the commercial devices are not right for everybody at least not in the current version. But even a lot of the users giving bad reviews believe in the potential. The fact that Google paid $3.2 billion also shows the trust they have in the commercial potential of nest and smart heating control. With the purchase of nest by Google, the possibilities for nest have also greatly improved due to much more resources. I think at the moment and also in the near future it will be important to keep allowing users to manually interact with the thermostat. This is due to the fact that none of the automatically ”smart” algorithms are guaranteed to work in every setting and for every situation. Users need to have to possibility to manually and easily override decisions taken by the algorithms.

CONCLUSION

Heating more efficient is and will be an important topic as it is an area where a huge amount of energy can be saved. While greater energy efficiency can be achieved in many ways, scheduling the heating devices in a smart way is definitely an important one. We have seen that there are several approaches in how to achieve these smarter schedules, also there are already commercial solutions which implement some of the ideas proposed in the reviewed research papers. These commercial solutions are currently not perfect but can achieve great savings under the right circumstances. Future research has to also focus on these different circumstances, lifestyles and habits in order to be able to create products that work for everybody and every setting or at least be able to define the ones their solutions work in. While the automatic solutions might currently not always achieve greater results than programmed thermostats, the fact that a big part of the people owning programmable thermostats do not use them, makes it clear that there is a necessity for smart heating and even simple algorithms like the one proposed in [1] can deliver better heating patterns. The relatively low correct usage of programmable thermostats also highlight another important fact, if a solution wants to be commercially successful it needs great usability so that people will actually make use of it and explore the potential.

ACKNOWLEDGMENTS

I would like to thank Wilhelm Kleiminger for his support for this paper. Also I want to thank everybody who participated in the lively discussion after my presentation.

REFERENCES

1. Gupta, M., Intille, S.S., Larson, K., Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges.

2. Guo, X., Tiller, DK, Henze, GP, Waters, CE, The performance of occupancy-based lighting control systems: A review. August 6 2010

3. Scott, J., Bernheim Brush, A.J., Krumm, J., Meyers, B., Hazas, M., Hodges, S., Villar, N., PreHeat: Controlling Home Heating Using Occupancy Prediction In *UbiComp'11* (Spetember 17-21, 2011, Bejing, China)

4. Krumm, J., Bernheim Brush, A.J., Learning Time-Based Presence Probabilities In *Microsoft Research*

5. Lu, J. et al., The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes In *SenSys '10* (November 3-5, 2010, Zurich, Switzerland)

6. Mozer, M.C., Vidmar, L., Dodier, R.H., The Neurothermostat: Predictive Optimal Control of Residential Heating Systems In *Adv. in Neural Info. Proc. Systems 9 (pp. 953-959)* (MIT Press, 1997, Cambridge, MA)

7. Artificial Neural Network `http://en.wikipedia.org/wiki/Artificial_neural_network`

8. ASHRAE `http://en.wikipedia.org/wiki/American_Society_of_Heating,_Refrigerating_and_Air-Conditioning_Engineers`

9. ASHRAE Standard 55 Calculation Tool `http://smap.cbe.berkeley.edu/comforttool`

10. Thermal Comfort `http://en.wikipedia.org/wiki/Thermal_comfort`

11. Nest Homepage `https://nest.com`

12. Google Press Release `https://investor.google.com/releases/2014/0113.html`

13. tado homepage `http://www.tado.com/`

14. Amazon Images `http://ecx.images-amazon.com/`